



Managing the data revolution: data streaming at Porsche

10/11/2020 Data is a valuable resource, which needs to be handled systematically. Sridhar Mamella works as a Platform Manager for Data Streaming Platforms at Porsche. He explains, why it's so important to streamline data and how their tool Streamzilla helps Porsche's engineering product teams to work more efficiently.

We are in the midst of a data revolution that is transforming how companies do business. Once a scarce resource, data has become abundant, fast and cheap. Yet this revolution isn't solely about the staggering pace at which we generate this information: it's also about new technologies that change the way we produce, collect, process, store, and analyse it.

With new streams of data being created every day, and with Industry 4.0 and the Internet of Things on the horizon, there is significant value in taking a strategic approach to so-called Big Data. Turning numbers and information into actionable insights is a key factor. In fact, organisations can gain a competitive advantage by having instant access to their data and by analysing, understanding, and managing it in real-time. Big Data and data streaming technology play a pivotal role in accomplishing

these objectives.

What is data streaming?

Devices, sensors, and interconnected systems continuously generate enormous amounts of data. Using data streaming technology, we can aggregate and integrate it from disparate sources into a single platform.

To understand data streaming, it is also useful to understand the difference between batch processing and stream processing. Batch processing is an efficient way to handle large volumes of data, but it is not suited to handling data while it is still in motion. Data streaming, in contrast, allows us to process and analyse the numbers and information in real-time and get immediate insights.

Data streaming at Porsche

The recent rise of data streaming has opened new possibilities for real-time analytics. At Porsche, data streaming technologies are increasingly applied across a range of contexts, including warranty and sales, manufacturing and supply chain, connected vehicles, and charging stations. For example, there are numerous sensors in the Porsche Taycan that continuously scan the vehicle's internal and external environments. Thanks to the data streaming technology, this connected car is now capable of processing the information provided by these sensors and assisting the driver in real-time.

Data collected from onboard telematics devices on a modern car can be categorised as either behavioral or diagnostic data:

- Behavioural Data is generated by or is in response to, the driver's use of the vehicle. For example, telematics data such as speed, steering, braking, and fuel-efficient driving is streamed into a safe and secure central system. This streamed data can then be used to issue alerts when the machine learning algorithms suggest driver fatigue.
- Diagnostic Data results from the ability to access a vehicle's data, this could for example enabling manufacturers to assess the health of a vehicle and notify drivers when a service is required with in-car notifications.

As a Platform Manager at Porsche, along with defining the data streaming strategy, my job has been to build and develop a highly available and highly secure central data streaming solution, which we named Streamzilla.

Streamzilla: the one-stop-shop for all data streaming needs

Streamzilla aims to be the one-stop-shop for all the data streaming needs within Porsche. It is an

internally managed service that enables different engineering product teams to build and run applications that take advantage of the low latency, high throughput and fault tolerance capabilities of the Apache Kafka distributed system along with Apache Kafka. We also use Apache NiFi for automating and managing the flow of data between systems. These are both highly scalable open-source platforms that enable real-time data streaming pipelines and applications.

Why do we need Streamzilla?

With Streamzilla, the Porsche engineering product teams can use native Apache Kafka & Apache NiFi APIs to populate data lakes, stream changes to and from databases, and power machine learning and analytics applications.

Apache Kafka & Apache NiFi clusters are challenging to set up, scale, and manage in production. When the individual teams run these complex clusters on their own, they need to provision servers, configure a distributed open-source technology manually, replace servers when they fail, orchestrate server patches and upgrades, architect the cluster for high availability, ensure data is durably stored and secured, setup monitoring and alarms, and carefully plan to scale events to support load changes.

Streamzilla makes it easy for our engineering product teams to run production-level applications without needing the open-source infrastructure management expertise. That means they spend less time managing infrastructure and more time building remarkable applications to enhance in-car services, improve real-time navigation, increase fuel efficiency, advance predictive maintenance, convert cars into 5G WiFi hotspots, refine driver assistance, and – most importantly – put a smile on our customers' faces when they drive their Porsche.

Info

Author: Sridhar Mamella, Platform Manager Data Streaming at Porsche.

Photographer: Richard Pardon

MEDIA ENQUIRIES



Stefan Mayr-Uhlmann

Spokesperson Finance and IT
+49 (0) 1523 / 911 7804
stefan.mayr-uhlmann@porsche.de

Consumption data

Taycan Turbo (2023)

Fuel consumption / Emissions

WLTP*

Electric power consumption* combined (WLTP) 23.6 – 20.2 kWh/100 km

CO emissions* combined (WLTP) 0 g/km

CO2 class A Class

*Further information on the official fuel consumption and the official specific CO emissions of new passenger cars can be found in the "Leitfaden über den Kraftstoffverbrauch, die CO-Emissionen und den Stromverbrauch neuer Personenkraftwagen" (Fuel Consumption, COEmissions and Electricity Consumption Guide for New Passenger Cars), which is available free of charge at all sales outlets and from DAT (Deutsche Automobil Treuhand GmbH, Helmuth-Hirth-Str. 1, 73760 Ostfildern-Scharnhausen, www.dat.de).

Link Collection

Link to this article

<https://newsroom.porsche.com/en/2020/innovation/porsche-data-streaming-revolution-tool-streamzilla-sridhar-mamella-platform-manager-22744.html>

Media Package

<https://pmdb.porsche.de/newsroomzips/7d61b805-0120-4701-9df3-2eec1d94f045.zip>

External Links

<https://medium.com/next-level-german-engineering>