



Streaming Integration and Intelligence

The Striim Platform Overview White Paper



November 2017

Streaming Integration and Intelligence

Table of Contents

- Executive Summary** 1
- Striim’s Core Platform Capabilities** 4
 - Continuous Collection of Structured, Semi-structured and Unstructured Data 4
 - In-Stream Data Transformation and Enrichment..... 5
 - Streaming Analytics 5
 - Alerting and Visualization 5
 - Real-Time Delivery 6
- Technology Solutions and Business Use Cases**..... 6
 - 1) Real-Time Data Integration 6
 - 2) Detecting Patterns and Anomalies 8
 - 3) Creating and Monitoring Real-Time Metrics 11
 - 4) IoT Edge Processing and Edge Analytics..... 13
 - Key Value Drivers 15
- Why Striim?** 16
- Conclusion** 19



Executive Summary

The exponential innovation in technology has forced organizations to re-think their way of running business operations and adapting to new market demands. Only the businesses that reform their processes and policies, and turn data into asset using next-generation data infrastructure are the ones that not only survive but thrive in today's fast-paced, innovative environment.

Similarly, the exponential data growth has forced IT organizations to re-think their data architecture. When data is generated in massive volumes and with high-velocity, to gain its full value businesses must accelerate how fast they capture, analyze, and act on this continuously streaming Fast Data. Combined with the competitive and market pressures, the data deluge is forcing businesses to modernize their data infrastructure by adding in streaming data architectures that provide them in-time, in-context insights for better operational decision making.

Striim is a complete, enterprise-grade software platform for real-time data integration and streaming analytics that makes working with modern, streaming data architectures fast and easy. It enables businesses to gain the maximum value out of their high-volume, high-velocity data by collecting, processing, and analyzing it in-flight, before feeding it in real time to the rest of the enterprise.

Striim can ingest data from virtually any data source including transactional data from enterprise databases — via non-intrusive change data capture (CDC) — log files, events from message queues, and IoT sensor data. While it's streaming, Striim immediately processes the data, filtering out what's irrelevant, aggregating and transforming the data, and enriching it with historical or reference data for full context. Striim also offers streaming analytics with multi-source correlations, outlier detection, pattern matching, statistical processing, predictive analytics, and other advanced analytics. With the data analyzed, Striim can issue real-time alerts, trigger workflows, provide push-based visualizations, and store processed data in memory for fast access via ad-hoc SQL queries.

Striim customers use the platform to achieve their key business and IT objectives, ranging from modernizing their data and analytics infrastructures to ensuring regulatory compliance, to delivering superior customer experience, and to optimizing the use of their scarce resources. The platform supports these critical initiatives with solutions including real-time data integration across enterprise systems and the cloud, creating and monitoring operational metrics with live dashboards, detecting patterns and anomalies, such as for preventing security breaches and fraud or offering location-based marketing programs, and more. By delivering a wide variety of innovative use cases, Striim helps businesses minimize operational risks and seize time-sensitive opportunities to grow profitably and surpass the competition.



Striim enables businesses to minimize operational risks and seize time-sensitive opportunities to grow profitably and surpass the competition.

Striim is designed from the ground up to meet the unique requirements of streaming data management in mission-critical, high-volume, and high-velocity data environments.

Striim is unique among other real-time integration and analytics solutions by offering an end-to-end solution that combines continuous unstructured, semi-structured, and structured data collection, with in-flight data processing and analytics, and data visualization in a single, enterprise-grade platform. Striim is designed from the ground up to meet the unique requirements of streaming data management in mission-critical, high-volume, and high-velocity data environments. It easily integrates, analyzes, and visualizes different sources and types of data — on-premises, at the edge, or in the cloud — with reliability, security, and scalability. By integrating and analyzing data across cloud, Big Data, and IoT devices, Striim helps companies manage their data across heterogeneous legacy and next-gen infrastructures without getting locked into a single topology.

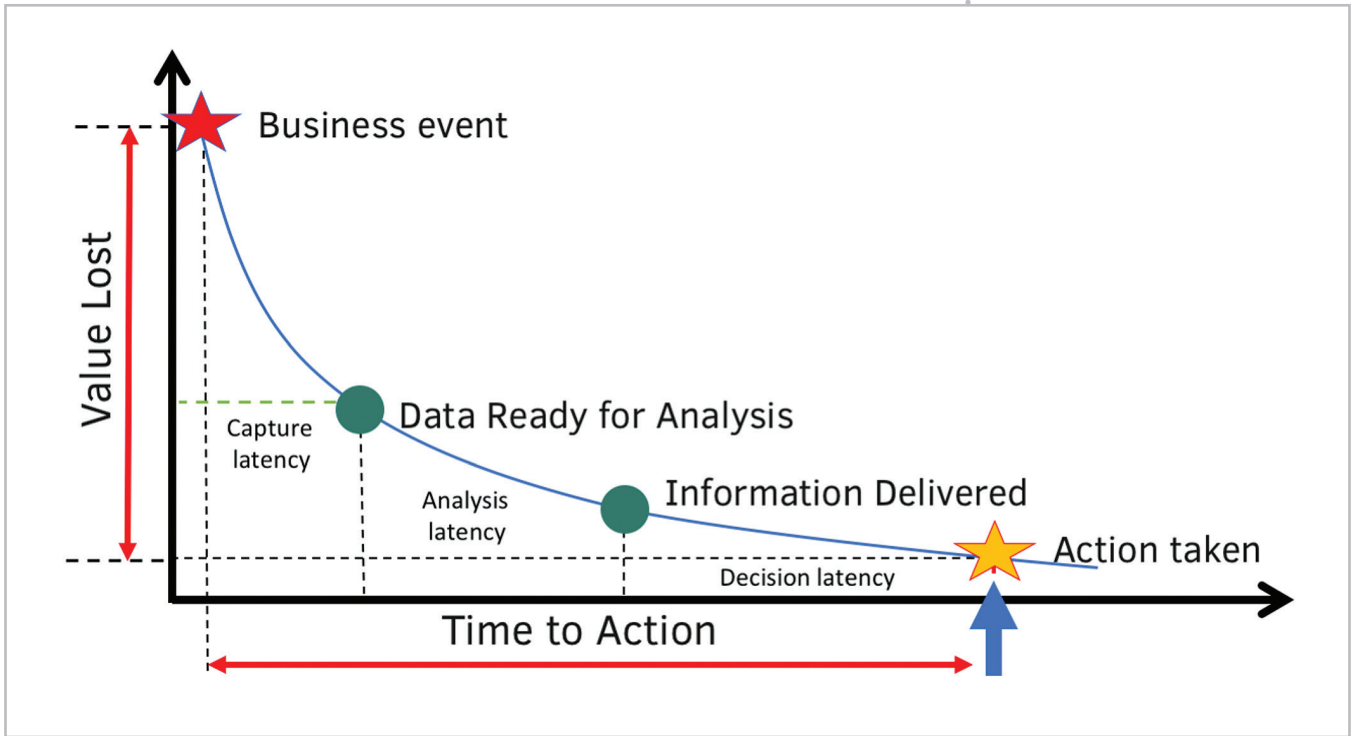
Application development using the Striim platform is SQL-based and GUI-driven, making it easy to build and modify integration and analytics applications. In addition to empowering business users to modify their analytical applications running on Striim, it effortlessly integrates with the existing data infrastructure to augment batch processing solutions with streaming data integration and analytics. With Striim, organizations can gradually expand their stream processing use cases as the business demands, without having to rip and replace previous IT investments. Striim makes the transition to a modern data infrastructure smooth and easy and enables businesses to thrive in fast-paced environments with in-time, in-context insights.

Introduction: Why Streaming Data Integration and Analytics?

The digital world moves fast with a machine-based speed. Disruptive technologies and innovations, dynamically changing market demand and supply, macro-economic and political environments, and competitive pressures require companies keep pace to control “the Now” in order to prevail. The only way to better control the Now in the business is to make the moment-to-moment operational decisions based on the best and most current insights available. Organizations can gain better understanding and control of their operations when they analyze *all* available data the instant it’s born, and act on urgent threats or opportunities immediately.

After-the-fact analytics with Big Data does not enable enterprises to control the Now. In real life, data is business-event driven and generated continuously as streams, not in batch. Because of the high-velocity nature of today’s data and the speed at which businesses operate, as time passes, “Fast Data” becomes less and less relevant to current reality, and less able to impact business operations. Unless analyzed immediately, most value from Fast Data diminishes. In some cases, such as with location data, its value is completely perishable unless acted on within a short time frame, typically less than a minute.

Richard Hackathorn depicts in his Value-Time Curve (shown below) the lost value for the enterprise due to latency in data capture, analysis, and action. To improve operational decision making and act on time-sensitive opportunities or threats, businesses need to acquire, analyze and act on data as close as possible to the originating time and place, while insights remain relevant and easily actionable.



The Value-Time Curve by Richard Hackathorn

Another key reason for moving toward streaming data integration and analytics is the need to build a smart architecture for the imminent data deluge. Per IDC's¹ projections, in 2025 the world will create and replicate 163ZB of data, representing a tenfold increase from the amount of data created in 2016. In the same study, IDC estimates that nearly 20% of the data generated will be critical to our lives and 10% of that will be hypercritical. It is estimated that more than a quarter of data created will be real-time, and IoT data will make up more than 95% of this segment.

Also, by 2025, the size of global data storage capacity relative to the amount of data generated is expected to decline significantly. According to IDC, "the quantity of data generation can and will continue to outpace any reasonable expectation of our ability to store all of the data. For example, it would take roughly 16 billion of today's largest 12TB enterprise HDDs to store the 163ZB data expected to be created in 2025. To put that into perspective, over the past 20 years, the disk drive industry shipped 8 billion HDDs and nearly 4ZB of capacity." Numerous predictions including IDC's point to the fact that only a small fraction of the data generated will be able to be stored.

Only by processing data the instant it's born can businesses determine what is important, and filter out the rest before it is stored.

With the exponential growth, we see in data volumes, efficient use of available storage space is more crucial than before. Since it is physically not possible to store all the data being generated, the only option is to process, filter, aggregate, and smartly extract information from the data in-memory, in a streaming fashion, before it ever lands on disk. Only by processing data the instant it's born can businesses determine what is important, and filter out the rest before it is stored. Streaming data architectures can also process data in batch, giving businesses the flexibility to determine the speed they collect and analyze their data.

Moreover, with the adoption of cloud computing, connected IoT devices and Big Data analytics, organizations now need to manage their data across different legacy and next-generation infrastructures, meeting varying requirements for data latency and speed in processing. Streaming data integration and analytics is the modern approach to managing data in today's digital age, enabling better resource utilization while providing a future-ready environment to meet ever-changing business needs.

Striim's Core Platform Capabilities

Striim is a complete, enterprise-grade software platform for real-time data integration and streaming analytics. It enables businesses to make smart decisions with full context by integrating and analyzing all of their data in-flight, while the data is still relevant. With the following core capabilities, Striim enables businesses to respond to key business events that need immediate attention intelligently so they can navigate the challenging roads with ease, at full speed.

CONTINUOUS COLLECTION OF STRUCTURED, SEMI-STRUCTURED AND UNSTRUCTURED DATA

Striim captures real-time data from a wide variety of sources including databases, log files, IoT devices, message queues, and does so for different data types including JSON, XML, delimited, binary, free text, change records. For data ingestion, the platform uses various mechanisms including Kafka, JMS, MQTT, OPC UA, UDP, TCP, HTTP, log files, and Hadoop Distributed File System (HDFS). The platform is unique in its ability to continuously ingest data from enterprise databases, such as Oracle, SQL Server, MySQL, HPE NonStop, MariaDB via non-intrusive change data capture (CDC), as well as to aggregate and correlate that data with other data streams while in flight.

Striim reads database transaction logs to capture changed data without requiring any changes to the source environment, minimizing any impact on source systems. Striim can also capture data from Oracle GoldenGate Trail Files. With its Publish API, users can deliver data into Striim remotely through Java code. Striim also provides users with live data processing metrics across different data sources, so users have visibility into their data flows across different environments.

ETL VS. STREAMING INTEGRATION

Traditional Data Integration (ETL)	Streaming Data Integration with Striim
<p>Focus on:</p> <ul style="list-style-type: none">• Data Modeling (defined schema)• Data Profiling• Data Quality <p>Processing:</p> <ul style="list-style-type: none">• Data written to disk multiple times causing material lag• Non-native real-time approximation, complex, and expensive to maintain• Brute-force processing techniques require additional cores	<p>Focus on:</p> <ul style="list-style-type: none">• Simple In-line Enrichment• Filtering• Aggregation <p>Processing:</p> <ul style="list-style-type: none">• Data processed the instant it's born• 100% in-memory, before ever landing on disk (including temporary tables)• Native, efficient real-time stream management

IN-STREAM DATA TRANSFORMATION AND ENRICHMENT

To add context and enable easier downstream analysis, Striim applies easy-to-use filtering, transformations, aggregations, and enrichment of streaming data with reference data. Continuous data streams are fed into continuous queries running on distributed servers to transform, enrich, and analyze the data in milliseconds. Striim is specifically architected to join high-speed, streaming data with both real-time and historical information and metadata, without slowing down processing. As data streams flow through the distributed cluster, the context caches are ready to enrich and transform the data stream on each node without bottlenecking the stream and impacting performance.

STREAMING ANALYTICS

Using flexible time series and windowing, Striim performs multi-source, spatial and time-based correlations, complex conditional pattern matching with case statements, predictive analytics, outlier detection and other statistical operations. Customers can detect events of interest and gain insight from all their data before it ever lands on disk. Striim stores processed data in-memory for fast access and advanced analytics. Users can run SQL queries on both the streaming data and results data for their ad-hoc and regular reporting needs. The Striim platform can be easily extended with Java-based, imported functions.

ALERTING AND VISUALIZATION

Striim issues real-time email, text, and web alerts to notify end users of events of interest. It offers real-time, push-based visualizations via easy-to-build/-modify dashboards. In addition, Striim applications can trigger workflows to enable specific actions for any critical event.

Key Features

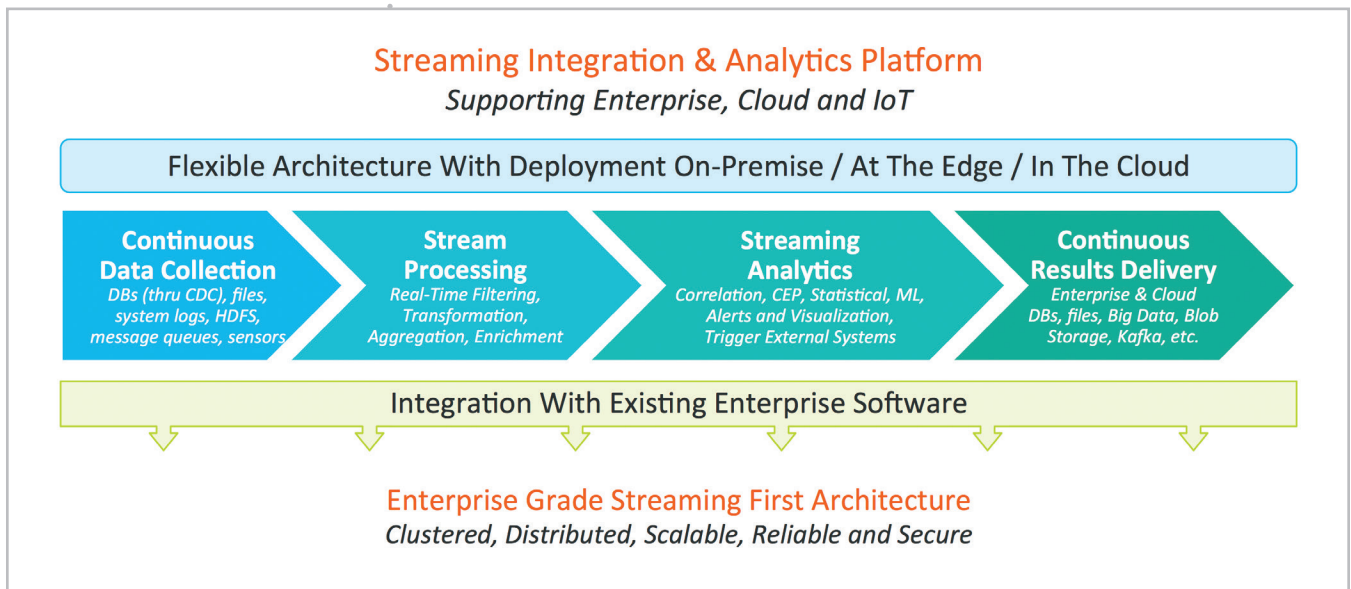
- 1. Continuous processing:**
All ingestion and processing is continuous as the data is generated.
- 2. Sub-second latency:**
Captures, processes, and delivers data and resulting insights in milliseconds.
- 3. Minimal impact:**
Log-based change data capture for relational databases minimizes impact on sources.
- 4. Multi-source correlation:**
Correlates across multiple and varied sources — including streams of enterprise data.
- 5. Flexible time windows:**
Provides time-based, event-based, and session-based windowing.
- 6. In-memory computing:**
Combines in-memory computing (IMC) technologies for High Speed Message Infrastructure, In-Memory Data Grid, Event Stream Processing and Analytics, and Data Visualization.

REAL-TIME DELIVERY

The Striim platform delivers data in real time to a wide variety of targets including:

- Hadoop environments — including HBase, Hive, HDFS — , MapR DB and MapR Streams, Cloudera, and Hortonworks, as well as NoSQL databases
- Cloud environments such as Microsoft Azure, Amazon Redshift and S3, Google BigQuery and Google PubSub
- Messaging systems such as Kafka, JMS
- Enterprise databases and data warehouses
- Files in delimited, JSON, XML, AVRO, template-based formats

By supporting a wide variety of targets, Striim ensures pre-processed data and real-time insights reach the rest of the enterprise to support existing analytics workflows and transactional systems.



Striim offers an end-to-end platform that offers continuous data collection, real-time data processing, streaming analytics, data visualization and continuous delivery with sub-second latency.

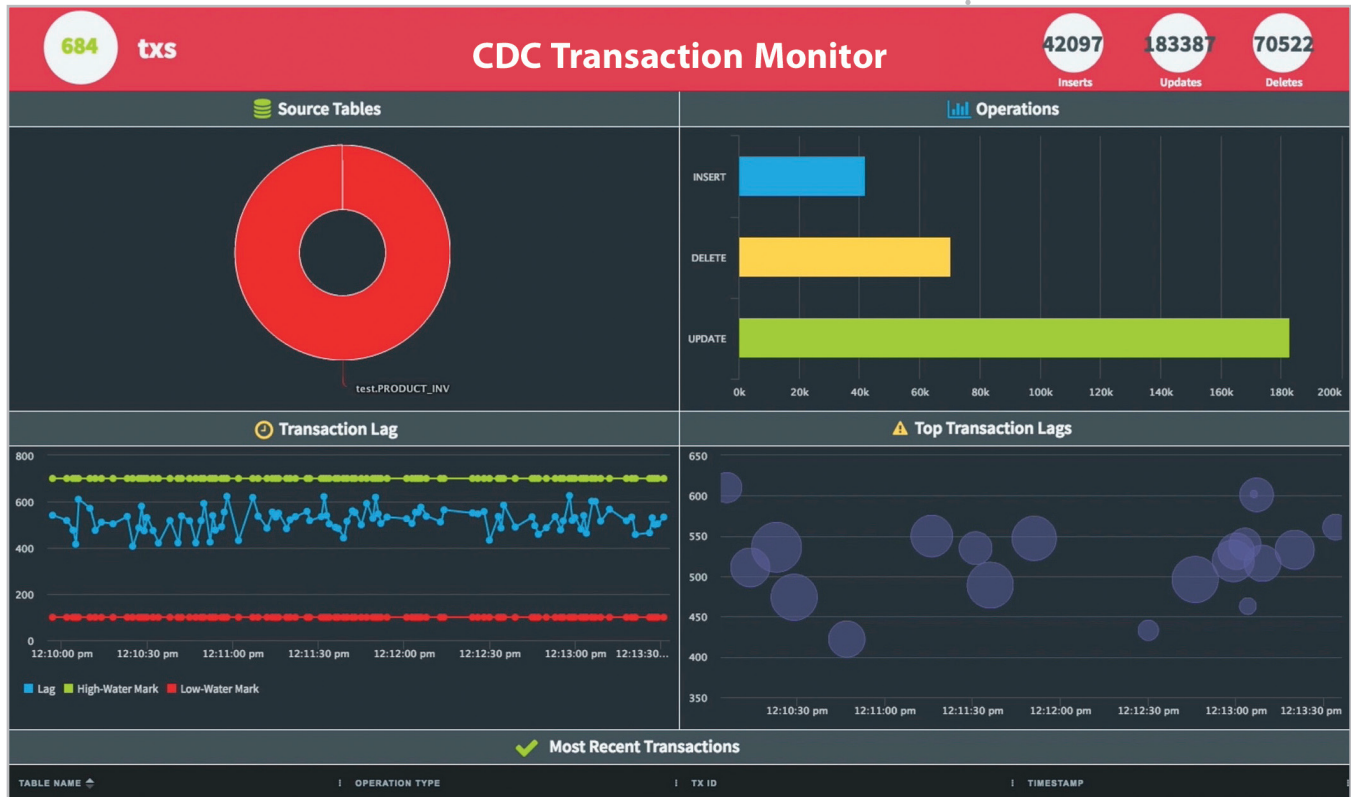
Technology Solutions & Business Use Cases

With its comprehensive streaming data integration and analytics capabilities, the Striim Platform enables numerous technology solutions that support a broad spectrum of use cases in different businesses and industries. Below are some commonly implemented technology solutions and sample business use cases.

1) REAL-TIME DATA INTEGRATION

In today's fast-paced business world with high customer expectations for speed and accuracy, access to real-time data is critical to support business operations. Striim enables real-time data movement from a wide variety of data sources to virtually any destination on premises or cloud. As described

in the Core Platform Capabilities section, Striim uses log-based change data capture (CDC) to move real-time data with minimal impact from heterogeneous databases and continuously collects real-time data from log files, sensors, and message queues. It performs filtering, aggregation, transformations, and enrichment as needed to provide the additional context for easier and better analysis.



With Striim customers can easily monitor their integration solutions in real time.

Real-time data integration can be leveraged in different topologies and solutions. Most commonly used technology use cases include:

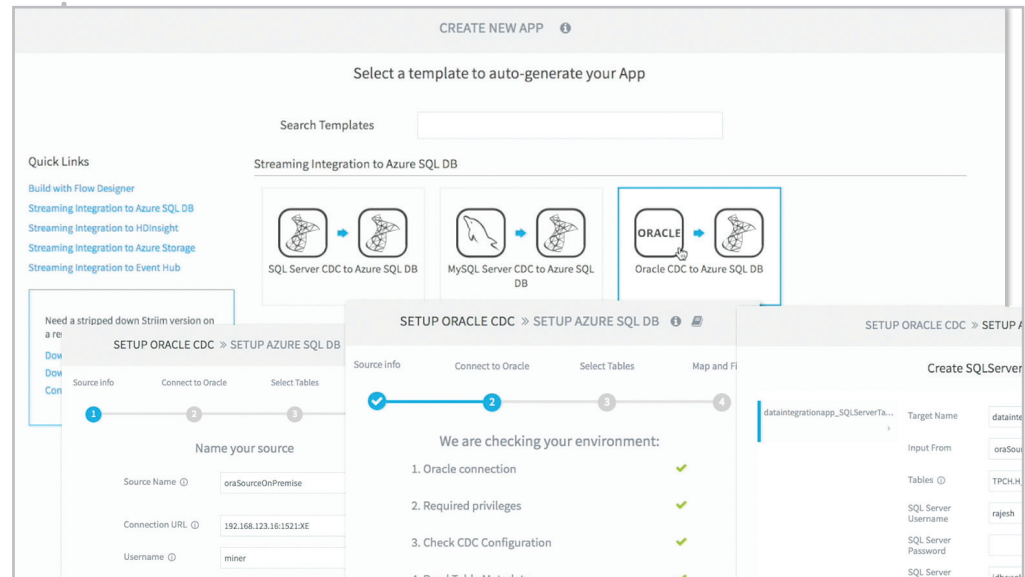
- Real-Time Data Integration for Big Data and Data Lake Solutions:** Using Striim customers deploy streaming data pipelines into Hadoop and NoSQL environments with pre-processed data. Striim offers structured, semi-structured and unstructured data ingestion and in-flight transformations, filtering, and aggregations, which significantly reduces storage footprint. Real-time data integration is also used to create data pipelines into cloud environments such as Azure, AWS, Google BigQuery. Customers also publish streaming data into message queues, such as Kafka before delivering to Big Data environments.
- Data Lake Optimization:** With Striim, users can perform aggregations, enrichment, and correlations before storing data, so OLAP queries are faster and easier to write. In-flight, pre-processing of different data sources optimizes and accelerates analytics for data lake users.

Striim enables real-time data movement from a wide variety of data sources to virtually any destination on premises or cloud.

Benefits of Integration with Striim

- Easily connect Big Data, IoT, on-premises and cloud data in real time for new services
- Simplify downstream analysis by pre-processing data for data lakes
- Extend the lifespan of your Big Data environment by filtering and aggregating data before delivery
- Increase productivity by automating data preparation tasks
- Use a single platform, and re-use streaming data pipelines, for multiple use cases
- Leverage and bolster existing open source investments with an enterprise-grade streaming data platform

- **Hybrid Cloud Integration:** Striim helps to make hybrid cloud architectures easy and reliable by enabling real-time data integration and interoperability among on-premise data sources, edge devices, and cloud environments such as Azure, AWS, and Google. Striim can also be deployed in the Microsoft Azure Cloud and Amazon Cloud as a subscription service.



With wizards-based development UI, Striim makes cloud integration fast and easy.

- **Integration for Operational Data Stores:** Striim feeds operational data stores (ODS) with real-time data to help companies centralize their data from multiple locations for a consolidated view, and offload real-time operational reporting from production databases. In addition, these ODS environments serve as the intermediate step for data warehousing solutions, eliminating the need to extract data from the source systems in batch mode while allowing the use of existing ETL solutions to load the data warehouse.

CUSTOMER EXAMPLE: HP ENTERPRISE

HP Enterprise uses Striim to accelerate its order management process. Historically, the company had no visibility into the order process for up to two days after an order was placed. Striim enables real-time data integration from ERP, MDM, and credit scoring systems into Hortonworks, Kafka, and HPE NonStop environments to enable fast order processing with access to real-time data. Now HPE has real-time visibility to order management activities and can check customer creditworthiness in real time to expedite the process.

2) DETECTING PATTERNS AND ANOMALIES

Today large data volumes are generated continuously from a variety of sources with high velocity. It is easy for organizations to miss critical events

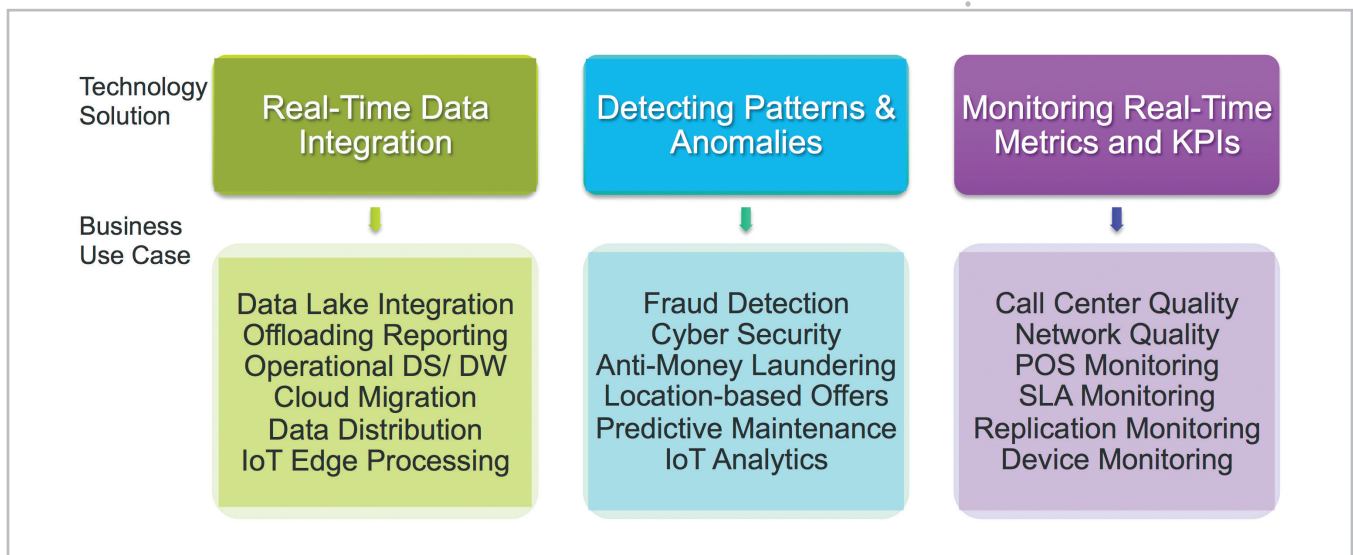
that can potentially harm the business, or forgo the opportunities that could bring profitable growth unless this high-velocity data is analyzed immediately. Striim offers streaming analytics capabilities for data-in-motion to detect patterns and anomalies, enabling immediate detection and action.

Striim uses geolocation data, sensor data, and other relevant enterprise data to obtain accurate insights, fast. Using Striim’s flexible time windows with filtering, multi-source correlation, pattern matching, predictive analytics, outlier detection, and other statistical processing capabilities, end users can easily build applications that support complex requirements in detecting anomalies and patterns.

Striim queries data streams continuously and notifies end users of any critical event immediately via email, SMS, and web alerts. Results are displayed via real-time, interactive dashboards. The platform can also trigger workflows automatically in related applications. The Striim platform can deliver results to various targets for storage and further analysis.

By using an SQL-like language and Striim wizards for development, it is very easy for business groups to modify the applications running on Striim as business requirements change.

Striim captures all security-related events from across the business, and from numerous and varied point security solutions (SIEMs) in real time.



The Striim Platform supports a broad spectrum of use cases in different business domains and industries.

Some of the most common use cases for the pattern and anomaly detection are as follows:

- Security:** Striim captures all security-related events from across the business, and from numerous and varied point security solutions (SIEMs) in real time. Striim performs multi-source correlation immediately after a qualifying event occurs, providing the security analysts with immediate information from combined logs and different alarms on the dashboard. By presenting the correlated data automatically via different views, Striim helps security analysts make accurate decisions

Benefits of Pattern and Anomaly Detection with Striim

- Respond to critical events while they are actionable to outsmart risks and seize opportunities
- Make use of all data, in-time and with context, for accurate event detection
- Correlate different sources to have a comprehensive view
- Gain agility to easily build and modify applications as business needs change

fast. The platform enables businesses to detect potential threats in real-time based on customized rules, and allows users to take proactive action automatically based on predetermined procedures.

CUSTOMER EXAMPLE: LEADING CREDIT CARD NETWORK

A major U.S. based credit card network uses Striim for an enterprise-level security solution. This leading credit card network was using 50+ IT security information and event management (SIEMs) solutions monitoring numerous aspects of their IT infrastructure. Each solution provided a siloed perspective, without a comprehensive view that could explain how a combination of alerts from various points solutions could more accurately signal a specific type of threat. Thus, the security analysts spent excessive time chasing false positives, and risked missing the real threats that needed immediate action. The company's security analytics solutions were using near real-time data from their operational data store, but this was not fast enough for the analysts to be alerted to a potential threat.

With Striim, the credit card network builds a user behavior analytics solution on top of a SIEM platform and achieves a comprehensive and real-time view into all security events to detect threats accurately and immediately. Unlike siloed SIEM solutions, Striim can collect, prepare, and correlate all types of security data, and distribute to various targets. The Striim platform ingests events from a wide range of sources including Syslogs, website logs with card management events, SIEM solutions such as CorreLog, end-point security solutions such as Crowdstrike, firewall logs from Palo Alto Networks, network traffic flow logs from Netflow — all in real time. It also prepares the data with transformations and enrichments with blacklisted IP for easy and accurate analysis. While publishing to the Kafka-based enterprise bus, the Striim platform tracks and records key metrics about all the events collected.

Using real-time security events, the company can immediately uncover endpoint hacking by detecting events such as file executing from the recycle bin, suspicious SVC host activity, PowerShell abuse, DNS abuse, suspicious scripting tasks, remotely-created services, etc. It can also determine immediately if a blacklisted IP is attempting to log in to the website. On the interactive dashboards, users can see live threat maps showing attacking IPs in real time. Striim provides the security analysts with relevant event data for assessing the situation further for a more comprehensive and proactive response and increases their productivity. The pre-processed security events are eventually delivered to their big data analytics environment for further and deeper analysis.

- **Fraud Detection and Prevention:** The Striim platform provides the ability to analyze all relevant data as soon as it is generated to detect potential fraud cases. It provides the results in real time via alerts and

dashboards to support accurate assessment and timely action to prevent fraud. With the ability to analyze different sources in real time, enrich data-in-motion with critical context, simplify the building and iteration of application logic, and trigger workflows, Striim not only minimizes risk for the business but also increases productivity for the analyst team by streamlining the whole investigation process.

- **Predictive Maintenance:** Striim can correlate machine data and sensor data to detect patterns that signal the need for maintenance before there is a malfunction or crash. This proactive approach minimizes potential safety issues, revenue loss, and customer dissatisfaction. Striim customers use the power of the platform to predict maintenance for ATMs, manufacturing machinery, truck fleets, healthcare devices, and more.
- **Real-Time Marketing Offers:** Businesses can use the Striim platform to analyze real-time information about customer's ongoing engagement and location, enrich it with historical data, and check for the patterns that indicate which products or services they might have a higher propensity to buy. With this timely insight, Striim can trigger the workflow to send the right offer to the customer, while the customer is still interacting with the business, or remains in the vicinity. Striim's in-flight analytics capabilities enable businesses to detect unique opportunities to improve customer experience with targeted and timely offers, and generate more revenue.

3) CREATING AND MONITORING REAL-TIME METRICS

To achieve any operational improvement, up-to-date and comprehensive understanding of the current state of operations is the first priority. That's why both lines of business and IT teams monitor key operational metrics and performance indicators to track how well they're tracking against their service level agreements and departmental objectives.

The Striim platform provides a flexible solution to create and monitor real-time metrics that help the business to know and control the Now. It enables operations teams to gain granular and up-to-the-second visibility into business, and perform time-based search and comparison with historical data. With Striim, end users can create sophisticated, yet easy-to-modify metrics using virtually any data source, filtering, aggregation, transformation, and enrichments as needed, along with mathematical computation and statistical processing with customizable, jumping and sliding time windows.

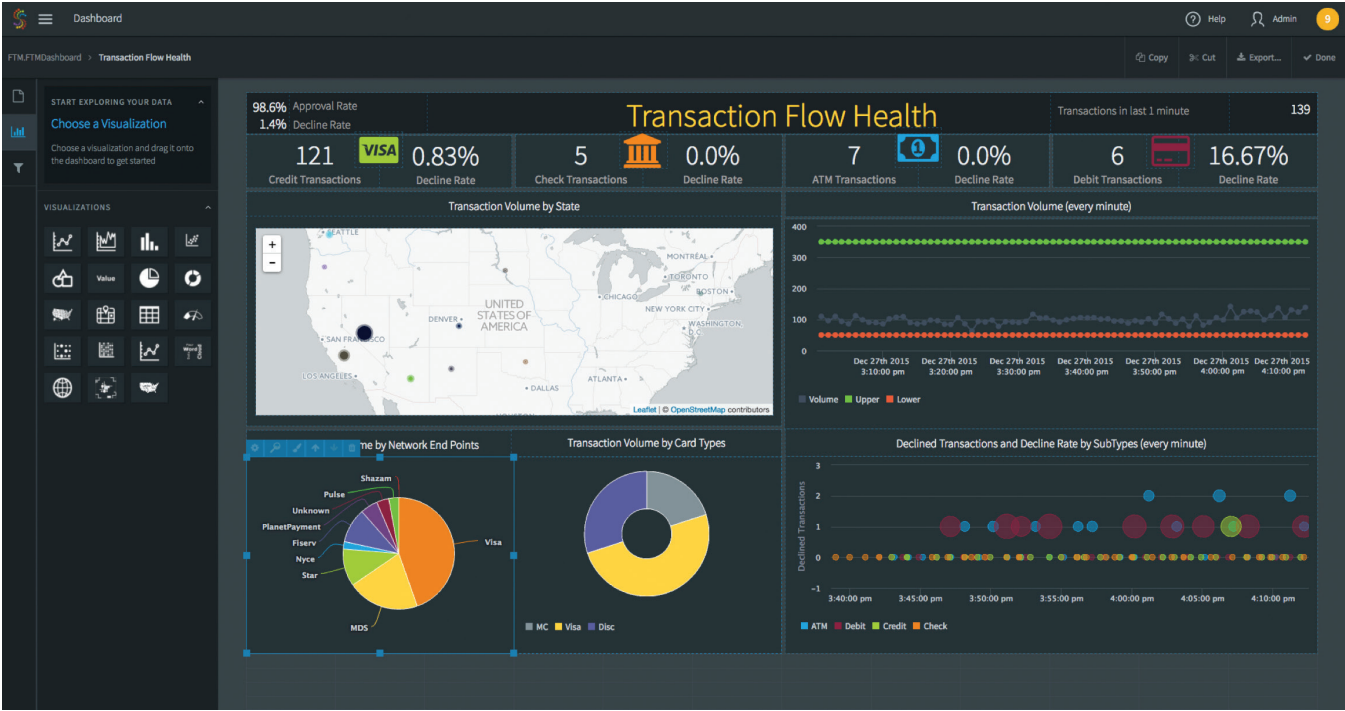
Thanks to Striim's streaming data integration and analytics capabilities and in-memory processing, the metrics are updated in real time as new data streams in from various sources. Using Striim's wizard-based dashboard builder, end users can build and modify their dashboards in seconds and view continuously and automatically refreshed data. The interactive dashboards display the metrics via tailored views to help end users make smart decisions quickly and with confidence.



By setting up customized metrics using a wide variety of data sources, Striim users can see the emerging trends as they happen — before they turn into a problem or missed opportunity — and work on a proactive response.

Real-time metrics monitoring can be used in any industry or operation, from telecommunication to banking to IT infrastructure. Some of the common use cases include:

- **Network Quality Monitoring:** By monitoring key metrics from network routers, customers can detect any potential performance issue before it affects customers. Viewing customizable metrics in real time allows businesses to be more proactive in improving service levels, and better adhere to SLAs with customers.
- **Transaction Monitoring:** Financial services companies can ensure high-quality service and continued revenue growth by closely monitoring how well they process customers' transactions throughout the day. Via real-time dashboards that display transaction volumes across different regions and metrics (such as approval rate and decline rates), companies can detect any outages and delays in transaction processing that may affect SLAs. With this timely insight, employees can resolve the issue before it damages the customer relationship. As business requirements change, analysts can easily modify the metrics and dashboards.



Striim offers real-time performance dashboards with flexible views.

- **Replication Monitoring:** Striim's non-invasive data capture and streaming capabilities make data validation in replicated environments easy and reliable. Striim uses change data capture (CDC) on both the

source and target database to continuously read database logs. In real-time, Striim ships a record signature from both the source and target into a predefined time window and uses correlation to match the primary and secondary transactions. For every record coming in from the primary system, if the secondary system does not show a commit within the time window, the transaction is flagged and an alert is created.

With this solution, Striim can track replication speed and volume, detecting issues when the performance degrades, before any crash, to prevent data loss. Unlike legacy replication monitoring solutions, Striim enables IT teams to be proactive and predict a possible crash or failure before it happens.

CUSTOMER EXAMPLE: RAYMOND JAMES

Raymond James had difficulty with its home-grown tool for monitoring 500+ Oracle GoldenGate replication groups that enable their HPE NonStop disaster recovery environment, and support the data warehousing solution. The monitoring tool was receiving 800+ alerts per day related to replication process performance issues and crashes. The after-the-fact emails were time-consuming to assess, and not actionable for the replication team.

Raymond James now uses Striim to monitor the GoldenGate replication environment. The Striim application sends real-time alerts when a GoldenGate replication process slows down, before any crash. Striim's real-time dashboards replaced the after-the-fact email alerts and provide the replication metrics that the IT teams need to manage the disaster recovery and data warehousing systems.

Using the Striim application, Raymond James' replication team gained 60% productivity. Development teams now have up-to-the-millisecond visibility into their replication systems and SLA reports.

4) IOT EDGE PROCESSING AND EDGE ANALYTICS

By connecting physical devices to the digital world, the Internet of Things (IoT) has transformed and enriched all aspects of business operations. By tapping into IoT, businesses can now better understand customers' needs and experiences, empower their employees to respond to customer needs, offer new products that unlock new revenue streams, and optimize their operations to deliver higher quality products and services faster to the market.

In addition to combining sensor data with other machine data and enterprise data to provide a complete view and accurate analysis, the Striim platform provides businesses the flexibility to distribute their data processing and analytics in a way that serves their business most effectively. With Striim, customers can perform real-time edge processing and edge analytics to

Benefits of Creating and Monitoring Metrics with Striim

- *Gain granular and real-time understanding on your daily operations*
- *Create and modify complex metrics easily*
- *Respond to emerging trends proactively to minimize risk, and improve outcomes*
- *View real-time metrics via customizable, interactive dashboards*

Benefits of Edge IoT Processing and Analytics with Striim

- *Minimize risks and improve operations by obtaining the insight right where it is needed*
- *Gain insight from high-velocity data before it loses value*
- *Adapt a smart IoT architecture for better storage management by filtering out, aggregating and enriching data for use-case-driven, downstream analytics*

build a smart architecture for better storage management, and accelerate insight to the point of data collection.

Via edge processing, Striim filters out repetitive and normal (i.e. uneventful) data, transforms, aggregates multiple data streams, and enriches with context before delivering to a centralized storage for further analysis. By pre-processing data-in-motion, Striim enables businesses to extract the signal from the noise, improve storage management, and deliver more use-case driven analytics downstream.

Edge analytics builds on the same notion and offers fast-path analytics and anomaly detection to trigger alerts and actions without having to reach the Cloud or an on-premise analytics environment. Edge analytics is used in a variety of industries — from manufacturing to smart cities to healthcare — to maximize the value of high-velocity data. Below are a few use case examples:

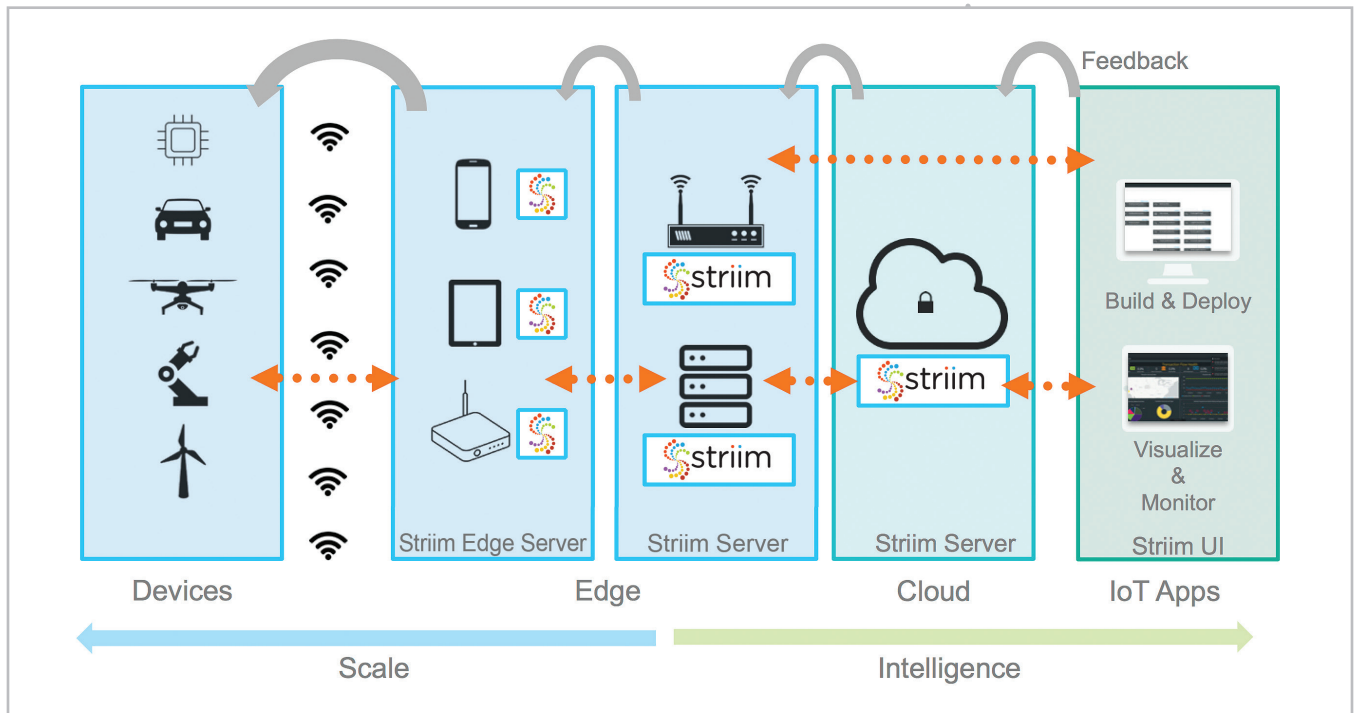
- **Real-Time Product Quality Control in Manufacturing:** Companies can detect product quality issues by comparing real-time machine and environment data to the historical data gathered from high-quality products. For example, if the size of the item as measured by real-time imaging is too large and quality failure is predicted, then Striim can notify operators to prevent that item from continuing to the next section where it may damage the machine and slow down the process significantly.
- **Patient Health Monitoring:** By analyzing device data in real time from patients, healthcare providers can detect health issues faster and act on immediate alerts with automated responses where possible. For example, the alert from edge analytics can trigger a care coordination protocol at the hospital, or even instantly dispatch emergency services if the patient's vitals drop dramatically at an outpatient clinic.
- **Remote Monitoring of Oil and Gas Operations:** By having real-time analytics at the edge, oil and gas companies can enable early detection and predict a major malfunction or disaster, and act immediately to prevent a catastrophe without having to wait for centralized analysis before taking action.

By using Striim for edge processing and edge analytics, companies can build a smart data architecture that enables insight where it is needed, and store only the required data in the form it is needed for use-case driven cloud analytics.

In addition to edge analytics, businesses can also use Striim for advanced real-time IoT analytics in the cloud. Striim partners with several leading IoT vendors to enable end-to-end IoT management solutions, including Microsoft, Fujitsu, Statistica, and Dell EMC.

Via integration with cloud solutions such as Azure IoT Hub and Azure Event Hub, Striim offers streaming analytics on data it streams into the Azure environment.

With this capability, companies can gain further insight into their IoT data by combining it with other relevant data coming from on-premises systems or the Azure environment itself. Striim can combine unstructured, semi-structured, and structured data to enable Striim can deliver the processed data and insights into other Azure solutions such as Azure Storage and Azure HDInsight.



Striim provides businesses the flexibility to distribute their data processing and analytics in a way that serves their business most effectively.

KEY VALUE DRIVERS

With its comprehensive and enterprise-grade platform, and solutions for a variety of business and technology initiatives, Striim offers a clear return on investment to its customers.

For the business, Striim enables enterprises to minimize risk factors to protect the customers and the business from a loss by enabling applications that:

- Detect and stop internal and external security threats
- Ensure compliance with SLAs and government regulations
- Prevent fraud, theft, billing leakage

In addition, it allows businesses to uncover time-sensitive opportunities that drive revenue growth and profitability by enabling applications that:

- Enhance customer experience with timely services and promotions
- Offer new, innovative services to better meet customer demand and gain competitive advantage
- Reduce inefficiencies and costs by streamlining operations with automated intelligent actions

Striim supports the gradual adoption of streaming architecture by enabling to start with small, ROI-focused use cases.

For IT infrastructure teams, Striim delivers the following value drivers:

A smart, modern architecture for fast innovation and time to market by:

- Making real-time data available across enterprise systems, cloud solutions, and IoT devices
- Delivering new analytical applications within days, without requiring complex coding skills
- Pre-processing streaming data before delivering into Hadoop, cloud environments or other targets to reduce storage and simplify downstream queries and analysis
- Enabling the rapid creation of dashboards for faster insight
- Supporting the gradual adoption of streaming architecture by enabling to start with small, ROI-focused use cases

Increased IT productivity by:

- Automating many time-consuming monitoring and data preparation tasks
- Empowering business groups to modify their applications, enabling IT to focus on strategic projects
- Enabling re-use of streaming data pipelines for multiple streaming data analytics applications, minimizing overhead and reliance of different skill sets

Reduced total cost of ownership by:

- Supporting multiple use cases most efficiently by standardizing on a single platform for streaming data processing and analysis across the enterprise
- Extending the lifespan of Big Data environments by pre-processing and filtering raw data in-flight
- Allows companies to easily scale with their existing commercial or open source investments, such as Kafka and Hadoop, with an enterprise-grade streaming data platform
- Future-proofing streaming data solutions by leveraging an enterprise-grade platform without additional cost for reliability, scalability, and security

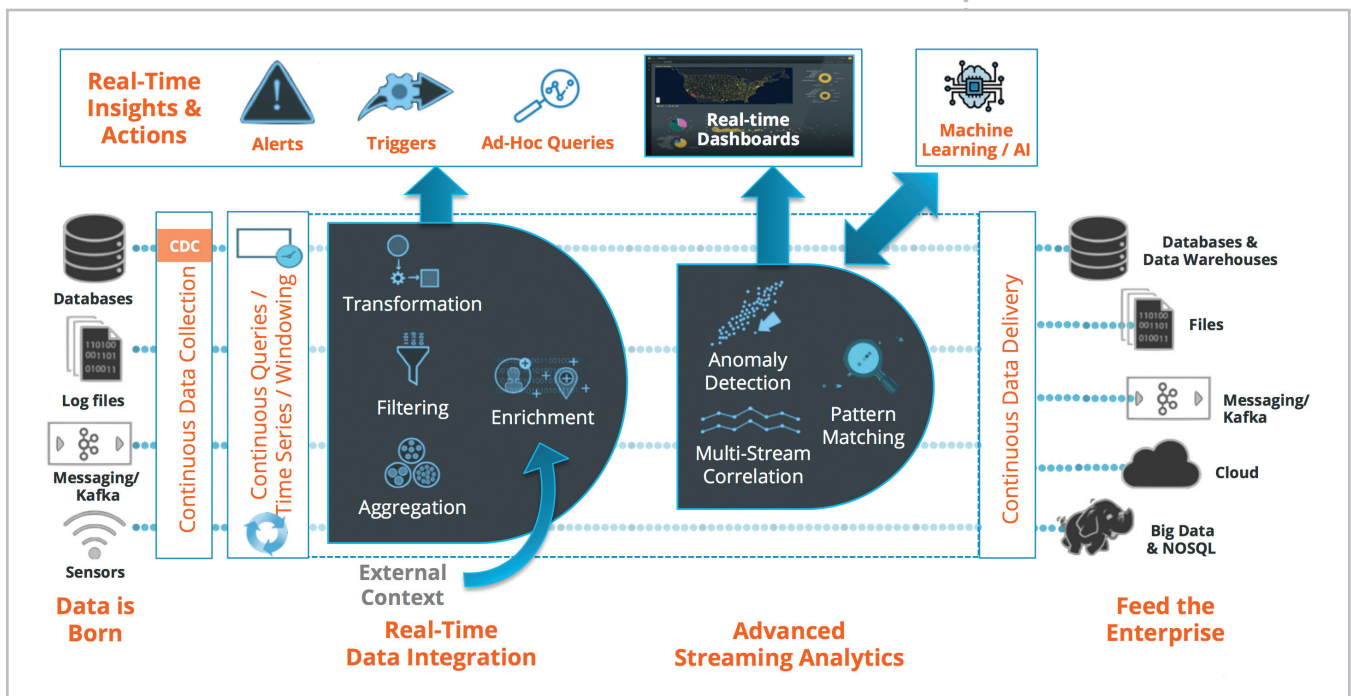
Why Striim?

Striim is unique among streaming integration and analytics offerings by providing all four of the following in a single platform.

- 1. Striim is End-to-End:** The Striim team designed the platform with the understanding that to deliver streaming analytics that transforms operational decision making a comprehensive real-time data processing and integration component is of utmost importance. By including streaming integration, analytics and visualization in a single platform, Striim has differentiated itself from other analytics or integration solutions

that address a narrow slice of the challenges involved with streaming data management and analysis. Starting with data ingestion from the sources, Striim filters, transforms, aggregates, enriches, performs streaming analytics, triggers alerts and workflows, enables visualization on dashboards, and delivers the results to the rest of the enterprise.

In addition, it offers support for wide range of data sources and targets, from heterogeneous databases via non-intrusive change data capture and log files, to message queues, sensors, Big Data, and cloud environments. It supports unstructured, semi-structured and structured data types such as JSON, XML, delimited, binary, text, and change records.



Striim provides an end-to-end platform with real-time data integration, streaming analytics, and data visualization in a single platform.

Striim incorporates best of breed open source technologies to bring end users the latest innovation in a comprehensive and cost-effective solution. For example, the Striim platform includes Kafka out-of-the box for high-performance, persistent messaging and Elasticsearch for indexed results store.

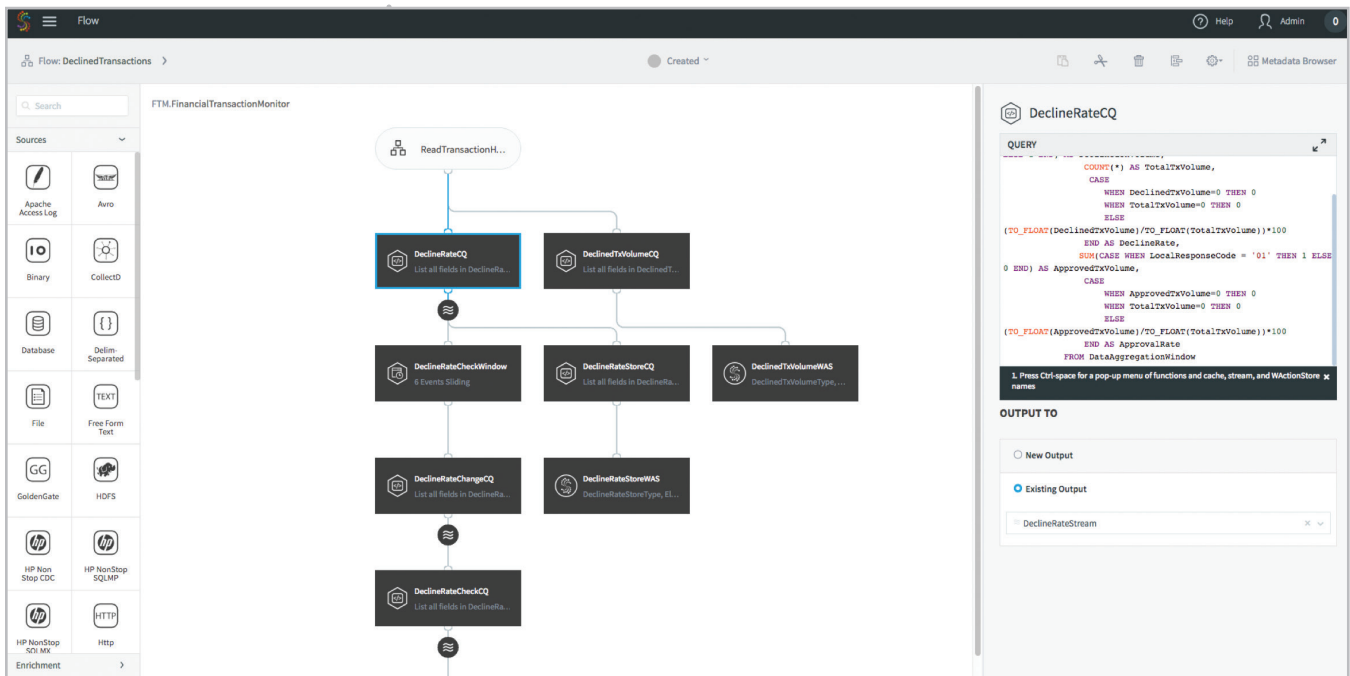
- 2. Striim is Enterprise-Grade:** Founded by the executives of GoldenGate Software, Striim was designed from the beginning to handle the demands of high-velocity data in mission-critical environments with security, high performance, scalability, and reliability.

Secure with built-in authentication, authorization, protection and encryption. Applications can be secured easily through roles, while specific items can be locked down to the stream level using granular security rules. All data can be encrypted in-flight.

High performance and highly scalable with a distributed, modern architecture that combines highly optimized data serialization and windowing with in-memory computing. Striim enables linear scale-out using a low-cost compute infrastructure to support extreme and varied data volumes and velocity.

Reliable with a fault-tolerant architecture and “exactly once” processing. Recovery ensures no events are missed or processed twice, and takes window contents into account.

- Striim is Easy-to-Use:** Striim provides an intuitive development experience for streaming data analytics applications with wizard-based configuration and coding and a drag-and-drop dashboard builder, and speeds time-to-deployment with prebuilt data pipelines. With its SQL-like language, the Striim development environment is familiar and easy-to-use both business analysts and developers, without compromising the ability for technologists to easily customize and code in Java. Additionally, sophisticated Java algorithms can be imported and made available for business users, offering unlimited extensibility.



Striim accelerates development with a drag-and-drop-style Visual Flow Designer, in addition to command-line interface.

Striim is also easy-to-maintain. Cluster nodes, stream analytics applications, and live streaming dashboards are managed via a central visual user interface. The administration window provides an overview of the cluster, stream analytics application deployments, and node health. Integration with existing monitoring solutions is available via a RESTful API.

- Striim is Easy-to-Integrate:** Striim helps businesses effortlessly modernize their legacy data infrastructures to integrate, analyze, and visualize streaming data. It integrates rapidly with existing data management

solutions in the enterprise (both open source and commercial), allowing users to extend the platform by bringing existing Java code, and integrating it with machine learning solutions to operationalize previously-defined models via stream processing. Striim streamlines the adoption of a streaming architecture by allowing customers to start with a single, small use case and expand their solutions with new sources and targets, and analytical applications as needed.

Conclusion

Striim is a software platform that combines real-time data integration with streaming analytics and enables companies to make smart decisions on Fast Data. By leveraging sub-second data ingestion, processing, integration and analytics, Striim arms businesses with the relevant data in-time and in-context to take the best actions at the speed of their business.

In today's fast-paced, innovation-driven world where legacy environments no longer suffice to push the business forward, Striim makes the adoption of a next-generation streaming architecture a seamless and virtually a risk-free process. It helps organizations manage their data across legacy and next-generation infrastructures including cloud, Big Data, and IoT, with flexible solution topologies. Thus, it streamlines the modernization of data and analytics infrastructure. Organizations that want to keep pace with market demands, and survive the unrelenting waves of change coming from disruptive technologies, can easily plug-in Striim to upgrade their data architecture. With in-time, in-context insights at their disposal Striim customers navigate the stormy waters of today's business environment with confidence and ease.

With its SQL-like language, the Striim development environment is familiar and easy-to-use both business analysts and developers, without compromising the ability for technologists to easily customize and code in Java.





Connect with us:

 www.striim.com/blog/

 www.linkedin.com/company/striim

 www.facebook.com/striim

 www.twitter.com/striimteam

 www.striim.com/youtube

For more information, or to schedule a free trial, please contact us at info@striim.com or at **+1 (650) 241-0680**

www.striim.com