

Zenuity, a joint venture of Volvo Cars and Autoliv, is racing ahead on the road to autonomous driving with software powered by deep learning. The company built their AI infrastructure on NVIDIA® DGX-1™ and Pure Storage® FlashBlade™, delivering the ultimate platform for their data scientists and researchers to make autonomous vehicles smarter and safer.



BUSINESS TRANSFORMATION

Due to a breakthrough in AI infrastructure, Zenuity has taken a leadership position in autonomous driving with deep learning. Data scientists can process the amount of data needed to ensure safety without compromising product timeline.

GEO

Global

INDUSTRY

Automotive

“We are in a global race to see who can deliver the safest solution for autonomous vehicles first. The combination of NVIDIA DGX-1 and Pure Storage FlashBlade provides a breakthrough AI infrastructure that enables us take a leadership position.”

Benny Nilsson,
Manager of Deep Learning

DEVELOPING THE VEHICLES OF THE FUTURE

As companies race to develop vehicles which will transform how we travel, how goods are transported across countries, and ultimately how our cities are designed, they are facing huge challenges.

At the core of successfully developing autonomous vehicles is the software that will be entrusted with making crucial, real-time decisions, as accurately as any human driver can. Ensuring the safety of passengers is of crucial importance, but companies are also looking to ensure their time-to-market is competitive. With estimates of the autonomous driving market being worth \$40 billion by 2025, the opportunity is huge.

Zenuity, a Sweden-based company with offices also in Germany and the United States, is tackling the greatest challenges faced by the automotive industry by developing a software strategy which places the company's expertise in artificial intelligence (AI) and deep learning at the centre. A joint venture of Volvo Cars and Autoliv, Zenuity will develop the software 'brain' of the car, whilst Autoliv will bring the technology to the global market.

ZENUITY USES DGX-1 AND FLASHBLADE TO ACCELERATE AUTONOMOUS DRIVING

The deep learning work at Zenuity has its roots in a project that began at Autoliv in 2015. Benny Nilsson, who is now Manager of Deep Learning at Zenuity, was part of that original project. “We spent a lot of time at Autoliv debating and testing for the best architecture for deep learning,” Nilsson recalled. “Our first system was an off-the-shelf server with GPUs (graphics processing units) and a typical enterprise storage system. We learned a lot from that system, both in terms of what works and what doesn't.”

This prior experience with GPU technology proved invaluable when thinking about what an optimal deep learning training platform needed to deliver at Zenuity. They knew they needed ground-breaking performance to handle the exponentially increasing volume of sensory data to be trained on their models. This would permit them to iterate on models developed rapidly without losing valuable development cycles. Additionally, their researchers couldn't afford to waste time playing the role of IT system integrators; designing, building and troubleshooting hardware and software or scouring open source forums in search of answers.

One compute platform emerged as the best fit for their needs — the DGX-1 deep learning platform from NVIDIA. “DGX-1 is clearly the gold standard when it

COMPANY:

Zenuity
www.zenuity.com

USE CASE:

- AI/Deep Learning – NVIDIA DGX-1™, Pure Storage® FlashBlade™, Tensorflow, PyTorch

CHALLENGES:

- Autonomous driving software is entrusted with making crucial, real-time decisions.
- The company needed to handle the exponentially increasing volume of sensory data.
- The industry faces a trade-off between processing more data — which is essential to deliver improvements in safety — and accelerating time-to-market.

IT TRANSFORMATION:

- With its compact form factor and simplicity, FlashBlade delivers a petabyte in just 8U.
- The solution enables the company to make ground-breaking progress in reducing training run intervals.
- The company reduced training times, which means IT investment is better utilized and data scientists are more productive.

“With the limitless scalability of the DGX-1 and the FlashBlade system, our deep learning infrastructure is future-proof.”

Benny Nilsson,
 Manager of Deep Learning

comes to AI and deep learning,” he noted. DGX-1 is engineered to deliver the highest deep learning training performance while offering the fastest deployment model with a plug-and-play deployment experience, and simplicity of management and day-to-day operations for data scientists.

Having decided upon a compute platform, the team turned to storage. “In our first system, we didn’t fully understand the importance of storage,” Nilsson said. “When we started testing the DGX-1, the storage system we had could not feed the GPUs with enough data to keep them busy. Data was the bottleneck and that’s when we realized we needed a modern storage system.”

Ola Eriksson, infrastructure expert at Zenuity, oversaw the selection process for storage. He identified ten storage vendors, and then developed a rigorous test regimen representing real-world deep learning workloads to compare their performance. “Our top priority was to ensure that all the DGX-1s were kept busy,” Eriksson said.

Added Samuel Scheidegger, Research Engineer on the deep learning team, “DGX-1 utilization directly translates to data scientist productivity. So before we purchased any storage system, we had to be able to test it in a hands-on environment.”

Eriksson said the key criterion was the cost per DGX-1 that Zenuity would rely on to provide sufficient bandwidth for deep learning workloads. “We evaluated how many DGX-1s each storage solution could deliver data to without throttling as the basis for calculating storage costs. On this measure, FlashBlade from Pure Storage was the clear winner.”

FlashBlade is architected with a modern, massively parallel architecture, eliminating serial bottlenecks that hold legacy storage systems back. It is engineered to deliver the performance essential for AI and deep learning while offering simplicity, so data scientists can focus on their data, not the infrastructure.

“With its compact form factor and simplicity, FlashBlade delivers a petabyte in just 8U,” Scheidegger observed.

DELIVERING THE POWER TO MAKE CARS SMARTER AND SAFER

Deep neural networks require a huge amount of computational power and tremendous amounts of data, which is particularly true with safety critical systems, like self-driving cars, where detection accuracy requirements are extremely high. In light of this, the computational demands placed on Zenuity’s IT infrastructure are considerable.

The primary workloads are neural network training jobs where GPUs are consuming huge amounts of data gathered from test cars. The data consumed consist of images, radar and LIDAR data, a surveying method involving light in the form of a pulsed laser. Data scientists at Zenuity are free to choose the framework of their choice, with the most popular frameworks being TensorFlow and PyTorch.

By processing increasing volumes of data with greater frequency, the autonomous driving software continues to learn and become more accurate over time, with each new version downloaded to vehicles. But these continuous iterations of network training are extremely time-consuming. Datasets for each training run can reach into the terabytes, and a single network training run can take weeks if the environment is not properly optimized. Zenuity is tackling this with the combined power of DGX-1 and FlashBlade, which is enabling them to make ground-breaking progress in reducing training run intervals, to the extent that they expect to be able to iterate on their models.

“FlashBlade is the only storage system that withstood the rigors of our deep learning benchmarks.”

Benny Nilsson,
Manager of Deep Learning

“The industry constantly faces a trade-off,” Nilsson said, “between the need for processing more data — which is essential to deliver ever-greater improvements in safety — and the desire to accelerate time-to-market. By combining NVIDIA DGX-1 with Pure Storage FlashBlade our data scientists can enjoy the best of both worlds — processing the amount of data needed to ensure safety without compromising product timeline.”

Zenuity currently has two FlashBlade systems — each equipped with fifteen 52TB blades — connected to multiple DGX-1s, which offer eight of the world’s most advanced data center GPUs — the Tesla V100, configured in an NVLink based hybrid cube-mesh topology for incredibly fast inter-GPU communications. FlashBlade keeps all DGX-1s busy with data, with significant performance headroom to give the Zenuity team the assurance that the combination of NVIDIA and Pure technology will meet their needs well into the future.

“Our training data sets are already in the hundreds of terabytes, and that number will increase exponentially as we collect more data and put more cars on the road,” Nilsson said. “With the limitless scalability of the DGX-1 and the FlashBlade system, our deep learning infrastructure is future-proof, able to scale in capacity by simply adding more storage blades to ensure we’re feeding enough data to keep the DGXs busy.”

ACCELERATING INTO THE FUTURE

The AI infrastructure of DGX-1 and FlashBlade has met Zenuity’s number-one priority — keeping its IT investment fully utilized. “If you’re not able to feed the GPUs with the data they need, then you are underutilizing your most valuable resources — data scientists,” Nilsson said. “The combination of NVIDIA DGX-1 and Pure Storage FlashBlade is helping us to reduce training times, which means IT investment is better utilized and data scientists are more productive and, ultimately, happier.”

For Zenuity, having systems which are simple to set up and easy to deploy has also been a crucial factor. “The DGX-1, and especially FlashBlade, are simple and require very little attention from day to day,” Eriksson said. “They both deploy quickly and you need very little support thereafter.”

Over time, Zenuity expects to grow its DGX-1 environment. This will be in direct response to the escalating size of their datasets; the need to tap into the power of multi-GPU training to process that data; and mounting business pressure to iterate faster and develop new, smarter autonomous driving neural network models. As Zenuity grows its DGX-1 based AI platform, the company’s broader infrastructure will need to keep pace, especially when it comes to storage.

“We are in a global race to see who can deliver the safest solution for autonomous vehicles first,” Nilsson noted. “The combination of NVIDIA DGX-1 and Pure Storage FlashBlade provides a breakthrough AI infrastructure that enables us to take a leadership position.”



info@purestorage.com
www.purestorage.com/customers

