

DKV MOBILITY

DKV Mobility Accelerates Innovation with Confluent Cloud


Headquarters

Ratingen, Germany

Industry

Mobility

Challenge

Streamlining management of data streaming pipelines to scale development of real-time digital products to deliver a more compelling, competitive experience for B2B mobility customers.

Solution

Using Confluent Cloud, a fully managed cloud-native solution with pre-built connectors, to overcome operational roadblocks and accelerate real-time innovation.

Results

- Enabled real-time payments and transaction processing for on-the-road insights
- Accelerated development with self-service cluster management capabilities
- Eliminated costly downtime during a six-month-long cloud migration initiative

Founded in 1934, DKV Mobility has a long history of providing its customers with mobility services and processes that are fundamental to their everyday routines. Whether businesses need energy, toll or mobility solutions or VAT refund, DKV Mobility continually transforms its platform business to meet the demands facing customers today and their challenges tomorrow.

DKV Mobility currently provides approximately 301,000 active customers, across more than 50 service countries, with intelligent solutions, reliably helping businesses efficiently and cost-effectively stay mobile. As the leading European B2B platform for on-the-road payments and solutions, DKV Mobility is renowned for putting the customer first. To continue this legacy, the company's Customer Product Services (CPS) department began leveraging real-time data to enhance their digital, customer-facing products.

Initially, DKV Mobility built real-time data streams with Apache Kafka®, which allowed the company to create some of the responsive digital experiences their customers needed. Over time, however, the CPS department's progress slowed as Kafka became a bottleneck instead of an advantage. To keep pace with the evolution of smart, sustainable mobility and execute the company's long-term growth strategy, DKV Mobility turned to Confluent for a state-of-the-art way forward.



The Challenge: Stalled Transformation with Kafka

To stay competitive amid the changing mobility landscape, DKV Mobility needed to improve the technical foundation of its customer-facing digital products. The company's CPS department realized that batch processing could not provide the real-time insights their customers expected from on-the-road mobility services.

Rising Customer Demand for Real-Time Applications

While the platform also supports other services, DKV Mobility is best known for providing large-scale logistics companies with a variety of fuel cards. DKV Mobility cards handle refueling, EV charging, toll solutions, and vehicle services across approximately 468,000 EV charge points, 63,000 fuel service stations, and 30,000 vehicle service stations.¹ These cards provide a convenient, scalable way to monitor and manage on-the-road service costs for small or large distributed fleets.

In recent years, customers have shown increasing demand for real-time capabilities for these types of mobility solutions. For example, businesses operating massive fleets of cars and trucks depend on having reliable, on-the-road access to the lowest cost fueling options available. Not only can employees use these cards to pay when refueling company cars and trucks, but they also have access to a portal and mobile app that shows them the nearby fuel prices.

Previously, DKV Mobility's digital products relied exclusively on batch processing through point-to-point, extract-transform-load (ETL) pipelines to provision and process data and then provide updated results to end users or systems. As a result, customers would be provided with outdated data when comparing fuel prices. Additionally, the DKV Mobility portal could take up to three weeks to ingest transaction data via a legacy enterprise resource planning (ERP)

system and then send customers a physical invoice via mail.

By adopting Kafka and building real-time data pipelines, DKV Mobility was able to transform the customer experience. Instead of having to wait hours to see updated fuel prices or weeks to see transaction data, DKV Mobility customers could access those insights in minutes or even seconds.

Stalled Progress: Operational Challenges with Apache Kafka

While batch processing still remained viable in other use cases, many of DKV Mobility's customer-facing digital products needed to incorporate data streaming and stream processing to deliver the real-time insights needed.

Within the CPS business unit, DKV Mobility's platform team is responsible for providing product teams with the scalable cloud

platform needed to develop high-quality digital services for DKV Mobility customers. With real-time data streams, managed by the platform team, DKV Mobility product teams could propagate updates across multiple channels in near-real time and deliver new features like contactless payment at DKV Mobility service points and power transfer point identification for EVs.

These kinds of real-time use cases became essential to DKV Mobility's continued success and approach to customer centricity. Adopting Kafka brought undeniable value to the business, but the open source data streaming platform soon became a productivity bottleneck for the CPS business unit.

The challenges of operating self-managed Kafka clusters occupied the platform team's time, taking away resources from other, high-value projects and lengthening time to value for new products and features.

The Solution: Back on the Road to Transformation with Confluent Cloud

Deploying and managing Kafka in production introduced significant complexity and toil to the CPS platform team's workload.

The CPS business unit had launched a new, microservices-based installation of its customer portal running on Kubernetes. To run Kafka on Kubernetes, the platform team used on-premises instances of OpenShift and Strimi, a combination that was difficult to manage. As a result, the environment within which the platform team had to manage Kafka was not stable, resulting in recurring, costly service outages.

At the same time, DKV Mobility soon began a year-long cloud migration initiative—migrating from private cloud to Microsoft Azure—as part of the company's long-term IT strategy. This ongoing migration increased and complicated the demands on the platform

team. DKV Mobility needed a way to ensure the uninterrupted performance of their real-time applications while also paving the way for the development of new real-time applications, features, and capabilities.

With an imminent product launch on the horizon, in 2019, the DKV Mobility platform team turned to Confluent to take advantage of the company's:

- Deep knowledge of [Apache Kafka and cloud operations](#)
- [Training and professional services](#)
- Pre-built [integration capabilities](#)
- Fully managed, cloud-native Kafka service, [Confluent Cloud](#)

"Confluent Cloud has become the technological success factor for DKV Mobility to drive the transition towards an efficient and sustainable future of mobility."

—TOBIAS GOCKEL, PLATFORM TEAM MANAGER, DKV MOBILITY

Why Confluent Cloud

To support product teams in delivering real-time experiences to customers, the platform team needed a trusted partner to manage the operation of DKV Mobility's Kafka infrastructure. But with Kafka already supporting critical customer-facing applications in production, the company also needed a solution that software teams could transition to without interrupting service availability for end users.

As DKV Mobility migrated from Apache Kafka to Confluent Cloud, Confluent's customer success teams provided critical support to ensure a seamless transition. Throughout the rollout of Confluent Cloud to production environments, Confluent consultants held numerous architecture reviews to help the platform team solve challenges unique to its microservices architecture.

While Confluent Cloud eliminated the operational burden on the platform team in the CPS business unit, [Confluent Connectors](#)—including the JDBC Source and Sink Connectors, the AzureDataLake2 Sink Connector, and the HTTP Sink Connector—and Confluent Replicator simplified the exchange of data between departments across the organization.

The Results: Accelerating Development and Minimizing Toil

Collaborating with Confluent throughout the migration process gave the platform team confidence that the CPS business unit would continue to make progress toward its strategic IT goals—keeping cloud costs low, solving long-term architectural challenges, and dedicating more time to developer enablement.

Eliminating Operational Roadblocks and Bottlenecks

Before adopting Confluent Cloud, platform team members spent close to 10% of their time (around eight hours per week) managing Kafka clusters for three product teams. Since that time, the number of teams relying on data streams in production has grown to eight—a number that could have easily doubled the amount of time the platform team spent managing Kafka.

But with Confluent Cloud, the platform team can now provide all eight product teams with the clusters they need in just minutes.

After migrating to Confluent Cloud, DKV Mobility spent six months migrating its cloud deployments from private cloud to Microsoft Azure. During that period, Confluent clusters in production saw zero hours of downtime, which meant an uninterrupted digital experience for DKV Mobility customers and less work occupying the platform team's time and attention.

Additionally, the cloud platform's Schema Registry capabilities allow the platform team to easily manage schema across clusters and teams, which they have found particularly advantageous when helping CPS software teams handle data migration.

Confluent Cloud's role-based access control (RBAC) and Access Control Lists (ACLs) streamline administration and security on the service and application level. Over time, continuous integration and continuous deployment (CI/CD) have become increasingly essential to DKV Mobility's product strategy. Together with RBAC and ACLs, Schema Registry allows the platform team to effectively apply its CI/CD approach to evolve the underlying data streaming infrastructure as needed.

“Choosing Confluent Cloud has allowed DKV Mobility's platform team to create a self-service tool for internal stakeholders developing customer-facing products. As a result, platform team members can spend less time managing clusters and more time working on higher-value projects critical to the company's success.”

—TOBIAS GOCKEL, PLATFORM TEAM MANAGER, DKV MOBILITY

Accelerating Product Development and Enabling Self-Service

When migrating from Kafka to Confluent Cloud DKV Mobility's CPS business unit had already been working toward using microservices to decouple decentralized product teams as part of their larger CI/CD approach. Alongside the other changes to the company's technology infrastructure, Confluent Cloud helped increase deployment frequency.

CPS product teams went from deploying product changes once a month to multiple times a week—sometimes within a single day.

The platform team also strives to increase internal self-service management among software teams. Adopting Confluent Cloud supported this effort as well, allowing developers to request the creation or deletion of a topic, cluster, or other resource. Previously, this took one day to process manually—now the automated, self-service process is completed in minutes.

Over the three years since DKV Mobility adopted Confluent Cloud, the platform team has been able to create additional services and solutions to accelerate developer productivity—including self-service data streaming capabilities.

Enabling Real-Time Payments That Let DKV Mobility “Lead in Green”

With Kafka Streams, DKV Mobility developers can take advantage of stream processing without having to worry about processing failures. This allows DKV Mobility developers to avoid the burden of traditional event-driven system design while still being able to process data flows in real time.

As a result, DKV Mobility customers benefit from more responsive and reliable digital products that help them manage their vehicle fleets more efficiently and sustainably.

Examples include:

- Fleet managers who use the DKV Mobility platform have real-time insights into the charging transactions among approximately 468,000 charge points for electromobility.
- Annually, billions of transactions from third-party platforms are integrated into the DKV Mobility platform from various separate third-party systems. With Confluent Cloud managing these integrated data streams, DKV Mobility customers receive payment authorizations and view new transactions on their accounts in minutes.

The Future: Cross-Functional Integration and Data Reusability

DKV Mobility continues to explore new opportunities to leverage real-time data and improve the customer experience. From operationalizing real-time insights from payment transactions or platform telemetry, the CPS business unit plans to use Confluent Cloud to continue innovating as the company onboards more electric charging and alternative fueling stations to their portfolio.

Departments and teams across DKV Mobility are increasingly interested in how Confluent Cloud supports cross-departmental data integration. Both the platform team and product teams are taking advantage of Confluent's self-paced training subscriptions to increase their technical knowledge of data streaming, allowing the company to further anchor Confluent Cloud as a central technology for building data streaming pipelines.

The increased use of Confluent Connectors continues to reduce the time and cost required for data integration across the organization. As a result, DKV Mobility is investigating how they can use [ksqlDB](#) to reduce development efforts, standardize data processing, and make data more accessible and reusable across functions.

Explore the Benefits of Confluent Cloud

<https://www.confluent.io/confluent-cloud/>

Learn More About DKV Mobility

<https://www.dkv-mobility.com/en/>