

SAS.COM

# SAS Powers Instant, Real-Time Omnichannel Marketing at Massive Scale with Confluent's Hybrid Capabilities



**Headquarters**

Cary, North Carolina, USA

**Industry**

Analytics Software & Solutions

**Challenge**

Scalability issues with self-supporting open source Kafka in-house were making it difficult to adjust to changing demand, and operational overhead and complexity were driving significant costs. Also, migrating the data to Kubernetes was a major hurdle.

**Solution**

The upgrade from Kafka to Confluent for Kubernetes solved scalability and improved performance while lowering the total cost of ownership (TCO). And Cluster Linking made the migration to Confluent for Kubernetes painless to unlock further agility.

**Results**

Confluent created radical improvement for SAS in three tangible areas:

- Ability to elastically scale on demand with improved performance
- Reduced TCO while ensuring reliable uptime and compliance
- Easy migration to Confluent for Kubernetes with Cluster Linking

SAS has been around for over four decades and has established a strong pedigree in marketing analytics, assisting more than 80,000 businesses in 150 countries in creating better downstream experiences for their own customers.

The vision of SAS is to transform a world of data into a world of intelligence by enabling decisions to be made at just the right moment. The key phrase here is *at just the right moment*—which is why SAS depends on a reliable, scalable data streaming foundation.

For example, with their popular SAS® Customer Intelligence 360 solution, SAS powers omnichannel marketing based on customer intelligence and ultimately makes large marketing campaigns more targeted, personal, and relevant at the right time. To optimally serve customers across industries from healthcare to smart cars to financial services, SAS must handle data in a way that's reliable and scalable—in real time.

Confluent acts as the central hub to ingest all customer data events for SAS Customer Intelligence 360 solution, powering a suite of state-of-the-art marketing software solutions. Here's the story of how the company's IT team unlocked the full value from Apache Kafka® with Confluent.

## Reliable scalability without performance degradation and skyrocketing costs

Some very big companies depend upon SAS to manage huge marketing campaigns that create a surge of data events. Being hyper-responsive to these customers means being able to scale campaigns on demand, despite a sometimes unpredictable ebb and flow of data traffic. As Justin Dempsey, senior manager, SAS Cloud, puts it, "Accommodating the sheer volume was challenging and negatively impacted the turnaround time on processing incoming events."

Although the IT organization at SAS was using Kafka, they were mitigating operating systems and software solutions that had reached an end-of-life state. Inability to scale—and its close cousin, degraded performance—were becoming problematic with managing open source Kafka on their own, and that was jeopardizing customer trust. Sherril Barnhill, manager, core services at SAS Institute, underlines the issue bluntly: "Historically, we were running on open source and were processing hundreds of thousands of events per second in each Kafka cluster with millions of events processed per second globally. Addressing performance anomalies and adding scale on demand for seasonal variation were challenges that we encountered."

*"With the containerized approach using Confluent, we now had ongoing vendor support and were able to meet on-demand scale with no disruption to our customers. The tools and dashboards provided by Confluent also helped us to discern performance anomalies in real time."*

— JUSTIN DEMPSEY, SENIOR MANAGER, SAS CLOUD

Kafka was a good initial fit for enabling a real-time data platform, but open source Kafka on its own wasn't scalable and reliable enough for SAS to use as part of a modern marketing automation solution. In addition, it lacked visibility and ease of use, and failed to provide a means to seamlessly migrate workloads across clusters and environments. SAS had run into limitations because of the sheer volume of data being ingested. The cost of ownership of Kafka was ballooning, and SAS needed a sleeker solution.

## Going beyond the limitations of Kafka

Confluent beckoned because of its ability to consistently and quickly scale both compute and storage on demand, while getting visibility into performance metrics benchmarks. The completeness of Confluent's data streaming platform, including enterprise-grade features many organizations consider to be prerequisites when

deploying mission-critical apps to production, provided further differentiation. "There were some specific factors that Confluent offered that made [competitors] look less appealing," Dempsey says, "specifically, the Role-based Access Control using group-based ACLs provided as part of Confluent's offering and the maturity of the Kubernetes operator, as well as the Cluster Linking component."

"Confluent is how we ingest all of our data events," says Barnhill. "They go through Kafka and are transformed from one microservice to another as they become a part of our ecosystem." Confluent enables SAS to go beyond what they could achieve with open source Kafka alone to deliver more targeted, personalized, and immediate marketing as part of their solution than ever before. It also ultimately saved the company significant time and money. Infrastructure cost savings alone totaled to hundreds of thousands of dollars per year.

## The Results

### The downstream effect of data in motion

#### The ability to scale on demand, with optimal performance

*"The value-add was to get to a point where we would be able to consistently and quickly add scale on demand for both compute and storage, and also provide visibility into performance metrics benchmarks."*

— JUSTIN DEMPSEY, SENIOR MANAGER, SAS CLOUD

Capturing and acting on customer interactions is the power of real-time data streaming for SAS customers. It's what allows them to push a digital coupon to a customer in the moment that customer is contemplating a purchase. If it's not happening at the right moment, if performance degrades, if there's a huge resource constraint happening—all of these things ultimately affect the downstream experience. In addition, with specific compliance and SLA agreements with customers in place, SAS needs to ensure that campaigns will reliably scale, in real time, without any sacrifice in performance.

In the past, preventing these issues required resource-heavy predictive analysis that took months and months, increasing the company's costs and talent resources. The IT team turned to Confluent because it allows them to handle an enormous and fluctuating load of streaming, real-time data without any performance degradation, and automatically scale up or down, on demand, without having to do a lot of complex calculations. With Confluent as the central nervous system of a streaming data ecosystem, SAS can guarantee real-time analytics and enable users of the platform to react to customers immediately and continuously—without risk of failure. Even as customer campaigns fluctuate wildly, scalability is consistent, and uptime and compliance are assured.

#### Significant TCO reduction

*"In terms of hard costs, Confluent reduced our bottom line from the infrastructure cost perspective, reducing licensing and physical hardware. We reduced our projected yearly cost by 69%. And the months of planning we would have to put into meetings and seasonal projections are gone, so our management oversight is reduced from a soft-cost perspective."*

— JUSTIN DEMPSEY, SENIOR MANAGER, SAS CLOUD

Improving the way data is processed and used company-wide has a profound impact on resource use. There are both direct and indirect ways that Confluent reduces the total cost of ownership (TCO).

Open source Kafka is free to use, but it often costs more to self-support in-house in the long run. In terms of direct costs, it often requires significant resource investments, including additional engineering staff to run and operate the complex manual processes. And Kafka often requires the purchase of additional infrastructure to deploy and operate at scale for critical business use cases. With Confluent, because the IT team can more efficiently provision resources, SAS can instead spend more on innovation and strategic initiatives rather than managing the underlying infrastructure. Confluent also reduces costs associated with paying people to oversee patch schedules and maintenance windows.

Indirect costs can be just as deadly to the bottom line: downtime, Kafka outages, and the ensuing loss of customer trust. Downtime often means lost revenue, reduction in customer confidence, diversion of valuable resources, data loss, and more. Confluent's ability to reduce the risk of downtime dramatically alleviates these indirect cost hits as well.

## A much easier migration thanks to Cluster Linking

*"Cluster Linking played a critical role in making the shift to Confluent for Kubernetes and our more cloud-native solutions a success, with no downtime and minimal dependencies."*

— JUSTIN DEMPSEY, SENIOR MANAGER, SAS CLOUD

To enable real-time data for their Customer Intelligence 360 suite, SAS wanted to deploy Kafka on their private cloud environment based on AWS EKS. But deploying Kafka on Kubernetes is a complex, challenging operation. Dempsey says, "Our transformation to a cloud-native, agile SaaS solution required a large-scale migration from open source Apache Kafka, which was complex and challenging to self-support."

Cluster Linking, a solution to easily link Confluent clusters together to form a highly available, consistent, and real-time bridge across environments, was a critical aspect of the project

because it streamlined the anticipated migration from the old Kafka infrastructure to the new. During the cross-divisional effort between SAS Cloud, R&D, and other SAS teams, they developed an approach using Cluster Linking to enable them to have effectively zero downtime. Dempsey's team was able to leverage Cluster Linking to seamlessly migrate existing clusters, enabling the team to build a private cloud Kafka service on their Kubernetes distribution of choice in order to properly power their customer 360 use cases.

By deploying on Confluent for Kubernetes, SAS was able to realize some of the benefits of cloud—scalability, ease of operations, and global resilience—while still on their private infrastructure. This enabled a smooth migration of Kafka clusters with minimal downtime and business disruption.

## What's Next for SAS?

With Confluent as the backbone of their next-gen Customer Intelligence 360 platform, SAS has been able to more effectively provide visibility for its customers to find insight into *their* customer data—which products they're searching for, whether they're opening emails, when they click on links, when they use a catalog code or purchase from a physical store—the moment they happen.

Also, one of the biggest wins has been the ability to radically improve customer experience with an ad hoc testing environment that enables rapid innovation. SAS will continue to create radically intuitive marketing experiences for downstream customers and end consumers with R&D efforts supported by Confluent, which is key to enabling the test environments and space in which they can continue to experiment and innovate.

Confluent will also help support and accelerate their journey to cloud, with the ability to expand from AWS into Azure. With Confluent's Oracle CDC connector, SAS will be able to reliably and cost-effectively offload event data from Oracle Database to Confluent. And later this year, Dempsey's team plans to explore Tiered Storage to further optimize costs, with additional focuses on security, auditing, and governance controls coming too.

## Learn More About SAS

<https://sas.com>

## Get Started with Confluent for Free