

Winning The ADAS Race with Tech Mahindra and Pure Storage[®]

Driving

Autonomous

Mode

WHITEPAPER

48

mm

Why innovation in ADAS and AD programs is fueled by the power of data

Abstract

The automotive industry is undergoing a period of rapid transformation, with an increased demand for advanced driver assistance systems (ADAS) and autonomous driving (AD) technology. The global ADAS market size is projected to grow from USD 28.4 billion in 2021 to USD 76.5 billion by 2030, at a CAGR of 11 to 13%.

In this white paper, we explore how suppliers can optimize their infrastructure and the way they ingest and manage growing volumes of data that ADAS and AD solutions rely on, while also under pressure to assign higher CAPEX budgets.

Table of Content

02

Introduction

04

Partners in Automotive Innovation

09

The Technology Driving Transformation

03

Challenges in Accelerating ADAS and AD Programs

05

The Technology Driving Transformation



Introduction

Automotive original equipment manufacturers (OEM) and Tier 1 suppliers are in the race to deliver ADAS and AD solutions earlier than their competitors. According to our findings, at least 30% of the investment is being made in data management and hybrid data centers.

Development programs typically follow five key stages:

- 1. Build and run test vehicles using test path and data collection
- 2. Pre-process and manage the data
- 3. Undertake vehicle training and inference development
- 4. Test the software and hardware modules
- 5. Complete vehicle integration and validation



Figure1: ADAS/AD Development and Testing Lifecycle

Challenges in Accelerating ADAS and AD Programs

In OEMs and Tier 1 companies, ADAS development programs are managed by engineering teams, and the data management and infrastructure are managed by the chief information officer (CIO). This creates a number of challenges for research and development infrastructure teams in ADAS development:

Key challenges faced by R&D and Infra teams in ADAS development



Figure 2: Key challenges faced by R&D and Infra teams in ADAS development

Challenges While Developing ADAS Features (Using Camera-based Data)

Object detection and classification is a computer vision technique that recognizes images, and it's the most valuable technology in enabling such systems. In modern cars, AI systems read data from cameras and other sensors, and then use pre-trained models to act. These systems use complex, neural networks to read the incoming video and sensor feed and generate inference data.

Developing these pre-trained models has a major challenge for compute and storage, as a training a system that is accurate 99.99999% of the time also requires a huge amount of video data. Multiple video feeds may be required from different angles and locations to compensate for a moving environment, and weather factors can also alter the appearance – and the training algorithm must account for such issues. In addition, there could be supplemental data in the form of radar, infrared images or light detection and ranging (LIDAR) to improve the accuracy of the model. So, the compute engine needs to be able to handle massive amounts of data from different sensors.

The training stage is currently performed on a server in a data center. The data is collected from a car in motion and the neural network model is trained for days or months, it is then downloaded to the vehicle periodically. Storage is a critical element of the infrastructure for both training and inference because the sheer volume of data in a vehicle could overwhelm a storage system. Take 360° cameras, which could require up to 17 video streams to be processed, stored, and then fed into the training system. And that volume increases further when considering additional data from elsewhere, such as radar sensors.

In vehicles, storage requirements are constrained by real-time decision-making needs. Automatic emergency braking (AEB), for example, must calculate the results within a permissible time, depending on the speed of vehicle, to avoid a collision. The storage system onboard must cope with multitudes of cameras and sensor data processing in real time.

Manufacturers are therefore seeking ways to enable training in real-time, so that vehicles can be updated on any issues immediately. For instance, a pothole could be identified quickly by a training system and the information passed on to all vehicles coming up to it to avoid an accident.

Partners in Automotive Innovation

As global partners, Tech Mahindra and Pure Storage can help address these challenges and power the development of ADAS and AD programs.

With its unique end-to-end engineering capabilities, Tech Mahindra has more than 18,000 associates in the automotive sector and partners with 80% of the top ten automotive OEMs globally. It delivers continuous innovation to OEM and Tier 1 customers in the automotive sector, with cutting edge solutions and turnkey product design and development programs. As well as smart solutions for infotainment, customer experience (CX), factory and aftermarket domains.

And Pure Storage is supporting the next era of human advancement with leading edge AI and data analytics solutions. Providing the world's first full-stack, AI-ready infrastructure—powered by NVIDIA and FlashBlade®, it helps customers reduce time to insight and drive AI at scale.

By combining innovative technology and industry knowledge, Tech Mahindra and Pure Storage can help customers drive key business outcomes, such as workload consolidation, improved TCO, reduced emissions, simplified architecture, shorter time to market, enhanced user experience, improved quality, and global presence. Let's explore how.

- Edge data centers and ingestion stations need to be established and managed globally. We have a global presence, so one can manage local regulatory requirements in different countries and acquire any necessary licenses
- Our knowledge and capabilities bridge the gap between disconnected IT and engineering teams
- We help customers reduce time to market by providing the latest technology
- Building their own infrastructure means customers face the risk of incorrect sizing and incompatible product selection. We deliver it as a service, so these risks are eliminated allowing the customer to focus on vehicle
- We implement multiple edge platforms which enable engineers to remotely access and analyze sensor and video data to develop their autonomous driving systems
- The process is simplified, and disruption is reduced, so developers and analysts to work more productively

The Technology Driving Transformation

High-level ADAS framework

Tech Mahindra brings end-to-end services from AI and ML simulation and app development to testing and comprehensive simulation, which are integrated within a vehicle.



Figure 3: ADAS Data center architecture framework

To enable feature development and training data is ingested from a vehicle, stored for labeling, and then processed through a test suite. The defined rules are applied and testing, and validation is carried out before deploying the model back to the vehicle. And that cycle is repeated.

Each of these steps require different software stacks and traditionally separate silos of storage, with data being copied and transformed between each silo.

	GET DATA Ingest	GET DATA Refine	TRAIN & ADJUST Training	DEPLOY & TEST Testing
	RAW	LABELING	GPU DNN TRAINING	SIMULATION
Media type	Mechanical/Hybrid	Mechanical/Hybrid	NAND/Performance Optimized	Flash/SSD
Processing mode	Batch	Batch	Real-time	Batch
I/O type	Sequential	Sequential	Random & Sequential	Sequential
Architecture	Scale-out	Scale-out	Scale-out/ Massively parallel	Parallel

Figure4: Autonomous Drive Phases

Pure Storage's data hub philosophy eliminates silos and drastically reduces data movement from the ingest stage to training. It can handle mixed access patterns, sequential or random, and work with any file size. And this parallel architecture ensures workloads are always serviced with low latency.

Unified fast file object storage





A traditional ADAS infrastructure is often hand-built with multiple components, delivered by different vendors. Not only does this bring several challenges, but it is also not sustainable and hinders innovations as needs constantly change and environments grow.

With Tech Mahindra and Pure Storage, the data hub seamlessly integrates with the ADAS framework, removing silos of storage and the need to copy data between different zones. A single Flashblade data hub can simultaneously handle data ingestion, perform cleansing and training without the need to copy data between various silos.

Our perspective on vehicle training, simulation and testing



Figure 6: Data Collection Center

Multiple application types can be run, from data lakes to data warehouses, on a single unified fast file and object platform, which allows everything from ingest to training simulation and testing to be completed on a centralized platform. It simplifies deployment, adapts to change, increases quality, and reduces time to market.

Efficiency is improved by eliminating multiple copies of data, distributed across separate DAS cluster nodes. By consolidating them into one data hub, the deployment of new clusters and future applications is simplified and can be achieved on-demand in minutes.



Keeping GPUs busy with TensorFlow & 100Gb Ethernet with RDMA

Figure 7: Eliminating bottlenecks

Data scientists generally run AI workloads from a single server, with multi-node scaling only used by AI experts who are presenting their data at conferences. Tech Mahindra's and Pure Storage ADAS framework removes the barrier, enabling any data scientist to complete multi-node training.

In this chart, three popular neural networks can be seen. It shows two patterns:

- GPUs are kept at their maximum throughput
- Pure Storage's unified fast file and object platform allows jobs to scale linearly as more systems are added to a single job

Data needs to get all the way from storage over networks through CPUs, across peripheral buses and into the GPU. In DIY systems, the input pipeline can easily stall the GPU, slowing down the entire training process. With Pure Storage's data hub and Tech Mahindra's ADAS framework, the GPUs are continuously fed with data.



ILLUSTRATIVE CALCULATION BASED UPON CURRENT PROPOSED CONFIGURATION AND IS SUBJECT TO CHANGE AND CONFIRMATION OF COMPETITIVE CONFIGURATIONS

* Assumes DC PUE (Power Usage Efficiency) of 1.7 European average, source: https://www.opencompute.org/documents/the-current-state-of-data-center-energy-efficiency-in-europe-ocp-whit e-paper ** KWh uses market standard pricing - €0.10/kWh.

± CO2 and vehicle Calculation Source: https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator. ~ Assumes 4TB SATA disks, replication factor 3.

Figure8: Workload consolidation for a future proof architecture

This approach also helps automotive manufacturers meet their net zero goals with significant environmental benefits. This example shows:

- A reduction of 69% in power and cooling
- 87% in rack space vs using traditional DAS clusters with open-source software

Future Proof ADAS Solutions for Modern Containerized Applications

Disaggregation with fast S3 has driven an architectural shift. Traditionally, data pipelines and analytics applications were built in a way that leverages direct attached storage, originating from Hadoop Distributed File System (HDFS) and Google File System (GFS)

Advancement in technology means that this model is no longer relevant like it used to be a decade ago.

These disaggregated models continue to gain momentum with key benefits such as operational simplicity as the data pipelines get larger and larger. The stateless compute and storage offer a denser, more environmentally friendly architecture.

As needs evolve and complexity increase, the demand for agile containerized software environments grows. Pure Storage meets this demand with a Kubernetes data management layer that delivers persistent storage to modern containerized applications. Portworx^R provides enterprise data services alongside migration and DR, allowing applications to utilize the most appropriate storage, whether it's DAS, block, or cloud.

Driving Outcomes for OEMs and Tier 1s

Partnering with Tech Mahindra and Pure Storage to develop ADAS and AD programs, enables customers to deliver key business outcomes:



Simplified Architecture – Minimizes data movement, improving the management process and scaling a single namespace from 100TB to multi-petabyte scale, with consistent performance and enhanced user experience.



Workload consolidation – Eliminates silos into a single data storage platform.

Boosts innovation with better insights – Removes storage bottlenecks and allows faster analytics, which improves quality through more training iterations and reduces time to market.



Improved TCO – Reduces management overheads through simplification and elimination of up-front investment.



Global presence – Tech Mahindra's and Pure Storage global presence means all local regulatory requirements are met.



Environmental sustainability – Supports customers net zero goals through reduction of power, cooling, and data center footprint



Conclusion

Together, Tech Mahindra and Pure Storage enable OEMs and Tier 1s to innovate in ADAS and AD with a high performing, highly scalable future proof infrastructure. Through a strategic partnership with two leading innovators, customers benefit from:

- The core competencies to focus on development and innovation, and faster time to market, while their IT is in the capable hands of Tech Mahindra.
- Redeployed infrastructure budget for better product innovation, and the opportunity to consume infrastructure in a pay-as-you-grow model.
- Centralized management of all data across multiple locations.

Together, we provide the vision, best-in-class solutions, and capabilities to unlock the power of data and fuel successful ADAS and AD journeys.

Authors

Krishnan Venkitaraman

Krishnan is a sales management professional with over 28 years of diversified experience in selling IT hardware and software solutions to leading global and corporate organizations.

His skills involve effective sales and partner management skills with proven track record of closing eminent deals both as an individual contributor and as a team player.

Prashant Khairnar

CXO Advisor | Global Digital Practice Lead – Smart Products

Prashant comes with 20 years of experience in the automotive industry. He is currently working as the CXO Advisor | Global Digital Practice Lead – Smart Products, developing capabilities in evolving technologies and enabling Tech Mahindra's customers developing ADAS, Autonomous driving technology, V2X, and connected products. He brings expertise in IT, infrastructure technology and application in R&D domain that includes engineering cloud, data and analytics, application of machine learning in design and development. He can be reached at Prashant.khairnar@techmaindra.com



тесн mahindra

Copyright © Tech Mahindra 2022. All Rights Reserved. Disclaimer. Brand names, logos and trademarks used herein remain the property of their respective owners.