

ACCELERATING AUTONOMOUS CAR DEVELOPMENT WITH READY ACCESS TO GLOBAL DATA

A leading automaker makes driving safer by building intelligent and autonomous vehicles

Industry

Automotive

Objective

Develop autonomous driving capabilities to improve vehicle safety

Approach

Build a platform to efficiently manage and analyze millions of miles of vehicle test data

IT matters

- Extracts data from test vehicles quickly and efficiently
- Makes data available to developers and data scientists around the world
- Enables direct access to data, including from legacy systems
- Leverages deep learning techniques to train algorithms
- Scales storage easily and effectively
- Ensures high availability and data resiliency
- Enables developers to run containerized applications

Business matters

- Accelerates the development of driver assistance systems and autonomous driving functions
- Enhances vehicle and road safety through improved driving systems
- Gains economic advantage by using a platform that does not require extra investment in tools
- Optimizes compute resources through its ability to accommodate multi-tenancy on its data platform
- Lowers barriers to entry for additional projects



Global automaker Mercedes-Benz AG aims to develop cars that are not only intelligent but also autonomous. To do this, the company needed a data platform to collect and manage massive amounts of data from test vehicles and make them available to developers across the world. In 2016, the automaker deployed HPE Ezmeral Data Fabric (formerly MapR Data Platform), accelerating its development of autonomous driving functions with ready access to global data.



The HPE solution has worked out of the box as the enterprise data platform for the automaker. It enables the company's data science teams to easily store, manage, process, and analyze data in one place.

CHALLENGE

Making intelligent and autonomous driving a reality

Imagine being able to let go of the steering wheel while driving and still be able to resume control of it—or to totally disengage from the driving process. Global automaker Mercedes-Benz AG is moving closer to making this a reality as it works to develop highly automated and fully autonomous vehicles.

The company believes that by developing such vehicles, it can provide accident-free driving and enhance vehicle and road safety. But achieving this requires rigorous testing and the ability to collect massive amounts of data from cameras and sensors attached to a car, to perceive its surroundings and detect issues. The automaker must then make this data available to developers and data scientists, who train algorithms known as deep neural networks. Through these algorithms, a car can learn to make smart decisions in real time using sensor data, enabling it to drive safely. Data scientists and developers also use this data to identify anomalies, helping them refine automated and autonomous driving systems.

The challenge was how to manage and analyze the enormous amounts of data it was generating. Because it has development sites all over the world, the automaker needed a solution that would allow it to access and share data globally with high performance on a massive scale, to accelerate its development of autonomous driving functions. The company also wanted to minimize replication and avoid data duplication to optimize hardware resources.

SOLUTION

Selecting a data platform

As its previous technology stack could no longer keep up with constantly increasing requirements, Mercedes-Benz AG considered options for its enterprise-grade

data platform, including HPE Ezmeral Data Fabric. The company's data scientists and developers particularly liked the way they could directly access data within HPE Ezmeral Data Fabric, even from legacy applications. The platform's direct access file system supports real-time reading and writing of data flows. Using a Windows or Linux® system, data scientists and developers can access the files directly without needing a Hadoop data adapter.

Performance was another key factor—the company's assessment of HPE Ezmeral Data Fabric showed that data would get to developers very quickly. The fact that the platform came with an entire data ecosystem also stood out in the evaluation. This offered the auto manufacturer a cost advantage compared to other solutions because it would not need to invest in additional tools, processes, or workforce resources.

During the selection period, the company had to deal with a lot of small files in its distributed file system, which slowed the performance of its clusters. Of the data platforms it considered, only HPE Ezmeral Data Fabric offered a complete solution to this problem.

The automaker implemented its first HPE Ezmeral Data Fabric solution in 2016. The solution supported its autonomous driving program and subprojects, which spanned automated driving, advanced driver assistance systems, and vehicle endurance testing. The initial footprint was a multi-PB platform comprising a hybrid solid state drive and hard disk drive solution that optimizes analytic performance and cost.

For the data platform's hardware infrastructure, the company deployed HPE ProLiant DL380 servers and HPE Apollo 6500 systems with NVIDIA® graphical processing units (GPUs). High-density HPE Apollo 4530 and HPE Apollo 4510 systems supported the storage needs of the Hadoop and HPE Ezmeral Data Fabric environments.



Customer at a glance

Software

- HPE Ezmeral Data Fabric

Hardware

- HPE Apollo 6500
- HPE Apollo 4530
- HPE Apollo 4510
- HPE ProLiant DL380

The HPE solution has worked out of the box as the enterprise data platform for the automaker. It enables the company's data science teams to easily store, manage, process, and analyze data in one place. To scale storage, the company only has to install another server—it can non-disruptively add or remove storage elements without impacting the rest of the cluster. With the platform's native triple data replication, high availability and data resiliency are guaranteed.

BENEFITS

Accelerating autonomous vehicle development

Mercedes-Benz AG has accelerated the overall development of its automated and autonomous driving functions, largely thanks to its ready access to global test data. With the help of HPE Ezmeral Data Fabric, data from test vehicles around the world is quickly synchronized across its cloud storage infrastructure. The process is boosted by the high bandwidth within the platform's clusters. This enables the company to promptly share data across development sites, so its data scientists and developers can analyze it efficiently and know what to test for or refine.

With direct access to data from current and legacy systems, the company's data scientists and developers can get the information they need without having to switch applications or use a new interface. This has made access more efficient and allowed the automaker to integrate data from legacy applications into its autonomous driving initiatives.

A significant advantage for the company has been the ability to reliably and securely support multiple tenants and project teams. From starting as a small project involving about 100 users, it now has several hundred data scientists and developers using its clusters without experiencing system breakdowns. Not only does this optimize resources and increase compute power for each data scientist, it also lowers barriers to entry for additional projects.

Leveraging deep learning techniques

With their ability to efficiently extract what they need from terabytes and even petabytes of data, developers at the company have leveraged deep learning techniques to train neural networks on representative data sets. To achieve this, they have built Hadoop data lakes using information stored on HPE Ezmeral Data Fabric to create a pipeline for moving data efficiently from test vehicles to clusters. Developers can then quickly and efficiently analyze a data set, divide it into segments, and organize training data for labeling and then deep learning.

For example, to enable a car to detect and recognize traffic lights, developers extract thousands of images of lights from the company's data platform. They label the images and use them to train vehicle neural networks, so these algorithms can recognize traffic lights by themselves based on data collected from the streets. Successes in training these algorithms are important steppingstones for the company in its quest to develop autonomous vehicles.

Supporting Kubernetes and other applications

The auto manufacturer has discovered unexpected further uses for its HPE Ezmeral Data Fabric. One of them is running Kubernetes applications on the system to analyze and process data natively on the same cluster. The Kubernetes integration provides a persistent storage layer for access to any data. The company uses the interfaces that HPE has made available to process data efficiently within Kubernetes clusters, enabling its developers to run their own containerized applications.

The automaker can also use the HPE solution to process data after it has been ingested. This means data is converted into a format that the platform can read more easily. The company can then leverage the compute power of the platform's servers to process information in a distributed way.

LEARN MORE AT

hpe.com/datafabric

hpe.com/ezmeral

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Get updates

SHARE

