

The ARM logo is displayed in a white, lowercase, sans-serif font in the top left corner of the page. The background of the entire page is a dark, teal-tinted image of a car's interior, showing the steering wheel and dashboard. On the right side, there is a semi-transparent digital overlay of a car's instrument cluster, featuring a speedometer with a needle pointing to 098 km/h, a temperature gauge showing 21 degrees, and a tachometer. The overall aesthetic is futuristic and high-tech.

arm

# How To Accelerate The Journey To Autonomous Vehicles


Flexibility, robustness and standard choices  
are key to the transformation

## Why you should read this

Autonomous workloads are changing the way we design products such as autonomous and semi-autonomous vehicles. It's an enormous opportunity to scale design with more flexible, cost-effective methods to deliver products that are more reliable, robust and safe. But there are roadblocks along the journey. Here are the drivers, challenges and potential solutions to enable this vision of the future.

## Introduction

---



**Autonomous workloads promise to deliver devices and systems that can safely decide and act for themselves to enable operational efficiencies working alongside humans or taking humans out of the loop, opening up valuable time that can be spent elsewhere**

What vertical segments are most keen to embrace autonomous workloads? Automotive and industrial. In the case of automotive, it's shaping up as a win-win embrace: Not only are designers pushing increasing amounts of autonomous functionality in their vehicles, but their factory lines and supply chain stand to gain just as much by leveraging autonomous workloads.

The design landscape is fragmented at the moment and requires a new way of thinking about how to design for these new workloads. Autonomous workloads will need the fundamental building blocks including heterogeneous compute, safety, security and machine learning, that are common across autonomous-based use cases.

## Section 1

# Why automotive companies look to autonomous workloads

---

**The need for autonomous decision-making workloads that are not dependent upon human interaction is clear, both for safety and efficiency reasons.**



Average time a driver takes to react and hit the brakes



A car equipped with vision systems, radar and LIDAR

Various sensors tied into a vehicle's braking system can react far more quickly than human reflex. It takes the average driver 1.5 seconds to react and hit the brakes when a child runs into the street; it takes a car equipped with vision systems, radar and LIDAR just 0.5 seconds. The combination of intelligent hardware coupled with the software running atop it – which does not rely on a connection to the cloud – will enable this not only in driverless vehicles but in the factories they emerge from.

Automotive factory managers want to leverage autonomous workloads to boost manufacturing productivity, improve KPIs and improve operational safety. They want to make their factories more flexible (and therefore more profitable) in building customized products with efficiency and adapting their product mix on the fly. Autonomy will lead to increasingly more efficient and intelligent robotic machines tied into ever more sophisticated building and networking systems to deliver gains not only in manufacturing but in the larger supply chain and logistics operations.

Managers want to transform their business by integrating their factories more tightly with the wider supply chain and logistics providers. Executives want to build fewer but more flexible factories to help achieve certain environmental objectives. Here, the goal is one autonomous chain from ordering, to raw material requisitions, to last-mile delivery, taking humans out of the loop or creating a safer and more effective human-machine collaboration.

Further enhancing the automation capabilities of ADAS (advanced driver-assistance systems) and IVI (in-vehicle infotainment), technology will make our current vehicles safer, while new technology advancements powering driverless systems will address the growing demand for last-mile delivery, enable fixed-route shuttles, and transform road-freight and ride-hailing services.

This is a tipping point for many companies in the history of compute. In automotive, already companies such as Waymo, Volkswagen, General Motors and Mercedes are taking advantage of autonomous workloads and new compute approaches. Competitors need to move quickly to reduce this first-mover advantage in a market that has tremendous long-term potential.

## Section 2

# Challenges to deploying the new workloads

---



### There are layers upon layers of complexity that must be addressed in designing for autonomous workloads.

Distributed sensors need to be unified to offer better insight and support the computational burden. Disparate systems and a dizzying range of software and hardware choices can make charting a path to autonomy challenging. For example, many current solutions for automated driving are based on bulky, high-power and costly standalone GPU chips. Others require proprietary solutions, which severely constrains design flexibility. However, to make it suitable for widespread commercial use, Tier 1 and OEMs are looking for a more power-efficient and cost-effective solution suitable for mass production.

To embrace the vision, we need to understand how fundamentally different autonomous workloads are from traditional workloads and confront the

challenges that complexity presents. Autonomous workloads are defined by receiving input and generating an output in their simplest terms – the ability to sense and take intelligent action. Thanks to the incorporation of AI and machine learning, the magnitude of workload is far greater when compared to something like the narrow scope of operation surrounding, for example, contemporary ADAS. This requires a fundamental reconsideration of design, embracing emerging technologies and methodologies.

Autonomous is also more complex than traditional workloads because of the overlapping and interconnected nature of systems in industrial applications. There may be a machine vision system with its own set of data and algorithms that needs to integrate with robot vehicles that are alerted by the vision system that a task has been completed and pickup is required. Integration may also be required for other points along the supply chain. For example, vehicles rolling down the manufacturing line will benefit from just-in-time parts arrivals

thanks to the more integrated, near real-time connection with the supply chain.

In both automotive design and industrial machines, there are the enormous computational requirements around capturing, fusing, analyzing and acting on enormous (and different – cameras and computer vision to LIDAR to GPS inputs) amounts of data. Layers of sensor inputs must be processed in real time in order to predict situations that machines can anticipate and react to. The computational power and efficiency to run the algorithms that process the data is what makes the difference.

Configuring these systems is costly. Compute alone can cost thousands of dollars, while the sensors themselves sometimes cost tens of thousands of dollars. Safety demands redundancy across the range of LIDAR cameras, radars, and that raises costs. As we move further away from driver control toward L4 and L5 autonomy, sensor and computational costs ramp significantly.

## Section 3

# How to scale hardware

### There are layers upon layers of complexity that must be addressed in designing for autonomous workloads.

Distributed sensors need to be unified to offer better insight and support the computational burden. Disparate systems and a dizzying range of software and hardware choices can make charting a path to autonomy challenging. For example, many current solutions for automated driving are based on bulky, high-power and costly standalone GPU chips. Others require proprietary solutions, which severely constrains design flexibility. However, to make it suitable for widespread commercial use, Tier 1 and OEMs are looking for a more power-efficient and cost-effective solution suitable for mass production.


To embrace the vision, we need to understand how fundamentally different autonomous workloads are from traditional workloads and confront the challenges that complexity presents. Autonomous workloads are defined by receiving input and generating an output in their simplest terms – the ability to sense and take intelligent action. Thanks to the incorporation of AI and machine learning, the magnitude of workload is far greater when compared to something like the narrow scope of operation surrounding, for example, contemporary ADAS. This requires a fundamental reconsideration of design, embracing emerging technologies and methodologies.

Autonomous is also more complex than traditional workloads because of the overlapping and interconnected nature of systems in industrial applications. There may be a machine vision system with its own set of data and algorithms that needs to integrate with robot vehicles that are alerted by the vision system that a task has been completed and pickup is required. Integration may also be required for other points along the supply chain. For example, vehicles rolling down the manufacturing line will benefit from just-in-time parts arrivals thanks to the more integrated, near real-time connection with the supply chain.

In both automotive design and industrial machines, there are the enormous computational requirements around capturing, fusing, analyzing and acting on enormous (and different – cameras and computer vision to LIDAR to GPS

inputs) amounts of data. Layers of sensor inputs must be processed in real time in order to predict situations that machines can anticipate and react to. The computational power and efficiency to run the algorithms that process the data is what makes the difference.

Configuring these systems is costly. Compute alone can cost thousands of dollars, while the sensors themselves sometimes cost tens of thousands of dollars. Safety demands redundancy across the range of LIDAR cameras, radars, and that raises costs. As we move further away from driver control toward L4 and L5 autonomy, sensor and computational costs ramp significantly.



## Section 4

# How to think about software

---

**For autonomous systems to succeed, an entire software stack and software ecosystem need to be developed.**




that abstracts many of the details about how the underlying systems work from the people who are thinking about the application level.

Traditional companies offer “sticky” proprietary software with no room for the developer to put his or her own programs on. Newer entrants have given these stalwarts a wakeup call with open systems, enabling the developer to have access to it or at least some portion of the control system to put their own control. Yet, autonomy needs to balance openness without losing all “stickiness” in order to head off the disruptors who propose totally open systems.

Customers that have much more demanding safety requirements will require safety-certified RTOS and middleware solutions. All of these elements will have some very demanding requirements as far as being able to achieve safety certifications such as ISO 26262.

As autonomous vehicles move from prototype to production or development to deployment, commercial software stacks will need to offer capabilities like real-time, fault tolerant, safety and security, and hence the commercial software ecosystem will be needed to replace or augment the open source and proprietary software that is currently in development. These stacks will need to scale across verticals as well. In other words, some portion of a software stack designed for a train can be used in a car. How that is standardized and secured is key.

Lastly, it’s critical to consider some common base architecture that allows your design team to pivot from open source to commercial SW as they reach deployment. This reduces complexity and improves time to market through quick migration of workloads.



# Section 5 How to think at a platform level

**To deliver on this vision and overcome the complexity issues, we have to not only bring down the cost of compute but also deliver choice and flexibility to developers ideally with a support ecosystem that can deliver everything from IP to algorithms to test and validation programs and protocols.**

New technologies in the design of CPUs, GPUs and ISPs will offer the crucial steppingstones to enhance ADAS and IVI capabilities in our vehicles, power new, driverless autonomy and enable more intelligent, configurable production lines in factories as we move into this new era of autonomous decision making. Machine learning, for example, will be a critical technology running on these devices.

CPUs need to offer the ability to run different workloads/application footprints, allowing the deployment of the same SoC compute architecture into

different domain controllers. They need to support features to help achieve both ASIL D and ASIL B safety requirements, allowing flexibility through its partitioning features that enable the highest safety levels when they're needed, or highest performance when they are not.

GPU technology needs to deliver "flexible partitioning" to allow the allocation of hardware resources at runtime - whether it's between safe and non-safe workloads or different types of workload. Such technology must be designed to ISO26262 and be ASIL B safety capable,

allowing the use of GPU resources for safe human-machine interfaces and for the heterogenous compute needed in autonomous systems.

Keeping with heterogeneity, developers need an ISP that can scale to support both real time and virtual cameras of varying types and offers a wide range of data output formats providing the flexibility to support both human and machine vision applications such as production line monitoring and quality control.

**Ultimately, this heterogeneous flexibility will deliver processing power that will enhance ADAS, IVI and autonomy through:**

- ✦ Enhanced safety
- ✦ Enhanced Security
- ✦ Machine learning
- ✦ Thermally and power efficient technologies
- ✦ Handling a wide range of data inputs and formats and processing with ease
- ✦ Shifting resources in real time

**This is the foundation for scalability that forms the basis of a standardized solution across applications, whether it's ADAS, IVI, autonomy, or the factories that manufacture the autonomous vehicles.**

## Conclusion

**The development occurring today and that lies ahead will transform the world as we know it. Heterogeneous and scalable compute will address the range of processing capability for differing workloads in a power-efficient manner, in ways that will transform how we design autonomous vehicles and factories.**

In automotive, by further enhancing the automation capabilities of ADAS and driverless technology and applying it to operational domains where routes, speeds, location and environmental conditions can be bounded, we can accelerate the journey to deploying autonomous vehicles.

For smart manufacturing, autonomy provides the intelligence to customize workflows, whether that's deciding what to manufacture next, or to optimize production line logistics. Factory production lines will become reconfigurable, enabling flexibility to meet the needs of differing products and their specific customisations, and this will add value and cut cost.

The promise of these advancements and the lessons from these early developments and relentless innovation from the ecosystem are the keys to delivering on the five key pillars of autonomous systems: Safety, scalable compute, flexibility, reliability and robustness.