

IFOOD.COM.BR

iFood Taps Confluent to Revolutionize Its Real-Time Data Flows and Support Their Growth in a Secure and Cost-Efficient Manner



Headquarters Osasco, Brazil

Industry Meal delivery

Challenge Inability to scale real-time operations alongside business growth due to heavy operational overhead and lack of internal AWS MSK expertise; this resulted in failure to flag production issues in a timely and cost-effective manner.

Solution Confluent Cloud provides a truly cloud-native and complete data streaming platform for event-driven microservices and streaming pipeline to a data lake on AWS to monitor the number of orders and track deliveries on iFood's platform in real time.

Results

- Instantly and cost-effectively scale real-time data availability with growing business needs
- Improved developer velocity and productivity by reallocating engineering resources to more strategic projects
- Lower TCO to seamlessly run mission-critical use cases at scale securely and reliably

For more than 10 years, technology company iFood has mainly focused on meal delivery, serving more than 60 million orders every month in more than 1,200 cities across Brazil. Its vision is to revolutionize the world of food, providing customers with a more practical and pleasurable way of life. It seeks to accomplish this by developing more intelligent ways to efficiently and conveniently produce, sell, and bring dishes to the consumer's table.

To forge a link between consumers, restaurants, and delivery services, iFood offers an online food delivery app/portal to customers where they can order food quickly with no hassles. On the back end, it's supported by a cloud-based data infrastructure.

The challenges with batch pipelines

After early growth and rapid operations expansion in 2016, iFood realized it needed to scale and migrate its data center-hosted back-end operations to the cloud. Fluctuations in demand (e.g., the peaks at Sunday lunch and dinner time) also contributed to this decision, which could be addressed by the elasticity offered by the cloud

iFood was also using a host of databases, services, and applications, which were all tightly coupled together, and batch pipelines to move between these systems. A simple error in any of the services affected other services; troubleshooting the source of the error was time-consuming and ultimately brought their business to a screeching halt.

To solve these problems, iFood decided to adopt Apache Kafka®. Without sufficient resources and expertise to self-manage Kafka, iFood initially went with AWS Managed Streaming for Apache Kafka (MSK) to build streaming data pipelines to monitor orders and track delivery.

However, very quickly, the iFood team found themselves continuing to spend a lot of time managing Kafka despite being on Amazon's cloud-hosted service. The data and platform engineering teams increasingly found themselves spending a majority of time performing operational tasks such as cluster sizing and capacity planning vs. focusing on higher value initiatives. Some of the challenges they faced included, but was not limited to, security authentication issues and manual Kafka software updates and bug patches. Ultimately, iFood decided to switch to Confluent Cloud, a fully managed cloud-native service for their Kafka projects.

Since then, iFood has built streaming data pipelines between 2,000 microservices, allowing the organization to fully decouple their legacy architecture and adopt an event-driven architecture.

They also expanded to a variety of other use cases. They built streaming data pipelines to a data lake on AWS to run analytics and machine learning algorithms. They also leveraged real-time data coming in from the geolocation of their delivery drivers to build tracking capabilities, so customers can track the status of their orders in real time. By leveraging ksqldb, they can run stream processing on their real-time data streams to monitor the number of orders on iFood's platform every second.

"Tracking the number of orders per second using ksqldb is very crucial to us, so we can monitor the use and health of the platform in real time."

— LUCAS VIECELLI, DATABASE RELIABILITY ENGINEER



Technical solution

By leveraging Confluent's fully managed data streaming platform, iFood built streaming data pipelines, which addressed the challenge of sending and processing real-time data prior to sending it to the data lake.

Facing lots of issues with managing Kafka, iFood's teams opted for Confluent Cloud, which provided significant ops burden reduction from its serverless offering and complete feature set, including security and Stream Governance.

Today, Confluent tracks and monitors all the phases that a piece of data goes through within iFood's system, improving the team's control and diagnostic capacity in relation to real-time data flows.

The adoption of Confluent meant the legacy architecture once used for food ordering was migrated to a decoupled, asynchronous model that allows for much greater resiliency—with an event-driven architecture, when a new order is generated through the iFood website or application, multiple parties are able to immediately use the information from this new event in real time.

Since Confluent makes it very easy to connect any other technology to this flow, the data can be enriched and processed with other tools before it reaches the final destination, the data lake.

This high degree of flexibility greatly increases the ability to create new applications and offer new uses of the data across the company. At the same time, there are risks. To avoid these, Confluent developers helped map out a series of best practices to ensure the security and governance of both the data and the system as a whole, something considered to be fundamental to the success of the deployment.

According to Viecelli: "By having the data in Confluent, we are able to seamlessly build streaming data flows with limited engineering resources using fully managed ksqlDB to power aggregations and real-time transformations. This helps us meet the real-time data needs of various teams across iFood and also helps us understand that everything is going smoothly, using metrics like 'how many orders are coming into the platform per second.'"

"By completely decoupling our data architecture, maintaining downstream compatibility, and leveraging ksqlDB for continuous stream processing for some of our streams. Confluent enables all our teams to have self-service access to data as if it were a consumable product. With Confluent, every team, every microservice, every system can continuously act and react on the most up-to-date enriched view of data the moment it's created, enabling us to meet the real-time needs of our customers."

— LUCAS VIECELLI, DATABASE RELIABILITY ENGINEER



Business results

Instantly and cost-effectively scale real-time data availability with growing business needs.

"The iFood system receives more than 60 million orders per month, and the big concern is making the data available to any internal team who wants to use it in real time," said Lucas Viecelli, database reliability engineer and manager of the Kafka implementation team at iFood. "Without Confluent Cloud, we had to manually and repeatedly monitor and scale our capacity up and down to meet demand, which was an impossible ask for our small and nimble teams. Now, Confluent Cloud enables us to seamlessly and elastically scale to meet demand without any burden on our platform operations teams."

Improved developer velocity and productivity by reallocating engineering resources to more strategic projects.

"With Confluent Cloud, engineers can now focus on the business logic instead of developing and maintaining low-level Kafka toolkits," said Viecelli. "For example, by leveraging fully managed Schema Registry, our developers can seamlessly connect to any data system while maintaining schema compatibility, version control, and quality assurance without having to manage the underlying infrastructure."

Lower TCO to seamlessly run mission-critical use cases at scale securely and reliably.

With the increasing reliance and footprint of Kafka to support iFood's growing business, Confluent Cloud helps significantly lower the total cost of ownership by reducing the engineering resources required, operational burden and risk, along with infrastructure spend.

Over time, the team learned more about Confluent and began to see opportunities to use it for other capabilities beyond sending data to the data lake, such as acting as the central mechanism for transmitting production data between microservices and aggregations. The teams are migrating more applications and data to Confluent beyond the critical flow backend. Gradually, Confluent has become the central nervous system at iFood, where all their data in motion is managed through a single platform powering 700 applications to help them unlock the full value of their data.

Last but not least, everyone at iFood agrees that during this process, the support and expertise provided by Confluent made all the difference to ensure a reliable and resilient data streaming platform that can support their entire business.

Learn More About iFood

<https://www.ifood.com.br/>