

How to Easily Get Real-time Analytics from Kafka



What is Kafka?

Apache Kafka is open source software designed to distribute data in high volumes among distributed systems including devices, applications, and databases. For example, it transmits Netflix real-time recommendations to subscribers; gathers Uber user, driver, and trip information from mobile devices and GPS signals to forecast demand for surge pricing; and at LinkedIn, where it was invented, it's used to identify and control spam by providing a real-time communications fabric to help users more safely make connections.



Why It's So Popular

Today's applications are everywhere — cell phones, IoT devices, global offices — and it's critical that they connect to each other. Kafka makes it easy to connect distributed applications and to communicate important business events between them, like a GPS signal that a supply chain company uses to see where delivery trucks are located. Businesses use the real-time data from these connected applications to improve their services and make decisions.


Why Getting Insights from Kafka Messages is Difficult

With the rise in digital business applications, Kafka use is quickly increasing, and that's a good thing. But unfortunately, most business intelligence, data science, and data management tools do not natively connect to Kafka; so, to use the data it passes for analytics-aware decisions, developers archive Kafka messages in a database. This data at rest creates a "too late" architecture in which continuous awareness of data currently being transmitted is impossible.



But Now, It's Easy

Getting analytical insights from Kafka usually requires custom coding, plus several complex components including Kafka Connect—which is a developer component that enables streaming data to be ingested into your tools—as well as a database and a business intelligence tool. It's not easy, not off-the-shelf, not real time, and it can be expensive. But TIBCO has changed this. Now, analytics agility can be achieved in minutes with a native low-code, or in most cases, no-code Kafka connection from TIBCO.





With TIBCO, business analysts and business users can now analyze Kafka data in minutes, without months of custom coding and without an intermediate database for storage. You just connect TIBCO Spotfire software to Kafka messages, design your visualizations, and go. In addition, for users looking for more advanced analytics, TIBCO makes it easy to add data science algorithms to data visualizations.

If You Don't Use Spotfire

If you don't use Spotfire analytics, you can still analyze Kafka data, but with certain limitations. Traditional BI tools aren't designed to accept data that's pushed to them, so TIBCO can't magically make Tableau or Power BI real-time. But through the TIBCO Data Virtualization engine for streaming data, we can convert Kafka topics to tables on-the-fly, in memory, without an intermediate data store. By adding caching databases, Kafka messages can be viewed by any BI tool without any coding and without complicating enterprise architecture. You just connect your BI tool to TIBCO Data Virtualization software for instant Kafka data ready for your BI tool.



Business User Access

In the past, Kafka data was only available to developers. Now, business users can gain data-driven benefits and self-service analytical insight into Kafka data, too. By building a bridge between open-source developers and business users, TIBCO helps you come together as a team to manage Kafka-powered business systems.




Formatted Kafka Messages for Business Users

With TIBCO Data Virtualization software, and the TIBCO Streaming four-time market-leading streaming analytics engine for streaming data preparation,¹ you can cleanse, aggregate, transform, and augment real-time Kafka messages however you choose. With a wizard driven, browser interface that connects seamlessly to all your dashboards, TIBCO's no-code streaming analytics tools help you quickly gain value from Kafka messages.

¹ Gualtieri, Mike with Srividya Sridharan and Rober Perdoni. The Forrester Wave™: Streaming Analytics, Q3 2019, September 23, 2019. <https://www.tibco.com/resources/analyst-report/forrester-wave-streaming-analytics-q3-2019>

Why TIBCO and Kafka Are Better Together

With over 25 years of messaging experience, TIBCO provides world class follow-the-sun enterprise support for TIBCO Messaging – Apache Kafka Distribution software. And the synergy between TIBCO and Kafka goes further. With low or no code, business users can easily get real-time insight into data being transported on Kafka, and more advanced users can apply data science and predictive insights in real time.





Examples of Kafka and TIBCO Working Together

TXODDS Delivers Real-time Sports Insights with Kafka and TIBCO

Otherwise known as “the Reuters of betting data,” TXODDS, a real-time aggregator and distributor of sports betting information, uses TIBCO and Kafka to deliver real-time data on sporting events around the globe. Betting organizations and players monitor the real-time state of games and consume the company’s real-time analytical insights to place bets and make their sporting experience more informed, fun, and profitable.

How It Works

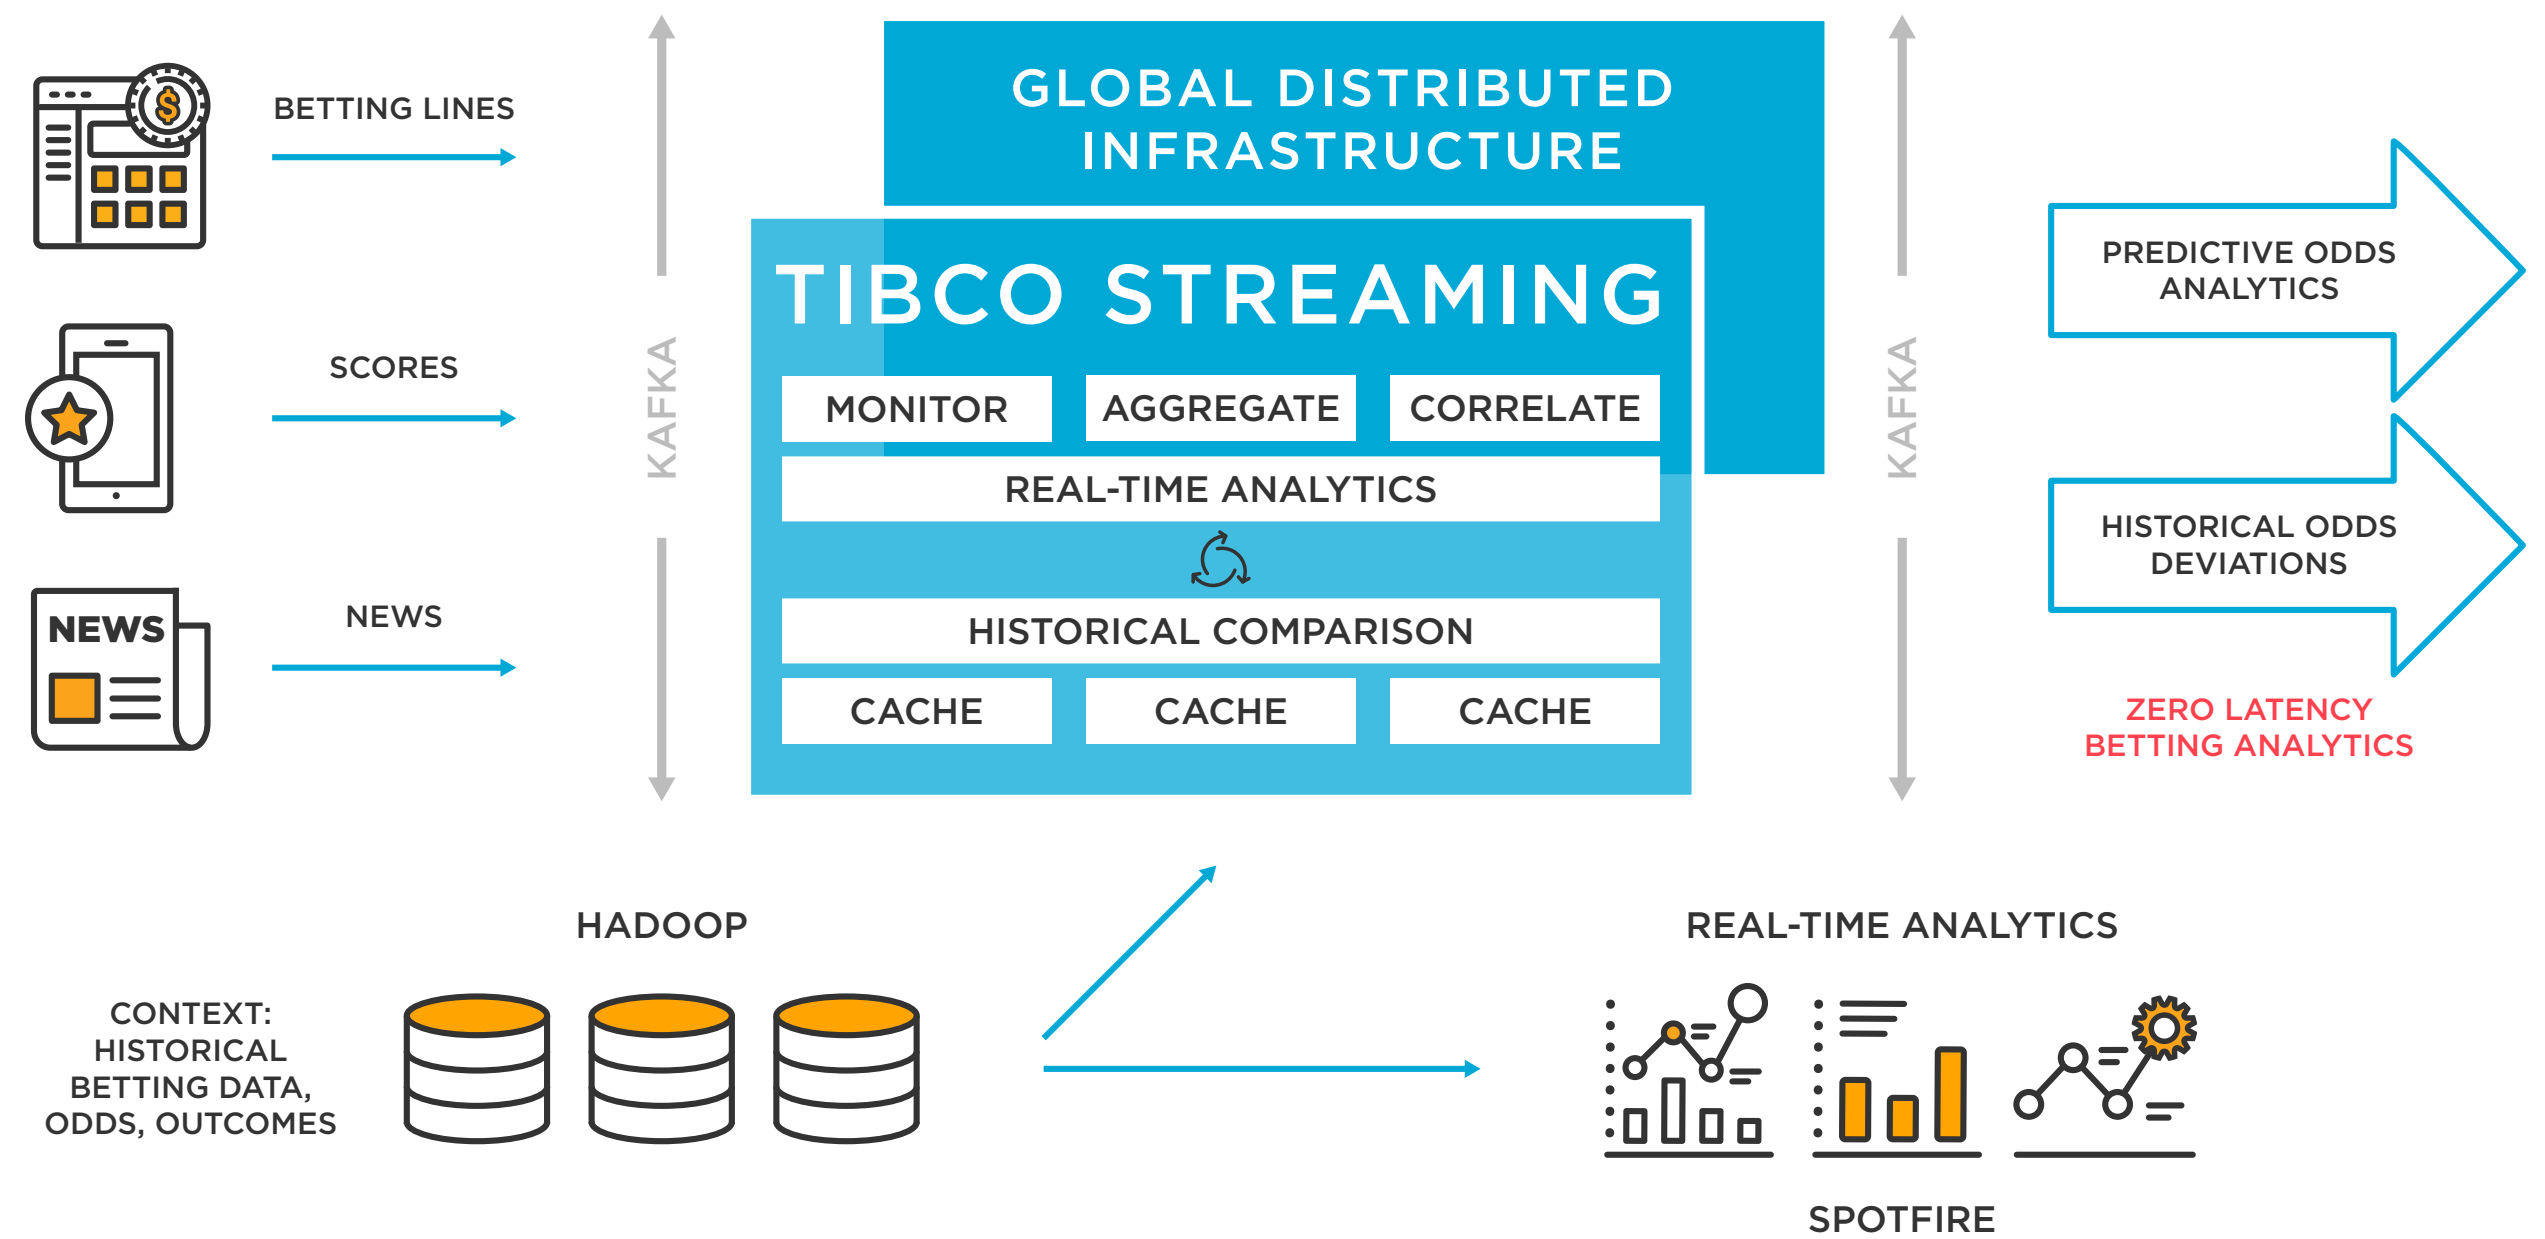
TXOdds absorbs data from systems monitoring thousands of live sporting events in real time, from football to rugby to cricket, using a network of Kafka messages. The messages carry real-time inputs, such as which players are on the field, the score, and even the weather. It’s millions of real-time datapoints — thousands of changes on thousands of simultaneous sporting events.

The Kafka messages are constantly fed to the TIBCO-powered TXOdds “brain,” which includes TIBCO Streaming and Spotfire software

providing sophisticated event-driven rules, AI, and continuously applied predictive models in real time. The brain can compare current game conditions to history, execute sophisticated AI learning models to predict which way the game is likely to go based on in-game data, and transmit a stream of data and predictions to their subscribers to help anticipate betting line movement, and, of course, game results.

Result

This scalable and easy-to-use platform enhances human intelligence with streaming, real-time analytics for the global sporting community, all transmitted via the cloud and API services in real time. By employing TIBCO's low-code and no-code analytics capabilities for processing streaming Kafka data, TXOdds can more quickly create advanced AI-driven streaming analytics, adjust to changes in the dynamic sporting market, and react to volatile conditions that impact bets.





Top Tier Bank Uses Kafka and TIBCO to Monitor Real-time Global Trading Activity

Based on messages flowing on Kafka, a top-tier financial institution uses TIBCO to monitor global trading activity in real time. With this solution, the bank gets a real-time view of the impact of client orders, trading activity, and IT infrastructure all at once.

How It Works

Similar to the TXODDS use case above, the TIBCO “brain” ingests streaming Kafka data from the bank’s entire IT infrastructure and applies AI in real time. This makes the data a single source of real-time truth. Depending on what the user wants, the system can alert, take action, and/or compare historical behavior and deviations from trends, providing continuous awareness for users and IT.



Results

The bank now has predictive banking operations with real-time control and visibility from IT through to the business. It can spot and stop problems before they happen, and fewer resources are required for citizen developers. Citizen development has also increased agility, added greater transparency, and increased the volume of applications deployed, from five apps a year to over 20.

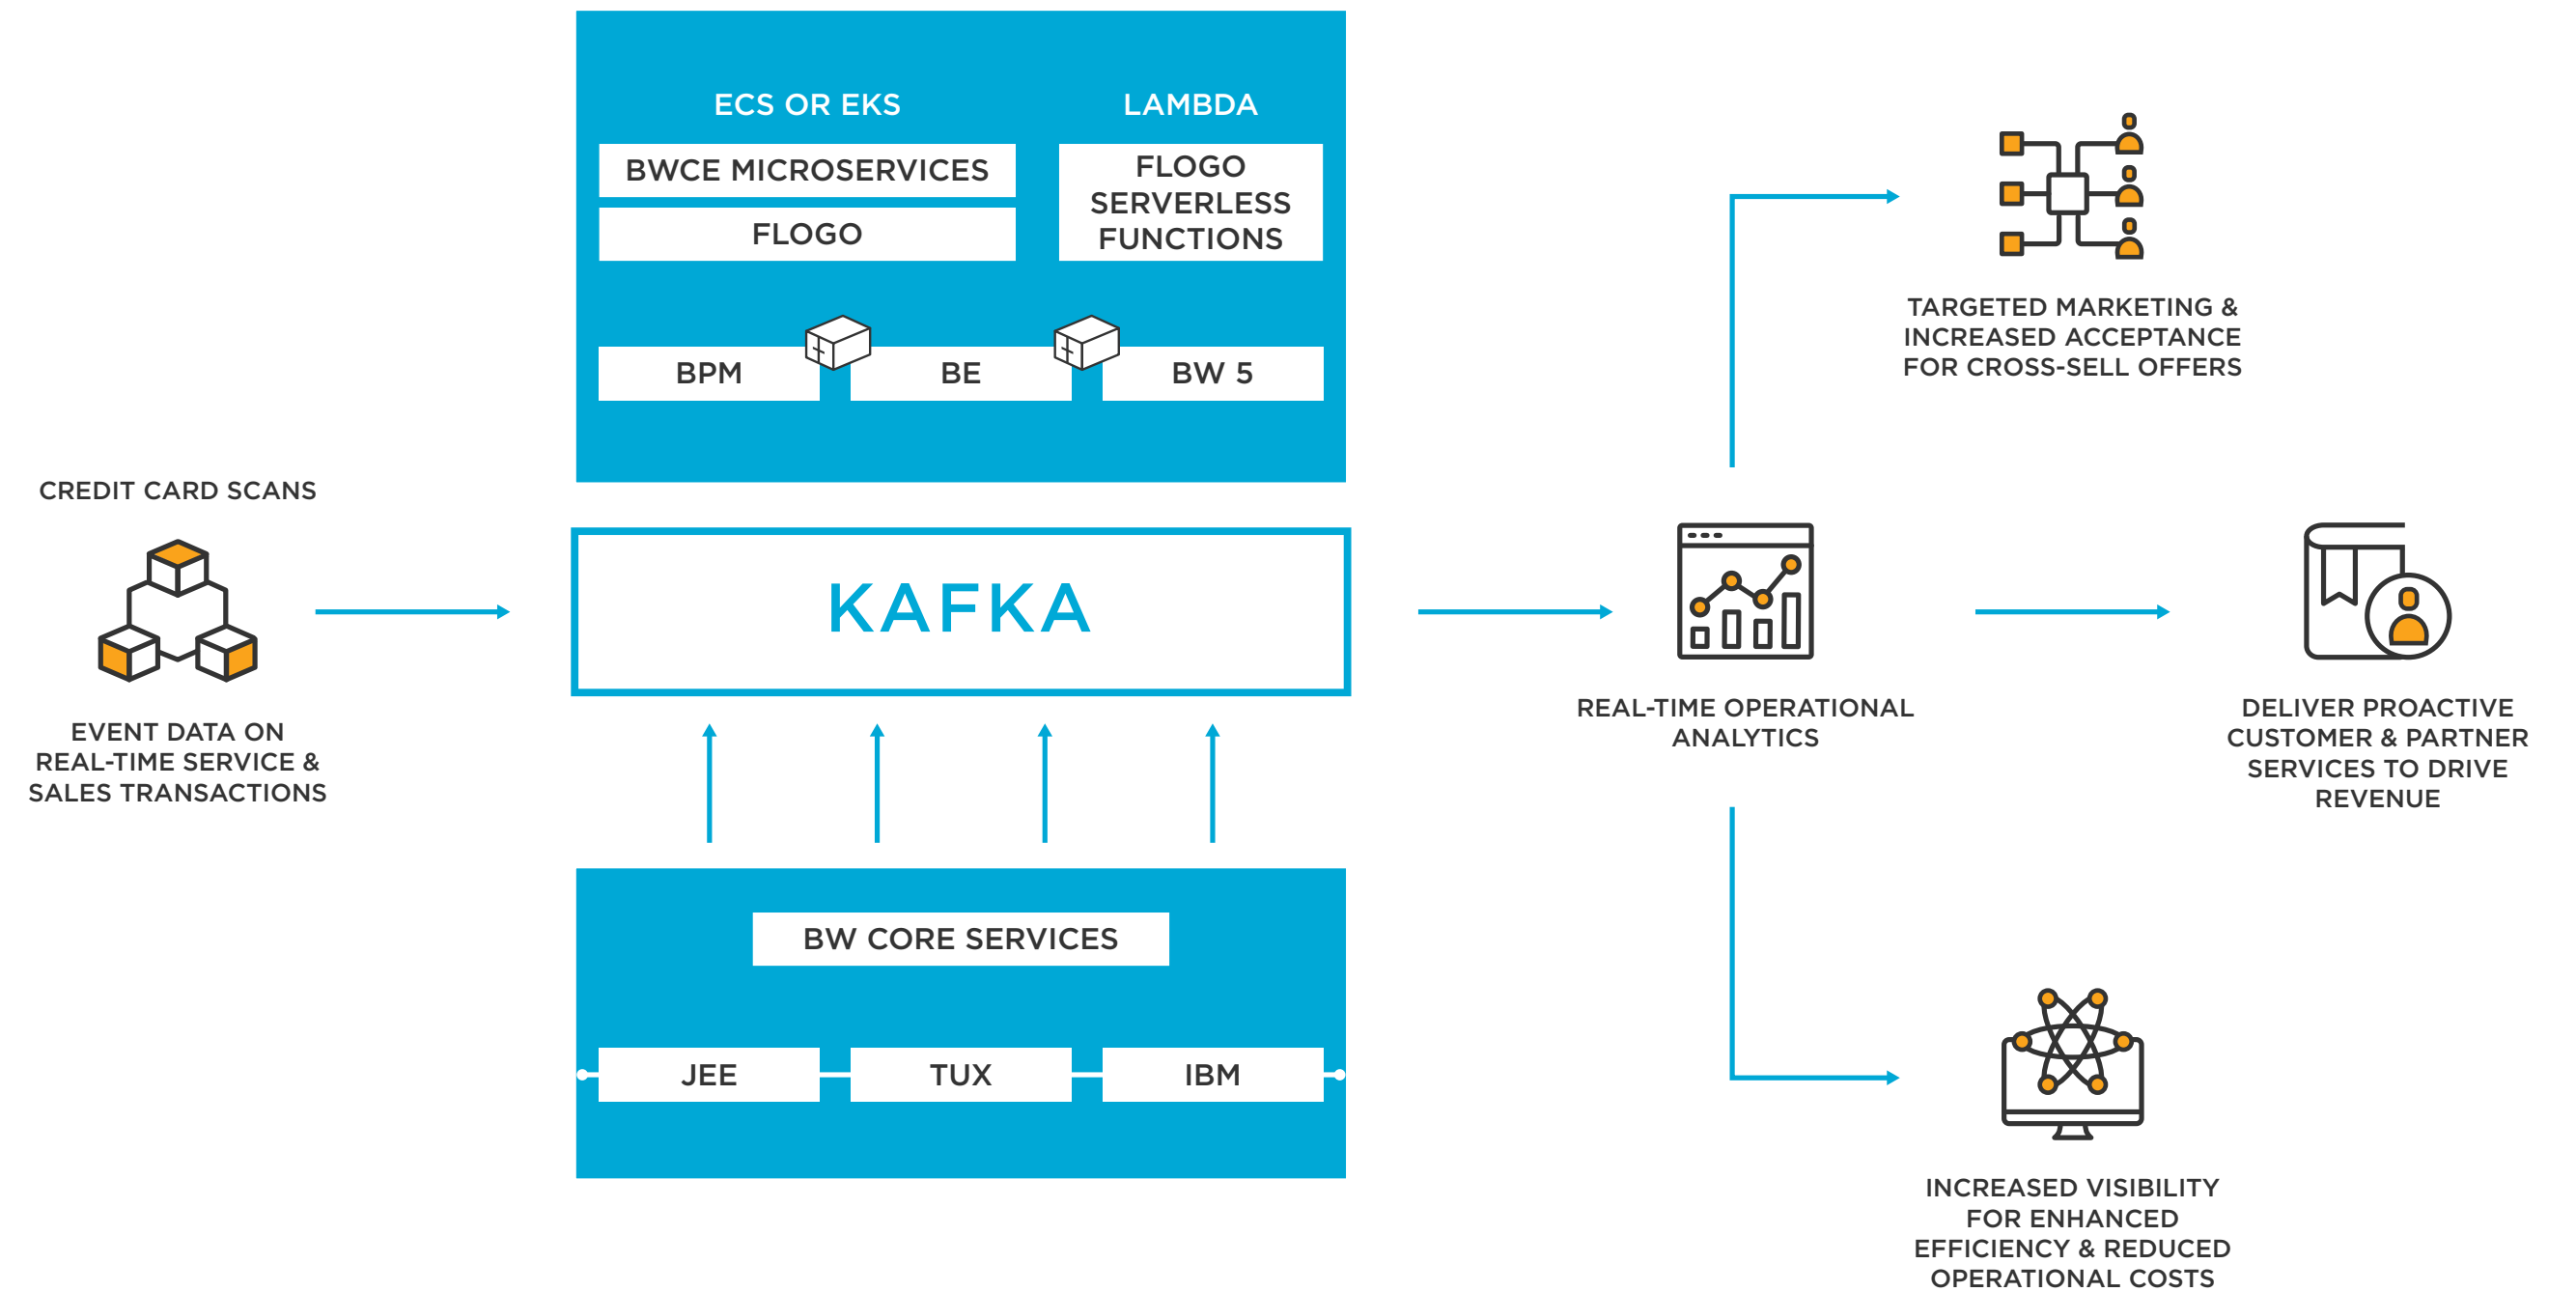
Managing Credit Risk in Real Time with Kafka and TIBCO

A well-known bank is using TIBCO and Kafka to detect credit risk in real time with a design pattern using event processing and analytics to detect fraud. All of this is done on the same solution.

How It Works

Credit card scans are captured as real-time triggered events that stream through Kafka.

The data is then run through TIBCO and compared against both supervised and unsupervised machine learning models. A score is then produced that indicates the risk probability, and another score is provided that indicates how much the transaction deviates from what could be considered normal. The system learns what a fraud-like transaction looks like and makes an autonomous decision based on its learnings, without any human intervention.





Results

The bank can process real-time credit card applications with only a 60-second approval window using a real-time pre-trade/pre-approval calculation and response. The TIBCO solution reduced the custom coding previously required to get data from Kafka to business users, in some cases, from months to hours. On average, that's 77% faster time to market, with more access to real-time data.

Key Takeaways:

TIBCO, a 25 year leader in messaging and a four-time leader in streaming analytics, can help you take Kafka to the next level.

Now, you can leverage Kafka for enterprise analytics with an enterprise streaming platform; a native, no-code connection to Kafka; and real-time dashboards, streaming data science, and data virtualization engine for streaming data.

With TIBCO's industry-leading expertise and innovation in messaging and streaming analytics, you can use Apache Kafka with confidence, backed by enterprise-class support.

[Get started](#) analyzing all your Kafka data in real time today.