Solution brief

Driving the future with AI and HPC

In-depth design guidance for autonomous driving architectures

Six levels of autonomous driving[1]

0. Driver support is limited to warnings and momentary assistance.
1. Driver support includes steering OR braking/accelerating.
2. Driver support includes steering AND braking/accelerating.
3. Car drives itself under limited conditions, such as a traffic jam.
4. Car drives itself under limited conditions, such as a driverless taxi.
5. Car drives itself everywhere under all conditions.

The road to level 5 autonomous driving is paved with data.

Autonomous driving (AD) holds the promise of safer, more comfortable and more efficient transportation. Many of today's consumers eagerly look forward to a day when autonomous vehicles can operate tirelessly at any time of day or night, with safety precautions and procedures ready to respond to any scenario.

At the same time, AD promises to provide more comfort. As drivers become passengers, they can relax or use their time for more interesting or productive activities. In addition, AD will lead to more efficient use of the roadways, as cars coordinate with each other and with traffic management systems. However, while advanced driver assistance systems (ADAS) such as lane assist and adaptive cruise control are already in wide use, the industry still has a long way to go to reach full AD capabilities.

To help frame the discussion, the Society of Automotive Engineers (SAE) divides the journey to AD into several levels. In levels 1 to 3, the car supports the driver, but the driver is still in control of the vehicle. This is considered ADAS. At the highest levels — 4 and 5 — the car takes care of everything, which is considered AD.

Several major steps are still needed for the industry to reach full AD capabilities. As each new level is achieved, moving to the next requires extensive development, test driving, and scenario development, with data processing and storage requirements growing exponentially at each level. In fact, level 5 AD will likely require collecting, processing and analyzing exabytes of data. This reality makes data the defining characteristic of AD development — both the sheer volume and the uncertainty about its growth.

To handle this data volume, High Performance Computing (HPC) systems, in some cases supported by artificial intelligence (AI), must provide high throughput to power many parallel streams of data analysis, simulation and correlation to deliver high-end simulations. Car manufacturers and suppliers on the road to AD must be ready to roll out IT infrastructure capable of supporting these needs at every level.

However, the IT infrastructure needed to support ADAS or AD development is large and complex, often consisting of thousands of servers and several software stacks. There is also no "one size fits all" approach for AD and ADAS architectures because each car manufacturer or parts supplier has its own requirements, approach and development roadmap.

However, there are some universals: The infrastructure needs to be performant, efficient, cost-effective and robust enough to span the development cycle for one level, and scalable enough to span several development levels. The design should also be able to incorporate new hardware and software as new insights emerge, new tools are developed, and new regulations come into place.

[1] Society of Automotive Engineers (SAE), SAE Standards News: J3016 automated-driving infographic update, July 2019.

Learn more

Read the Guide
DellTechnologies.com/HPC

A custom approach to simplifying solution stacks

Whether the focus is on ADAS or full level 5 AD, Dell Technologies has the expertise and experience to help you design HPC and AI solutions to support your current level and scale to the next.

Dell Technologies has decades of experience designing powerful, flexible and cost-effective systems for HPC and AI, and delivering them in a simplified and customizable building block model. "The ADAS/AD Architecture," a white paper published as part of the Dell Technologies *HPC & AI Innovation Exchange* series, provides an in-depth technical analysis of a range of solutions for manufacturing and automotive companies working to develop AD and ADAS. It takes a deep dive into the various options available for specific use cases and workloads, including discussions of remote site and data center infrastructure. It also includes considerations for software, services and infrastructure design with complete architectural design examples. These include:

- **Dell EMC Ready Solutions for AI** built with Dell EMC PowerEdge R740 and C4140 Servers with NVIDIA® GPUs, Dell EMC PowerSwitch S3048-ON and InfiniBand® SB7800 networking, and Dell PowerScale Scale-out NAS storage.
- **Dell EMC Ready Solutions for HPC** built with a choice of Dell EMC PowerEdge R640, R740, R840 or R940xa Servers for simulations, or Dell EMC PowerEdge R740, C4140, and the DSS8440 for deep learning.

Dell Technologies and NVIDIA

Dell Technologies and NVIDIA work together closely to deliver unprecedented acceleration and flexibility for AI, HPC and data analytics to help our customers tackle some of the world's toughest challenges. For Dell Technologies solutions for AD and ADAS, multiple diverse AI deep neural networks and algorithms for computer vision, localization, and path planning run on a combination of integrated NVIDIA GPUs, CPUs, deep learning accelerators and programmable vision accelerators.