



# Apache Kafka in the Automotive Industry

# Apache Kafka in the Automotive Industry

In November 2019, I had the pleasure to visit "Motor City" Detroit. I met with several automotive companies, suppliers, startups and cloud providers to discuss use cases and architectures around Apache Kafka. I work with companies related to the German automotive industry for many years. It was great to see the ideas and current status of projects running overseas in the US.


I am really excited about the role of Apache Kafka and its ecosystem in the automotive industry. Kafka became the central nervous system of many applications in various different areas related to automotive industry. Machine Learning also got more and more impact on these use cases.

## A Long Journey – From Car Production and Sales to Digital Services



My trip to Detroit started with some private events. First, I watched the amazing football college game of the Michigan Wolverines hosting the Michigan State Spartans in front of 111000 fans in Ann Arbor. The next day I went to the Ford Piquette Plant in Detroit where Henry Ford and his team manufactured the first cars. The Henry Ford Museum was another great visit to learn more about the mastermind Henry Ford and the history of car manufacturing in the US.

## The Car Industry is Changing...



I talked to many local people. The main excitement was around three topics in the state of Michigan: Universities, sports and cars. I have to admit that talking about car companies created a much more frustrating and negative mood for many people. They work in the factories and offices of car companies. They were more happy about discussions around sports and universities. Many people are worried about the future of the automotive industry.

Thus, let's think in more detail about the car business... A car company produces and sells cars. A supplier produces parts for the car. A vendor shop or independent company provides maintenance and repair services.

That was the situation for many decades. Welcome to a new era where car companies have to do much more than just manufacturing and selling cars to stay competitive.

I think people should not be worried, though. This is very exciting. While many old jobs will be done by machines and robots in the future, new tasks and job roles will be created for everybody.

## Customer Behaviour has Changed...

### Customer behaviour has changed significantly:

- Customers expect digital services and integration with their own IT ecosystem (smartphone, music streaming service, smart home, and many other apps)
- Customers look for cheap repair services away from the car vendors because it is much cheaper but still a good service
- Repair shops use 3rd party suppliers for car parts because the original supplier is too expensive and does not add any value or better product quality
- Tesla, Uber, Volkswagen, Daimler and many other automotive and non-automotive companies work on autonomous driving (in several steps and releases: Automatic distance measurement on the highway → Automatic parking in the parking lot → Driving back to the entrance of the supermarket when you call your car via your smartphone → Real autonomous driving to the final destination)
- Zipcar, car2go, drivenow and other companies provide car sharing services where car usage is paid by minutes and miles: You use your smartphone to locate the next car at the airport, drive to your home in the city, and leave the car in front of your house so that the next customer can pick it up
- Uber, Lyft, Grab, Free Now (the most ridiculous company name you could have for ride sharing) and other services provide cheap ride sharing with fantastic user experience: You order a taxi via mobile app, select guest services (like choice of music, wish of conversations with the driver or not), see expected estimated time of arrival and expected cost, monitor the status of the taxi location in real time, get transported, pay automatically by the integrated payment service, collect points at the integrated and partnered loyalty systems (like Hilton in case of Lyft or MilesAndMore like Free Now to name a few examples)
- In short, companies in the automotive industry have to change their business model fundamentally to stay competitive. Otherwise, they will become the next "hardware supplier" for a tech giant. This would reduce profit, establish tough dependencies and create market pressure. Even the existence of the company is at risk. No matter if you are on the market for 10 or 100 years, already.

### Digital Use Cases in Automotive Industry

The background about the history and new requirements for the automotive industry is important to understand. Every car-related company can add more value and innovation to the own business model.

The use cases can be separated in infrastructure, manufacturing and customer interaction (while there are some overlaps, of course).

## Infrastructure for Connected Cars

A fully connected infrastructure with real time communication is required to build innovative use cases such as:

**Connected cars:** The "Hello World" example for discussing cutting edge use cases in automotive industry. It includes more or less all challenges you have to solve in Internet of Things (IoT) scenarios, including connectivity to millions of devices, large scale throughput, reliable communication with bad and low network connectivity, real time processing requirements. Ingest and process sensor information from millions of cars in real time to correlate events and. Provide bi-directional communication between cars and other applications. Do big data analytics in the backend to get new insights. Build additional services to the manufacturer, customer and partners.

**Fleet Management:** A specific example of real time correlation of events from various different systems like mobile apps, hardware integrated into cars and trucks, and backend systems like CRMs and payment systems. Various examples exist in different industries, including logistics of truck fleets, track&trace of package delivery, ridesharing, food services, and many more.

**Emergency system:** If your car shows abnormal behaviour or crashed into a tree, the car automatically sends a crash notification (including details about the crash scenario) to the police and hospital next to you to send help immediately.

**Smart City and Smart Driving:** Connect and correlate data of cars with other devices like traffic lights. Automotive companies partner with cities and other data providers to offer better security and more comfortable driving experience. Collect and share basic safety messages. Monitor cars and send speed / slow-down warnings, wrong-way detection, congestion. Provide crowdsourced data with applications such as Waze, Google Maps, Twitter, Nextdoor.

## Manufacturing and Industrial IoT (IIoT)

Producing cars is a very complex task. This includes the production of the car itself (the "hardware") and the intelligent logic (the "software") of the car. Both integrate deeply with each other to implement use cases such as:

**Industrial IoT (IIoT) and smart manufacturing:** Predictive maintenance, early part scrapping and improved product quality are some examples to improve the quality of the manufactured parts, reduce costs and risks. The analysis of car sensor data also allows to implement additional digital services on top of the plain car parts.

**Autonomous Driving:** One of the most interesting use cases. Real time processing of big data sets and using cutting edge technologies such as Neural Networks allows car companies to build car with autonomous driving features. This is still in early stage and will probably take many years before it is available in Europe. But some US states and China are pressing ahead with self-driving cars on public streets, already. This will take some more years in most places. In the meantime, you can already use less powerful but still impressive features like Tesla's "summon car" feature. This is not working perfectly today, but this is just a matter of time.

## Customer Interactions and Data Analytics

Connected infrastructures and self-driving cars are awesome. Another added value comes through the possibility to communicate directly with the drivers to provide additional services, for instance:

**Remote Diagnostics:** Analytics can happen at the edge (e.g. in the car) or in a data center / cloud. Autonomous driving needs to detect persons in real time. Thus, this needs to be applied in the car. Cross selling is an example where you correlate data between different backend systems and remote users or apps. Therefore, this should happen in the backend. You just send the result of the correlation back to the car or mobile app of the driver. Predictive maintenance

can happen directly in the car or machine. Or you correlate the events in the backend and just send the alerts to the car or machine to stop it before the engine breaks. No matter where you deploy the analytics logic, it typically has to happen in real time. Usually within milliseconds or seconds, but sometimes minutes or even hours are sufficient.

**Identity Management:** The days are over where you use a key (a piece of hardware) to open your car and start your engine. Car sharing services provide a mobile app. This is used for registration, payments and handing over the key (for the time of usage). Security (Authentication, authorization and encryption of data) are getting more and more important in IoT. This is even true in Industrial IoT where security discussions did not exist at all in the past. Car sharing is a great example where this is a key requirement to deploy an infrastructure and service successfully. The next step several companies are working on (long live China!) is image recognition instead of using a physical key, or smartphone to open your car and start the engine. Other use cases include improving the customer experience in the car. For instance, default configuration of seat, radio and air conditioning can be based on the driver's historical behaviour (stored forever in your favorite database) and the outside ecosystem (e.g. integration with a weather service to know how warm it is). You will probably also not need to carry your driver's license with you in the future because police will also just scan your face. I was pretty impressed when I just had to scan my face at Global Entry entering the US border coming from Germany last week, instead of scanning my passport and entering all the annoying details about my travel and private information.

**Aftersales:** Car companies and manufacturers need to sell additional services to make more money and / or to make customers happy and loyal. This includes use cases like upselling (provide 100 extra horsepower via digital download for 24 hours to have some fun on German Autobahn) and cross-selling (provide a 20% coupon for the partnered steak house coming at the next stop 10 miles always from your current driving location).

**Payment Integration:** In the US, it is common to pay everything by credit card. In Europe, many shops still just accept cash. In the future, payments will be done automatically. You just go to the gas station or electric charging unit and refuel respectively load your car and leave. Similar to existing an Uber or Lyft car today. The payment is done via the integrated partner payment system of your choice. Once again, loyalty systems are integrated, too. Authentication and authorization are done via cameras and neural networks doing image recognition.

**Data Monetization:** Connected car infrastructures produce a lot of data. Depending on law, privacy options and other regulations, automotive companies can and will do its own analytics, but also sell the data or parts of the data to partners (hopefully anonymized properly). Like Google or Facebook today, the automotive industry will be able to make a fortune with the data of the cars and all the connected partner systems. Most customers will probably agree to this "feature" because they get added value out of this, too (like discounts or "free" additional services).

Wow, this is a lot of use cases for digitalization of the automotive business, isn't it? And there are many more to come...

## Challenges and Requirements to Build Automotive Infrastructure at Scale

Let's now quickly discuss the challenges and requirements to solve with all those exciting use cases:

Integration between many different consumer apps and backend systems (including many legacy and proprietary interfaces, and different communication paradigms like real time streaming, request-response, batch processing)

Big Data – sensors and millions of mobile apps from users produce a large set of (structured and unstructured) data continuously

Hybrid integration and deployments at the edge (e.g. cars), local proxies in different regions, backend applications in data centers or cloud, and SaaS applications

Global scale with zero downtime, rolling out the applications on different infrastructures (there is no AWS / Azure / GCP in China but only Alibaba, and there is no cloud at all in Russia (just private clouds on premises).

Real time information is required in most scenarios to provide a good digital experience and added business value



This is a lot of challenges to solve. If you know Apache Kafka already well (and maybe also other traditional middleware), then you can probably guess why Kafka is such a good fit here.

### Apache Kafka in the Automotive Industry

After we understood the various innovative use cases and its related challenges in the automotive industry, let's now discuss (quickly) why so many automotive companies and suppliers use Apache Kafka as central nervous system for the discussed use cases.

### Kafka for Real Time Streaming at Scale

Car sensors produce data continuously. The more cars you connect, the more data you get. Kafka provides messaging and streaming at scale, providing a reliable, highly scalable infrastructure with zero downtime. Rolling upgrades, backwards compatibility, up- and downscale at runtime, and other features are built-in.

### Kafka for Handling Backpressure and Decoupling of Services

Cars are decoupled from the backed systems and mobile apps. Kafka provides storage for backpressure and decoupling of services. It is not just a messaging system! Kafka also stores data as long as you want. For instance, a Kafka topic is stored for a few days for log analytics. Another Kafka topic is stored for years to analyze customer and payment transactions of the past. Order guarantees, timestamps and high availability are provided out-of-the-box. Connected applications and data stores can build a materialized view to leverage the data.

### Kafka for Integration with Legacy and Modern Technologies

Car data needs to be integrated. Kafka is the backbone for backend systems (legacy like ERP, Mainframe + modern SaaS

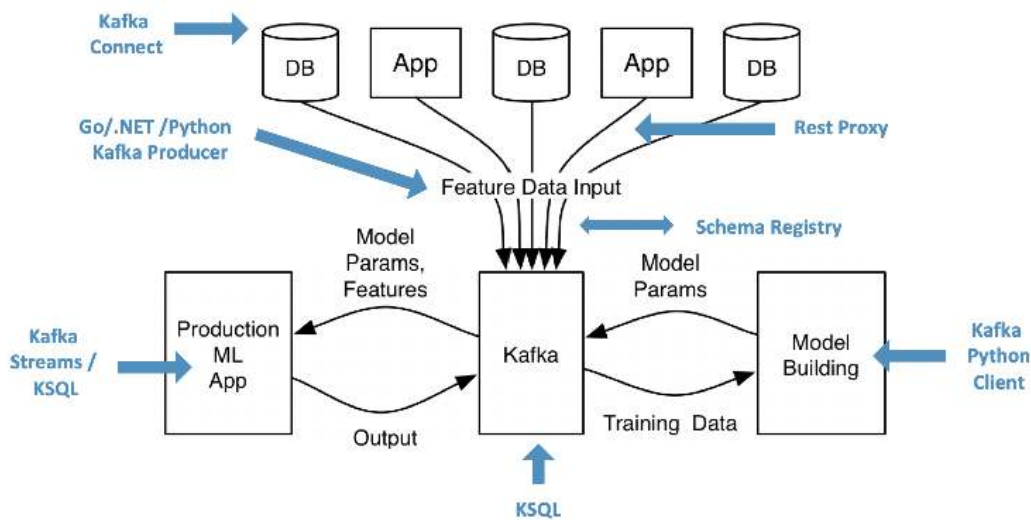
like Cloud CRM, big data analytics). Connectivity via Kafka Connect provides Kafka-native integration at scale to any legacy or modern technology and communication paradigm. Either directly from Kafka clients or via other technologies as proxy like MQTT to cars, Websockets to mobile apps, via proprietary protocols like Siemens S7 or open standards like OPC-UA to machines in plants). [Check out my comparison between Kafka and Middleware \(MQ, ETL, ESB\) for more details.](#)

### Kafka for Continuous Stream Processing with Kafka Streams and ksqldb

Car data needs to be processed and correlated in real time at scale with other backend databases and applications. Kafka Streams or ksqldb provide Kafka-native capabilities to do process data continuously, for use cases like Streaming ETL, stateful aggregations (sliding windows) and building business applications with their own state (leveraging rocksdb under the hood – no need for another database).

### Machine Learning and Kafka in Automotive Use Cases

Kafka is used more and more in Machine Learning infrastructures. Some examples from tech giants are Uber, Netflix or Paypal. The big challenge about Machine Learning is the deploy at scale in a reliable way (for both model training and predictions). I covered this topic in details in various blog posts, videos and demos. Check out this blog post to get started: ["How to Build and Deploy Scalable Machine Learning in Production with Apache Kafka"](#).



### Kafka for Data Engineering, Model Training + Deployment and Monitoring

In the Automotive Industry, Machine Learning needs to be applied at scale. Predictions often need to happen in real time. Therefore, more and more automotive-related use cases discussed in the above section leverage Kafka together with Machine Learning for different aspects like:

Data ingestion and processing of data with Kafka from cars and other systems for both model training of historical data and real time predictions on new events – ideally using the same kafka ingestion and pre-processing pipeline for both.

[Streaming model training with Kafka](#) without an additional database in the middle (e.g. leveraging TensorFlow I/O and its native Kafka integration) or the storage in a data lake (e.g. HDFS, AWS S3) for model training.

Model deployment with Kafka at the edge (e.g. in autonomous driving for stopping the car via image recognition because

a person is on the street in front of the car) or in the data center / cloud (e.g. predictions for cross selling by correlation of real time behaviour together with correlated information from the loyalty and CRM systems in backend). Check out the video and slides discussing ["Event-driven Model Serving: Stream Processing vs. RPC"](#) from Kafka Summit for more details.

Monitoring with Kafka (e.g. real time alerting, or distributed tracing to detect anomalies and the root cause of problems).

One huge advantage of using [Kafka for Machine Learning infrastructures is solving the impedance mismatch between the data scientist \(loving rapid prototyping with Python and Jupyter\) and the software developer / production engineer \(loving scalable and reliable Java applications\).](#)

### Kafka Architecture – From Edge Deployments to Global Replication

The architecture for Kafka deployments depends on the use cases, SLAs and many other requirements. Automotive companies and suppliers typically have to think globally:

Using Kafka as global nervous system for streaming data typically means you spin up different Kafka clusters. The



following scenarios are very common in the automotive industry:

Local edge Kafka clusters in the factories: Each factory has its own Kafka cluster to integrate with the machines, sensors and assembly lines. But also with ERP systems, SCADA monitoring tools, and mobile devices from the workers. This is typically a very small Kafka cluster with e.g. three Brokers (which still can process ~100+ MB/sec). Sometimes, just one single Kafka broker is deployed. This is fine if you do not need high availability and prefer low cost and very simple operations.

Central regional Kafka clusters: Kafka clusters are deployed in different regions. Each Kafka cluster is used to ingest, process and aggregate data from different factories in that region or from all cars within a region. These Kafka clusters are bigger than the local Kafka clusters as they need to integrate data from several edge Kafka clusters. The integration can be realized easily and reliable with Confluent Replicator or in the future maybe with MirrorMaker 2. Don't use Mirrormaker 1 at all – you can find many good reasons on the web). Another option is to directly integrate Kafka clients



deployed at the edge to a central regional Kafka cluster. Either with a [Kafka Client using Java, C, C++, Python, Go](#) or another programming language. Or using a proxy in the middle, like [Confluent REST Proxy](#), [Confluent MQTT Proxy](#), or any MQTT Broker outside the Kafka environment. Find out more details about [comparing different MQTT and HTTP-based IoT integration options for Kafka here](#).

Global Kafka clusters: You can deploy one Kafka cluster in each region or continent. Then replicate the data between each other (one- or bidirectional) in real time using Confluent Replicator. Or you can leverage the [multi-data center replication Kafka feature](#) from Confluent to spin up one logical cluster over different regions.

This is just a quick summary of deployment options for Kafka clusters at the edge, on premises or in the cloud. You typically combine different options to deploy a hybrid and global Kafka infrastructure. I will speak about different ["architecture patterns for distributed, hybrid and global Apache Kafka deployments"](#) at DevNexus in Atlanta in February 2020.

For example, you might leverage [Confluent Cloud, a fully managed Kafka service with usage-based pricing and enterprise-ready SLAs](#) in Europe and the US on Azure, AWS or GCP. But you need to use Alibaba cloud in China and used [self-managed Kafka there deployed to Kubernetes and operated by Confluent Operator](#). In Russia, no public cloud is available at all. You need to deploy on premises and manage the cluster by yourself. Then you replicate and aggregate some anonymous sensor data from all continents to get new aggregated insights. But some specific user data might always just stay in the country of origin and local region.

### Live Demo – 100000 Connected Cars

We built a live demo (available on Github) to show how easy it is to set up a connected car infrastructure at scale (including a cutting edge streaming machine learning example for predictive maintenance in real time): [Streaming Machine Learning at Scale from 100000 IoT Devices with HiveMQ, Apache Kafka and Tensorflow](#):

Such a setup can be the foundation for most of the discussed use cases in this blog post.

[You can also take a look at the following link to see the slide deck and video recording. They discuss the demo architecture and walking you through the live demo.](#)

