

On the Radar: Striim converges streaming intelligence and integration

Publication Date: 04 May 2017 | Product code: IT0014-003268

Tony Baer



Summary

Catalyst

Having covered fast data, and its predecessors such as CEP (complex event processing), we have noted that the full process of collecting, aggregating, correlating, and analyzing data in motion has typically required a chain of multiple tools. Striim, founded by veterans of the real-time database replication space, takes a different approach. It provides a streaming analytics engine, but complements that intelligence with upstream real-time data integration that delivers an end-to-end solution from data source to insight.

Key messages

- Striim provides a SQL-based end-to-end solution that is broader than what is offered by other streaming analytics engines and tools.
- Striim's real-time integration capability leverages expertise that company founders built from experience designing change-data-capture replication.
- Ultimately, Striim's biggest competition will come from cloud providers offering real-time data pipeline services.

Ovum view

In a market where it has become difficult to know what can properly be labeled as "streaming," Striim's prime challenge is communicating how its streaming integration back end differentiates it in an ecosystem comprised of analytics, data flow engines, and message broker providers. Ultimately, its prime competition will come from Amazon and Google, whose one-stop-shop approaches to real-time streaming are *not* exact matches of what Striim offers.

Recommendations for enterprises

Why put Striim on your radar?

Consider Striim if your organization is looking for a single solution that integrates and ingests multiple – and varied – sources of data at scale to deliver real-time streaming analytics.

Highlights

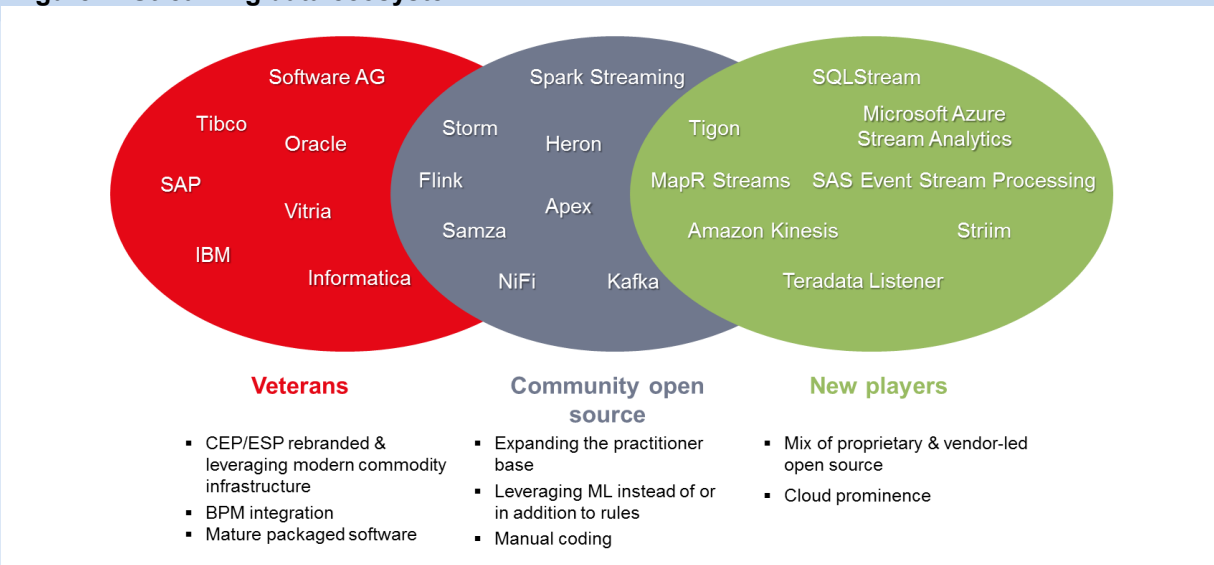
Background

Over the past year, Ovum has seen a sharp upswing in client queries about real-time streaming technology. The use cases are numerous: real-time IoT ingestion, log correlation, location-based applications, emerging connected car applications, and the rapidly mutating nature of security threats. These use cases are real and have brought streaming analytics to the forefront for a growing number of organizations across sectors as diverse as retail, transportation, financial services, digital services, and others. Not surprisingly, we have seen an explosion of streaming-related vendors. However,

comparing vendors within the ecosystem is like comparing apples to oranges (see Figure 1). Some tools provide streaming analytics while others perform broker functions, while others focus on downstream data flow. And there is a mix of open source and proprietary; so in many cases, the competing technologies in the streaming ecosystem will not necessarily have the tooling for development and administration that would be associated with enterprise software.

But there has always been a hurdle with streaming: it typically required a best-of-breed stack of point tools for building end-to-end, real-time data pipelines that are essential to these applications. Data must be aggregated and integrated, sent across a messaging pipeline, and then parsed and analyzed to provide actionable insights. In many cases, additional analytics tools were also required for performing more sophisticated analytics or event correlation.

Figure 1: Streaming data ecosystem



Source: Ovum

Not surprisingly, Kafka, an open source technology, has become popular for imposing order when you have multiple incoming streams and targets for data. But Kafka only provides brokering; you still need additional tools to ingest, transform, and analyze the data streams.

Striim was founded to deliver an end-to-end solution that starts with real-time continuous collection, aggregation, and integration, and ends with producing timely insights and triggering workflows. It was started by veterans of data replication, with many of the founders coming from Oracle GoldenGate. Their design goal was to develop a scalable end-to-end solution for capturing, monitoring, transforming, and analyzing data in motion. It includes a blend of proprietary and open source technology. As for the branding, the two i's of "Striim" stand for "integration" and "intelligence."

Striim is an in-memory engine for building the data pipelines. Reflecting on the GoldenGate background, Striim has connectors that perform change-data-capture (CDC) from database logs; they also have adapters for sourcing data in multiple formats, from CSV to JSON, Avro, and others. Additionally, Striim can source and deliver data to and from Kafka and the Hadoop Distributed File System (HDFS). It uses continuous in-memory queries to process and analyze data in transit, and a distributed in-memory cache to load data for enrichment purposes. For developers, it includes a declarative programming environment for designing integrations; filters, transforms, enrichments, and joins; time-windowed SQL queries; complex event processing; multisource correlation, and anomaly

detection. For the messaging backbone, Striim uses ZeroMQ, an open source distributed messaging protocol that was designed for embedding into applications, as well as a built-in Kafka cluster for cases that require persistent messaging.

Striim can be applied to the usual mix of use cases associated with real-time streaming. A good example of its capabilities can be shown in a security information and event manager (SIEM) use case applied to a credit card network for providing cybersecurity. Striim analyzes the log files of 50+ SIEM solutions (e.g., firewalls, proxies, and monitoring servers), combines the output with session data, and conducts pattern matching to identify outlier or suspect events. By taking a real-time view of data across the SIEM silos, the Striim solution was able to significantly improve the accuracy of alerts.

Current position

Founded in 2012, Striim has raised \$42m in three rounds of funding, including angel funding. From a technology standpoint, Striim spans several product segments, including streaming analytics, messaging, real-time data integration, and data visualization. Its solution runs both on-premise and in the cloud; Striim recently launched a streaming analytics service that is hosted in the Microsoft Azure cloud. In fact, one of Striim's most common use cases is supporting data flows that bridge hybrid cloud/on-premise environments. Examples include feeding log-based change-data-capture from on-premise databases to cloud targets such as Amazon Redshift, Microsoft Azure Storage, and Google Cloud PubSub.

Striim's capabilities touch areas that are covered by household names in analytics and integration middleware, such as IBM, Tibco, Microsoft, and Software AG; but the main distinction is that Striim's rivals address these capabilities with separate products. Although not a direct match, Striim's biggest competition ultimately will come from Amazon and Google; they offer end-to-end, cloud-based streaming data analytics that provide programming and analytic environments for data pipelines. But unlike Striim, they are only available in the cloud and do not *directly* accommodate streams from on-premise databases; for that, a best-of-breed approach involving third-party tools (such as Striim) would be required.

Striim's differentiator stems from the company's heritage with real-time integration, data replication, and CDC; its solution ventures further upstream to capture and integrate streams from data sources that would otherwise require separate tools or more loosely coupled API integrations.

Data sheet

Key facts

Table 1: Data sheet: Striim

Product name	Striim	Product classification	Streaming integration and analytics
Version number	3.7	Release date	April 2017
Industries covered	All	Geographies covered	North America, EMEA
Relevant company sizes	Midsized to large	Licensing options	Subscription and perpetual
URL	www.striim.com	Routes to market	Direct, partner
Company headquarters	Palo Alto, California, US	Number of employees	50

Source: Ovum

Appendix

On the Radar

On the Radar is a series of research notes about vendors bringing innovative ideas, products, or business models to their markets. Although On the Radar vendors may not be ready for prime time, they bear watching for their potential impact on markets and could be suitable for certain enterprise and public sector IT organizations.

Further reading

Fast Data 2015–16: The Rebirth of Streaming Analytics, IT0014-003064 (October 2015)

"Fast data analytics requires even faster governance," IT0014-003158 (October 2016)

"MapR Streams harnesses the confluence of IoT speed and scale," IT0014-003229 (February 2017)

Author

Tony Baer, Principal Analyst, Information Management

tony.baer@ovum.com

Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at consulting@ovum.com.

Copyright notice and disclaimer

The contents of this product are protected by international copyright laws, database rights and other intellectual property rights. The owner of these rights is Informa Telecoms and Media Limited, our affiliates or other third party licensors. All product and company names and logos contained within or appearing on this product are the trademarks, service marks or trading names of their respective owners, including Informa Telecoms and Media Limited. This product may not be copied, reproduced, distributed or transmitted in any form or by any means without the prior permission of Informa Telecoms and Media Limited.

Whilst reasonable efforts have been made to ensure that the information and content of this product was correct as at the date of first publication, neither Informa Telecoms and Media Limited nor any person engaged or employed by Informa Telecoms and Media Limited accepts any liability for any errors, omissions or other inaccuracies. Readers should independently verify any facts and figures as no liability can be accepted in this regard – readers assume full responsibility and risk accordingly for their use of such information and content.

Any views and/or opinions expressed in this product by individual authors or contributors are their personal views and/or opinions and do not necessarily reflect the views and/or opinions of Informa Telecoms and Media Limited.

CONTACT US

www.ovum.com

analystsupport@ovum.com

INTERNATIONAL OFFICES

Beijing

Dubai

Hong Kong

Hyderabad

Johannesburg

London

Melbourne

New York

San Francisco

Sao Paulo

Tokyo

