Revving Up IoT Use Cases with Real-Time Hybrid Cloud Data Architecture

Headquarters
Munich, Germany

Industry
Automotive

Challenge
BMW Group needed to make all data generated by its 30+ production facilities and worldwide sales network available in real time to anyone across the global business.

Solution
Using Confluent, BMW Group makes Apache Kafka deployments enterprise-ready, creating a central hub for real-time data streaming and data flows across the global organization.

Results

- Removal of manual provisioning and data infrastructure management through elastic scaling with Confluent's fully managed service

- Ubiquitous access to real-time data

- A reliable, zero downtime data streaming platform that connects edge IoT data from the production plants to cloud applications and services

- Full integration with SAP solutions, industrial IoT protocols like OPC UA, and cloud applications and services

Founded as Bayerische Motoren Werke GmbH in 1916, today, the BMW Group is the world's leading premium manufacturer of automobiles and motorcycles. The Group operates 31 plants in 15 countries, with a global sales network in 140 countries.

Inside BMW Group's state-of-the-art plants, data is collected from all the critical minutiae of manufacturing, including machine temperatures, machine performance, and location of parts on the assembly line. Data is also continuously produced and consumed by sales and the research and development (R&D) operation.

Making the abundance of real-time Internet of Things (IoT) data available to all teams that want to put it to use is a critical role for BMW Group's Integration Platforms organization to drive innovation through use cases such as enabling operational applications that streamline factory production and providing real-time data to the R&D team as they create and test new products.

BMW Group needed scalable, highly reliable, real-time data streaming and turned to open source Apache Kafka®, yet it lacked the scalability, reliability, and simplicity to deliver ubiquitous, real-time data for its vast, global operation. While Kafka allowed BMW Group to solve specific data use cases, its self-managing nature meant costly and specialized maintenance.

Confluent is working with BMW Group to deliver a data streaming platform that would accelerate the internal use of Kafka and support continuous innovation and new use cases by delivering easy access to data as real-time streams.

The Challenge
Removing the operational burden of data streaming

To handle an abundance of data, BMW Group initially replaced its old batch model with Kafka. However, while open source Kafka allowed BMW Group to add more applications without having to tack on additional clusters within their data streaming platform, it proved to be an operational burden.

Kafka was costly and not reliable enough to meet the demands of a global business, where applications must run 7x24x365. BMW Group needed a way to seamlessly integrate all its data—generated by machines, sensors, and other sources—and make it accessible in real time as a self-service product across the business.

BMW Group recognized the importance of shifting to a central, cloud-based data system to enable teams across the organization to leverage real-time data at scale, and build a future-proof data streaming architecture that would support innovative use cases.

By partnering with Confluent and Microsoft Azure, BMW Group harnesses their data streaming capabilities and offers these to hundreds of different teams and applications across the organization.

The Solution
Data in motion with Confluent Cloud

Initially, Kafka was deployed in BMW Group's on-premises cloud platform, Red Hat OpenShift, to provide a solution building block so application teams could "help themselves"—running their own Kafka clusters and taking advantage of ZooKeeper and other components of the Kafka ecosystem. This expanded to operating Kafka as a centrally managed service, with clusters operated by a central technology team.

Confluent provides an easily accessible platform that enables the effortless exchange of data between all applications and systems. The organization can directly ingest sensor data from the shop floor in the production plants into Confluent as a central hub in the cloud. That data is then bridged into other Azure cloud services for consumption by different teams and applications using Confluent's data integration capabilities.

Data flows continuously through Confluent via a combination of commercial solutions and custom code, using streaming ETL—the process of continuously transforming, filtering, and enriching data as it moves. As Confluent's Infinite Storage persists the incoming events, BMW Group can also replay historical data later to enable additional use cases. Any asset that produces data is connected to Kafka, with multiple different application types using the same data to trigger different, specific processes.

The migration from self-managed Kafka to Confluent Cloud allowed BMW Group to shed the operational burden and cost associated with running Kafka. This move enabled the organization to mitigate the risk of only having a small group of internal Kafka specialists within the IT infrastructure team, delivering a fully managed solution where 7x24x365 Kafka expert support is included, even in business-critical use cases.

BMW Group now has douzends of clusters, hundreds of connected applications and over one billion events produced per day—all with Confluent functioning as the central nervous system for this continuous river of critical business data. Highly scalable teams at BMW Group are onboarding multiple new applications and use cases every week.

The Results
The global data streaming standard

After putting its data into motion using Confluent Cloud, BMW Group is seeing considerable benefits:

A fast track to innovation
With Confluent, BMW Group enables continuous IoT innovation through the ability to process the immense amount of data generated by tens of thousands of sensors in its plants. With more and more machines, robots, and assembly lines pushing out more and more data, having a data processing platform that can swiftly distribute well-processed data to any internal consumer enables BMW Group to grow its IoT use cases, as well as after-sales, marketing, and customer engagement.

As a worldwide leader in the design and manufacture of high-performance vehicles, the highly scalable nature of Confluent enables BMW Group to accelerate project timelines and increase opportunities for innovation.

Removal of manual provisioning and management
From an operational perspective, the centralizing and decoupling services delivered by Confluent makes life easier for BMW Group's technical team by eliminating multiple interfaces and making the IT landscape significantly more manageable.

A lot of resources were needed to maintain the infrastructure and avoid downtime for Kafka, but switching to Confluent keeps BMW Group's data streaming constantly, without the burden of a self-managed service.

Confluent removed manual provisioning and data infrastructure management through elastic scaling with Confluent's fully managed service, Confluent Cloud, which offers 7x24x365 support against strict service level agreements (SLAs). As a result, BMW Group benefits from a reliable, zero downtime data streaming platform that operates more efficiently and mitigates infrastructure risk.

Ubiquitous access to real-time data
Confluent gives BMW Group a powerful data streaming platform that expands the benefits of Kafka and enables the business to easily and continuously access, store, and manage data. DevOps teams no longer need to deploy Kafka to benefit from its data streaming capabilities.

By moving data processing to the cloud, master data is immediately made available centrally through Confluent, and as many applications as possible can benefit from that architecture. Teams no longer have to manually set up a one-to-one interface contract between the central master data system and the consuming application in order to start innovating.

Real-time business insight

From a business management perspective, Confluent brings better insight into data across business units and use cases, which drives informed decision making.

Teams can instantly access data in the analytics hub to find correlations, identify opportunities, and develop solutions that are relevant for their business department, for example, ordering parts, reviewing customer sales, or monitoring the R&D process. The result is fast decision making that accelerates the business, driven by access to real-time information.

Full business integration

Confluent makes it easy for BMW Group to connect its existing apps and data systems, which operate both on premises and in the cloud. Full integration with SAP is particularly important, alongside new cloud applications and services to support BMW Group's cloud strategy.

BMW Group operates a complex infrastructure, with a large number of SAP systems, a lot of legacy software, plus third-party developed connectors and tools. Confluent enables the organization to integrate all aspects of its data architecture into one platform, marking a significant move forward in its evolution as a data-driven business.

The Future
Innovation and business acceleration

BMW Group's successful implementation of Confluent has already resulted in increased speed, flexibility, cost savings, and simplicity. Yet the company is ambitious and keen to expand the use of the data processing platform.

Confluent has an important role in the BMW Group's data strategy, acting as the central nervous system that integrates with different data platforms and sets the standard for an event-driven architecture. Through the cloud-based approach, the organization is ideally placed to expand the Confluent platform across the global BMW Group to connect all plants and locations.

Looking ahead, the ultimate vision is a fully integrated, transparent, and centralized approach to the way data is exchanged, accessed, and analyzed.

Learn More About BMW Group

www.bmwgroup.com/en.html