

BUZZVIL.COM

Seamlessly Manage Massive Volumes of Real-Time Ad Data with Confluent Cloud



Headquarters South Korea

Industry Technology (AdTech)

Challenge Establish a single source of truth for real-time ad transaction data in a rapidly growing AdTech company.

Solution Confluent Cloud for a powerful, flexible, and simplified event-driven architecture that seamlessly manages massive volumes of real-time ad data while saving money.

Results

- 18x faster data updates for ad bidding
- 20-30% infrastructure cost savings
- Better customer experience with a scalable and fault-tolerant data platform
- Optimized ad operations through better ad budget control

Confluent Features Used

- Confluent Cloud
- Connectors (S3 sink)
- Schema Registry

Founded in 2012, Buzzvil offers a platform that helps companies optimize their ad spend through automation, behavioral analytics, audience targeting, rewards programs, and more.

The SoftBank-backed company has more than 300 global publishers, including major South Korean mobile carriers and membership operators such as OK Cashbag, L.Point, CJ ONE, Happy Point, Hana Members, and Liiv Mate.

Prior to adopting Confluent, Buzzvil relied primarily on Redis and RabbitMQ to store and move data to other systems and fuel its platform. They started by self-managing Redis and RabbitMQ but as Buzzvil's customer base grew, the company began to encounter challenges with achieving their desired scalability, and preventing data loss while scaling.

After switching to Confluent's fully managed, cloud-native data streaming solution, Buzzvil was able to fully capitalize on its data in motion and establish a powerful, flexible foundation for event streaming that leverages real-time data.



The Challenge

Move Away from Batch Processing and Self-Managing Kafka

Via its rewards-based advertising platform "BuzzAd," Buzzvil helps companies run ads that reward mileage points to customers through widgets, in-apps, and lock screens. It also helps companies achieve the best possible ad experience and results, encouraging end users to interact with the ads.

As Buzzvil continued to expand its primary service offerings, the company began to understand the importance of having an event-driven architecture, since everything it was doing—from ad reports, to publishers, to performance optimization, to budget controls—required real-time data processing. Buzzvil's primary need was to build a single source of truth for all real-time ad transactions and move away from batch processing.

Buzzvil was using Redis to store data about their customer interactions with online ads. Buzzvil found that Redis couldn't handle massive volumes of data, leading to time gaps between multiple datasets.

"In the past, there were multiple complaints from the clients as there were no immediate rewards after running the ads," said Jueun Seo, Chief Technology Officer at Buzzvil. "We also found that implementing customer support service on top of user interaction events resulted in tight coupling between the bidding engine and the customer support service, so we started to look at adopting an event-driven architecture to allow them to run independently and achieve greater flexibility."

Buzzvil was also finding it hard to adapt to change and support a growing customer base using RabbitMQ. As the amount of data generated by customers continued to increase, the need for event streaming and real-time data processing became clear.

As Buzzvil began to look around for a data-streaming solution, Apache Kafka® immediately appeared on their radar.

"There were a growing number of people on the data science team expressing interest in Kafka, and we also wanted to learn why tech communities are eyeing Kafka," Seo said.

"Initially, it was a no-brainer to choose RabbitMQ and open source Apache Kafka" Seo shared. "However, the outcome after choosing RabbitMQ was different from our expectations. We faced issues with managing it because the SDK only offered basic features, and our team had to spend extra hours to build add-ons or more features like retry, batch, and offset management."

The programming language difference was another roadblock. Python and Go are Buzzvil's primary languages, but the official RabbitMQ client library is only available with Java.

Buzzvil also found itself combating data loss issues that began to occur during large traffic spikes that RabbitMQ and Redis were unable to effectively manage.

All of the above eventually led Buzzvil to look at a real-time data streaming platform to set their data in motion.

"Startups in the early stage tend to choose only a single programming language and simple infrastructure with monolith architecture," Seo said. "But things change when the company starts to expand and there's a growing number of people, teams, and business requirements, with multiple teams working on independent projects and moving fast. At that point, you really need a flexible IT architecture because go-to-market velocity becomes crucial."

"We believe Confluent is a customer-centric company. While we were doing the proof of concept with the sales team, we were pleased with how Confluent engaged with us and shared solutions to fix problems. The team constantly asks about our feedback and improves product offerings. The fact that the co-founders of Confluent created Apache Kafka gives us confidence in Confluent's credibility, encouraging us to learn about technology and innovate our offerings to the clients."

— JUEUN SEO, CHIEF TECHNOLOGY OFFICER AT BUZZVIL

Technical Solution Confluent Cloud for a Single Source of Truth

In the world of online advertising, being able to manage, process, and use real-time data is key for ad targeting because it allows platforms to recommend the best possible ads to the right clients at just the right time.

After a series of meetings, proofs of concept, and trials, Buzzvil was impressed with Confluent's fully managed, cloud-native data streaming service that completely offloads the operational management of Kafka for enterprise use cases. For Buzzvil, two of Confluent's most compelling characteristics were its ease of use and the platform's late stage of maturity, which convinced them that Confluent would be a strategic partner for them in the long run.

Buzzvil was also impressed by the Python Library support that Confluent provides. This allowed them to quickly pull together and successfully test the proof of concept to validate that the integrations between their existing Python libraries (which in turn provide access to upstream systems) and Confluent worked seamlessly.

The Confluent Cloud sink connector to S3 is a powerful feature that helps Buzzvil to pull the data from other upstream systems and save it to Amazon Simple Storage Service (Amazon S3). Schema Registry, another powerful feature, eliminates inter-service bottlenecks and dependencies so developers can seamlessly connect to any data system using other connectors, while maintaining schema compatibility, version control, and quality assurance.

With Confluent, Buzzvil not only had a single source of truth for all real-time ad transactions but a platform that could help them achieve massive scale without the operational overhead.



The Results

Speed, Cost Savings, Improved Customer Satisfaction

18x faster ad bidding data updates.

"It used to take up to three minutes when the ad bidding engine refreshed or loaded the most up-to-date data. Now it takes less than 10 seconds." — Jueun Seo, Chief Technology Officer at Buzzvil

20-30% infrastructure cost savings.

"We're expecting to increase our bottom-line revenue growth via data infrastructure cost savings of 20 to 30 percent, and we expect to greatly reduce the complexity of our data infrastructure thanks to Confluent." — Jueun Seo, Chief Technology Officer at Buzzvil

Better customer experience via higher fault tolerance and reduced errors.

"We're expecting to enhance customer satisfaction as we reinforce the fault tolerance and reduce the amount of error occurrences." — Jueun Seo, Chief Technology Officer at Buzzvil



What's Next?

Buzzvil just completed their first Confluent deployment and their short-term plan is to make Confluent accessible for internal services within its environment to support Buzzvil's product offerings.

"We want to see if other teams would be interested in data streaming for business purposes," Seo said. "After taking baby steps, we want the benefits of real-time data streaming to be available to the entire organization."

Learn More About Buzzvil

<https://www.buzzvil.com/en/>