



Fast-Track the Future of Driving

Distributed Deep Learning for Autonomous Driving on Dell EMC Infrastructure

Artificial Neural Networks

The state-of-the-art ANN models used in DL for autonomous driving have more than 100 layers, requiring computations to be able to scale across many compute nodes.

Power autonomous driving and advanced driver assistance workloads with a deep learning (DL) solution built on Dell EMC PowerScale storage and Dell EMC PowerEdge servers with NVIDIA® GPUs.

With this Dell Technologies solution for DL, designed in collaboration with NVIDIA, automotive equipment manufacturers can harness the power of artificial intelligence (AI) to develop deep learning (DL) algorithms for autonomous driving (AD) and advanced driver assistance (ADAS) workloads.

This jointly engineered solution is built on GPU-accelerated Dell EMC infrastructure and includes support for containers. It gives data scientists and developers the solution they need to automate and accelerate deployment of applications at massive scale across diverse environments to support next-generation automotive capabilities.

Overcome roadblocks to innovation

The combination of increased computational power and access to large, comprehensive verification data sets is increasing the accuracy and sophistication of AD/ADAS algorithms and driving the industry ever closer to realizing the vision of fully autonomous vehicles.

However, training and validating the deep neural networks used in AD/ADAS development requires processing large amounts of sensor data coming from multiple sources such as cameras, lidar, radar and ultrasonic sensors. In addition, these advanced algorithms must operate within a complex set of circumstances including variable weather conditions, visibility and road surface quality. This creates key challenges for training DL workloads for AD/ADAS, which include:

- **Massive data sets** — Safety requirements demand higher image resolutions than other industries, which increases the size of unstructured sensor data sets. Typical total storage requirements average from 50–100 petabytes of data per vehicle model.
- **Fast training cycles** — To accelerate time to market, neural network training must be as fast as possible. The ability to scale ANNs, train large data sets across GPU servers and scale-out storage is critical.
- **Intense test and validation requirements** — Safety standards demand exhaustive testing and verification of the trained algorithm to represent diverse variables.
- **Data labeling** — Labeled data yields better model performance for AD/ADAS DL training. Labeling massive collections of training data with high quality is a tedious task that requires significant effort.

Learn more

- [Dell Technologies Solution: Distributed Deep Learning Infrastructure for Autonomous Driving](#)
- delltechnologies.com/ai
- [Dell Infohub](#)

Solution highlights

DL algorithms have a diverse set of requirements with various compute, memory, I/O and disk capacity profiles. The right infrastructure — including ample compute power, fast networking speeds and scalable storage capacity — is crucial.

The Distributed Deep Learning Infrastructure for Autonomous Driving is built on high-density Dell EMC PowerEdge C4140 Servers, designed to handle the most demanding DL workloads. Each server supports up to four NVIDIA V100 GPUs, the first GPU to break the 100 teraflops barrier of DL performance. Dell EMC PowerScale Isilon F800 all-flash scale-out NAS powered by PowerScale OneFS provides low latency, high throughput and massively parallel I/O to shorten time for training on data sets up to thousands of petabytes. The complete solution also helps you achieve AI at scale in a cost-effective manner by leveraging PowerScale for DL as well as other AD/ADAS workflows.

Components

The solution [whitepaper](#) demonstrates the design of a large-scale DL infrastructure for the development of AD/ADAS use cases and includes performance testing results.

Configuration options

Servers	Storage	Networking	Software
Dell EMC PowerEdge C4140 with NVIDIA V100 SMX2 GPUs and NVIDIA NVLink™ (compute and head nodes)	Dell EMC PowerScale Isilon F800 all-flash scale-out NAS	<ul style="list-style-type: none"> • Dell EMC PowerSwitch S3048-ON (management) • Mellanox® SB7800 100GB/s EDR InfiniBand® • Mellanox SX6710G gateway switch 	<ul style="list-style-type: none"> • Data Science Provisioning Portal • Bright Cluster Manager® • NVIDIA Container Toolkit • NVIDIA Data Loading Library (DALI)

Dell Technologies and NVIDIA

Dell Technologies and NVIDIA have been working together closely for many years to target some of the world's most complex computing problems.

Accelerate autonomous driving

Get started building a solid foundation for large scale, enterprise-grade DL solutions with a future-proof scale-out architecture that meets your needs today and scales for the future.

