Autonomous Transportation

Stable and Scalable Infrastructures for Autonomous Vehicles

A world-leading peer-to-peer ridesharing company selected DDN to create, deploy and optimize a global AI-enabled IT Infrastructure for its autonomous vehicle program.

A leading US university is leading a joint project between the private sector, government and academia together to advance transportation safety, sustainability, and accessibility.

A premium auto manufacturer deployed a DDN all-flash solution to simplifiy data aggregation at scale, and uses an open storage system that works with a variety of compute architectures.

Whether its planes, trains or automobiles, the move towards autonomous vehicles is revolutionizing the transportation industry. Powered by big data, AI and millions of dollars in research, the creation of safer, smarter, more reliable transportation systems is quickly being realized.

Autonomous vehicles engage some of the toughest workloads in AI and at an unprecedented scale. They require the handling, ingest and delivery of a broad range of dataset types and sizes, generated from many different sources such as video cameras, radar, lidar, sensors, and in-vehicle software. Very large datasets captured over millions of miles undergo many cycles of processing, labeling, sub sampling and categorization before being presented to deep learning (DL) applications that ultimately deliver the best-in-class self-driving vehicle.

However, to do so, self-driving vehicles require the maximization of the number of testing scenarios to improve vehicle perception accuracy and operational autonomy. This requires a reliable data storage framework that scales to Terabytes per second of throughput and hundreds of Petabytes of capacity. Ride-sharing companies, auto manufacturers, as well as freight and shipping companies leverage DDN solutions to support massive data collected by extensive and complex DL frameworks that are trained, tested and refined for the autonomous driving.

DDN enables infrastructures to harness data at an immense scale, which successfully and reliably build an advanced AI framework. DDN storage platforms effortlessly handle concurrent ingest of massive data streams, organizing and structuring the underlying datasets. This allows millions of GPU cores to continuously access DDN systems while executing extensive and complex training processes, and continuously refining self-driving capabilities.

Managing Data at the Endpoint, Edge, and Core

		Workload	Features
Endpoint	Vehicle Collection	Modest ingest rates, small and large files, extremely high sequential egress rates. Between 40-200TB per day per vehicle.	Low-power, high-data rates or removable, simple management.
Edge	Regional Data Center	High concurrent sequential ingest rates. Steady egress rates in syncing data to core(s). 10s of PBs.	Data management and tracking of data sets to core.
Core	Core Data Center(s)	Transform: Random and sequential read and write intensive. Large variety of file sizes. Train: Compute/GPU intensive, read heavy, random, small files common. Inference: Compute heavy, read intensive. 100s of PBs	Hybrid Cloud, full data management between CORE systems and EDGE datacenters. Auditing and control.