RISC-V for Automotive AI Use Cases

Opportunities and Challenges

Introduction

As Artificial Intelligence (AI) is increasingly adopted across wide and varied use cases, the technology to deliver it must evolve to address new performance and capability needs. The automotive segment is one key area where AI deployment can yield substantial benefits, but requires modular scalable technology platforms and a supporting ecosystem to successfully adopt it. This paper will explore how RISC-V based systems are perfectly placed to enable AI across the automotive segment.

Market Dynamics

The automotive industry is undergoing a profound transformation driven by artificial intelligence. As AI continues to advance, the demand for more intelligent and efficient vehicles grows. According to Informa[1] the global automotive AI processor market is valued at $5.2bn in 2025 and is expected to grow, at a compound annual growth rate (CAGR) of 19.4%, to $12.6bn by 2030.

Shift to Zonal Architectures:

Laying the Foundation for Scalable Automotive AI

To meet these demands, automakers are shifting towards zonal architectures, where electronic control units (ECUs) are consolidated into fewer, more powerful nodes. As highlighted in an interview with Rivian, nearly all vehicle systems are expected to incorporate AI capabilities in the future. RISC-V, an open standard instruction set architecture, offers a promising solution. Its modularity, flexibility, and security make it ideal for developing scalable and adaptable AI platforms that can power the future of automotive.

By adopting a zonal architecture approach, automakers can deploy AI-powered features across various vehicle systems, e.g. from basic in-vehicle infotainment (IVI) voice control, to advanced driver assistance systems (ADAS) to fully autonomous vehicles. RISC-V's ability to be customized for specific AI workloads and right-sized for cost, coupled with the availability of multiple vendor options, make it a compelling choice for powering these zonal architectures.

Features:

- 64-bit high performance CPU
- 32-bit embedded CPU
- Multi-core, multi-cluster
- Vector computing
- CPU & I/O virtualization
- ASIL B, ASIL D
- Security

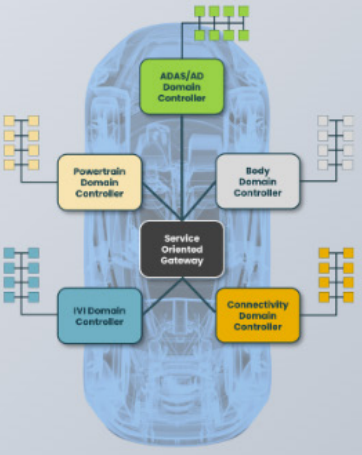

Domain architecture

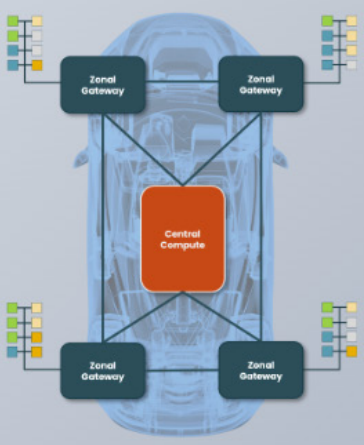

Zonal architecture

1 Omdia research data 2025 AI-equipped processors in automotive applications

Why RISC-V for AI in Automotive?

Architectural Flexibility in a Changing Landscape

In traditional system design, companies typically prefer incumbent processor technologies to mitigate development risks and control costs. However, the transformative potential of AI in automotive applications is compelling system architects to fundamentally reimagine computing platforms.

RISC-V represents a paradigm shift in processor architecture, offering unprecedented flexibility and efficiency for AI/ML workloads. Unlike proprietary architectures, RISC-V provides an open, scalable extendable standard that addresses the complex computational demands of modern automotive AI systems.

Technical Advantages of RISC-V for AI/ML

Architectural Scalability and Standardization

RISC-V's architecture delivers industry-leading performance through a unique combination of features:

- **Standard Instruction Profiles:** address the need for portability across vendors and enable development of a robust ecosystem. Each Profile specifies which ISA features are mandatory or optional, providing a common target for software developers and ensuring consistent compatibility across different implementations.
- **Custom Instruction Capabilities:** Allows radical performance optimizations for specific AI algorithms, creating unprecedented hardware-software co-design and product differentiation opportunities

The modular nature of RISC-V means the base instruction set remains compact, reducing implementation complexity while maintaining exceptional computational efficiency. Critically, the architecture's native support for variable bit-width processing makes it inherently well-suited for diverse AI/ML computational requirements.

Ecosystem and Adoption

Major technology leaders, including Google, Meta, and NVIDIA, have already recognized RISC-V's potential. The architecture's alignment with standard AI frameworks like TensorFlow, ONNX, and OneAPI enables precise hardware optimization, promising substantial performance and power efficiency gains. The open nature of RISC-V enables a collaborative ecosystem that depending on specific competencies and expertise contributes across the stack. Ranging from processor IP, through operating systems, development tools and end-user facing applications a diverse open ecosystem maximises choice while fostering competition and innovation.

Beyond Technology: Business Impact

The adoption of RISC-V in automotive AI at scale is not limited to a technological shift; it will deliver a strategic and transformative business impact to vendors across the supply chain. The approach addresses multiple industry challenges pertinent to the automotive sector, including:

- Complex and evolving supply chains
- Rigorous cost management requirements
- Stringent cybersecurity and safety standards
- Sustainability and electrification pressures
- Rapid regulatory compliance adaptation

As AI algorithms continue to evolve, RISC-V's flexible architecture ensures that hardware can rapidly adapt to emerging computational paradigms. But RISC-V represents more than an alternative processor architecture; it is a strategic enabler for automotive AI innovation, providing the technical flexibility and ecosystem support necessary to navigate an increasingly complex technological landscape. The table below indicates how RISC-V addresses the challenges faced by organisations in the automotive sector.

How RISC-V Addresses Automotive Business Challenges

Business Challenge	Industry Context	RISC-V Solution	Strategic Implications
Supply Chain	• Platforms must remain viable for decades • Evolving auto supplier ecosystem with shifting value capture dynamics from the traditional OEM → Tier1 → Tier2 flow • Frequent acquisitions and geographic realignments	• Open standard with no single corporate ownership • Consistent Profiles ensuring binary compatibility across solutions • Cross-geographic ecosystem vendors and providers	• Reduces dependency risks • Enables flexible, long-term platform strategies • Supports multi-vendor collaboration
Cost Management	• Escalating material costs continue to be a major issue • Pressure to optimize resource utilization • Need for competitive procurement strategies • The constant drive for cost reduction	• Competitive vendor marketplace driving innovation • Modular design reduces unnecessary complexity and promotes software reuse • Smaller foundational technology footprint with no legacy software overhead	• Enables more granular cost optimization • Reduces per-unit technology expenses • Facilitates custom, right-sized solutions
Cybersecurity	• Increasing vulnerability to digital attacks • Requirement for long-term software maintenance • Need for consistent security across vehicle lifecycles • Complexity of maintaining software updates	• Standardized, consistent security architecture • Fundamentally secure design principles • Transparent, verifiable security model • Ecosystem tools to support verification and certification	• Provides robust protection mechanisms • Simplifies compliance and certification • Enables more predictable security maintenance • Reduces long-term security management costs
Electrification & Sustainability	• Energy-constrained battery-based systems • Intense focus on efficiency, weight, and cost • Performance requirements for electric vehicles	• Purpose-optimized component design • Lightweight architectural approach • Flexible instruction set for power optimization • Enhanced energy efficiency capabilities	• Directly improves battery range • Reduces overall system power consumption • Power efficiency enables more processing capability in a constrained environment

Nevertheless, many of the challenges faced by the automotive segment are common with or analogous to those faced when adopting AI technology in adjacent or even unrelated industry segments. The RISC-V ISA is applicable wherever AI may be enabled, including use cases in data centres, cloud, high-performance computing, consumer, embedded devices, and edge AI IoT.

A common RISC-V ISA permits users to share similar challenges and reuse development efforts across multiple industry segments. A vehicle no longer exists as a standalone 'device', to optimise the advantages of AI it must interact as a component of a much wider AI ecosystem. Such an ecosystem could include the cloud, service digital twins, road infrastructure IoT and/or related consumer devices. Deploying a single RISC-V ISA across the board enables a common software stack, promotes reuse, simplifies integration and enhances performance. Wider interdependencies across the RISC-V ecosystem can also reduce the development burden and crucially, improve the end-user experience through slicker integration. Member companies typically contribute to software offerings where they have expertise and are not obligated to contribute across the whole stack. This provides users with an improved choice of optimized solutions.

Software Implementation for RISC-V Automotive AI

The modular nature of RISC-V enables AI developers to deliver innovative software-focused AI systems tailored for automotive applications. Vendors are now creating comprehensive software stacks for RISC-V-based AI deployments.

Large language models (LLMs) have become essential to numerous applications thanks to their powerful capabilities in natural language understanding and text generation. However, their large model sizes and high computational demands present significant challenges for resource-constrained edge platforms. The ML compiler plays a crucial role in addressing this challenge, allowing seamless integration of LLMs into the diverse hardware found in modern vehicles.

Automotive systems cover a diverse range of performance, power and cost requirements — from standard microcontrollers to specialized AI accelerators. The ML compiler must be able to optimize models for this heterogeneous hardware landscape, offloading performance-critical operations to the most appropriate processing elements.

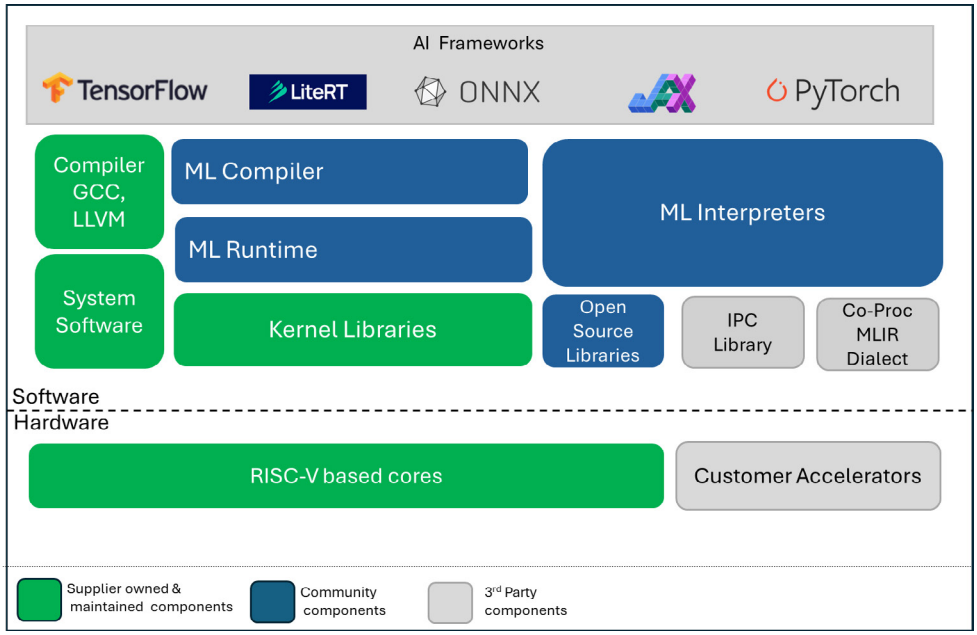

The **green blocks** represent the fundamental RISC-V building blocks typically provided by the silicon vendor:

- **Compiler:** Provides RISC-V C/C++ compilation and optimization, RISC-V Vector (RVV) intrinsic programming, auto-vectorization, and efficient RISC-V backend code generation for ML Intermediate Representation (MLIR).
- **System Software:** Linux or RTOS operating systems with KVM, Xen, and bare-metal Hypervisor virtualization
- **Kernel Libraries:** A highly optimized C/C++ library with a set of tuned routines that maximize algorithm throughput for a specific RISC-V processor. Some hot-spot operations in IREE will be offloaded to specific software libraries in order to maximize performance from specific RISC-V processors.

The **blue blocks** represent the components from Open Source projects. IREE is an open-sourced MLIR-based compiler and runtime.

ML Interpreters: For customers requiring a more lightweight framework, the following types of capabilities exist:

- Customized LiteRT with RVV optimizations.
- Upstream XNNPACK with RVV optimizations.
- ONNX Runtime and delegated to XNNPACK.
- Customized llama.cpp with additional RVV optimizations
- Other Open Source Libraries like libyuv, libc, OpenSSL, zlib-ng to accelerate other non-AI/ML domains in RVV and fine-tuned for specific processor microarchitectures.

The **gray blocks** represent the components provided by third-party vendors or communities, which customers can leverage to create a comprehensive solution.

Co-Proc MLIR Dialect: Allows users to target attached custom coprocessors with minimum effort.

RISC-V AI-Enabled Deployments in Automotive

The following section examines three case studies from industry leaders that illustrate the deployment of RISC-V automotive AI solutions in addressing real-world challenges. Each case details a specific customer problem, outlines the strategic deployment of RISC-V technology to address it, and highlights the tangible outcomes and benefits achieved.

Mobileye adopts RISC-V to Enable Advanced Features and Lower Development Cost

Modern vehicles are gaining ever-increasing levels of self-driving ability from the integration of advanced driver assistance systems (ADAS). Mobileye is a successful provider of ADAS, delivering solutions to Automotive OEM vendors such as Volkswagen, Audi, and Porsche. Mobileye utilizes a combination of artificial intelligence (AI), computer vision, localization, mapping, and machine learning technology in the development of its SoCs. The features offered can vary by car model and manufacturer, ranging from basic functions such as lane departure warnings, forward collision alerts, automatic emergency braking, and traffic sign recognition to more advanced features like adaptive cruise control, lane change assistance, and traffic jam or highway assist. Vendors and OEMs need to be able to introduce more advanced semi or fully autonomous driving assistance functionality, which requires the implementation of sophisticated AI algorithms.

At CES 2025, Mobileye presented their vision of moving from simpler ADAS systems that assist the driving experience to those that will enable the driver to be 'hands-off' the wheel, 'eyes-off' the road for at least part of the journey. Those systems will be based on the RISC-V ISA together with Mobileye's special AI technology. The Mobileye RISC-V system is based on the MIPS P8700 core, bringing high-performance, low-power compute and taking advantage of the RISC-V CMO that enables enhanced hardware-based data movement between the CPU and AI engines. Mobileye reports that implementing RISC-V will enable it to reduce deployment costs, as a single system can replace the multiple systems that are currently deployed in the field.

The switch to a common open ISA will also enable their customers to accelerate their ADAS software development, complemented by their use of automotive-safe Linux. Furthermore, applying automotive safety qualification processes to open-source software such as Linux and LLVM improves the safety approach and will further benefit from multiple entities all independently reviewing the code.

Mobileye's RISC-V based products will be rolled out from 2027 to 2029 to replace the existing solutions. They expect that by 2030 the majority of their ADAS systems will be RISC-V-based.

Customer Problem	How RISC-V helped	Outcome/Benefits
Mobileye required high-performance low power compute capacity to implement a new ADAS solution with 'Hands off/Eyes off' capabilities	They implemented the RISC-V-based MIPS P8700 core that improves performance by enabling enhanced hardware-based data movement between the CPU and AI engines	• Reduce deployment costs by replacing multiple systems with a single system • Open ISA with auto Linux accelerated software development effort. • Improved safety qualification due to more entities verifying software

How SiFive Addresses the Challenge of Automotive Data Handling

One of the overarching trends underway in the automotive industry is a shift to zonal architectures, where the information from multiple sensors in a particular area of a vehicle are aggregated and, where appropriate, sent on to other computer systems across the Ethernet-based network.

The amount of data produced in new vehicles is enormous, and therefore, zonal systems must sift through the raw data and only transmit useful data across the network.

In this application, a tier 1 car manufacturer was developing a sensor fusion platform that reviewed information from several sensors, including lidar, cameras, and radar. The company has created its own hardware accelerator but needed an engine that would feed the data into the accelerator in the form of wide vectors.

The customer decided to use RISC-V due to its open architecture and the availability of industry-standard development tools This meant that the customer could create their own variant for future designs or continue to select from a diverse set of software-compatible options from industry suppliers.

The customer selected the X280 Intelligence Core from SiFive. This processor was selected due to its support of wide vectors (512 bits), a tightly coupled port used for communication with the hardware accelerator, and a high-bandwidth memory port that could be used to share data between the accelerator and the X280.

At the time of writing this paper, this system is in development and is expected to appear in vehicles in the 2028/2029 timeframe.

Customer Problem	How RISC-V helped	Outcome/Benefits
Tier 1 auto vendor required an efficient engine that would feed data into the hardware accelerator in the form of wide vectors	They selected the SiFive X280 Intelligence Core due to its wide vector support, its tightly coupled communication port and high bandwidth memory port	• Industry standard development tools and open architecture allowed them to create their own variants • Could select from diverse set of software options from ecosystem

Siemens Drive Automotive RISC-V Adoption and Development with Dedicated EDA tools

Electronic Design Automation (EDA) comprises a suite of software, tools and related services that enable users to accelerate the design, modelling and verification of their electronic system developments. The specific intersection of mechanical and electronic technology found in the automotive segment, plus enhanced safety requirements, calls for advanced AI-enabled EDA tools that optimize development resources and ensure safety verification.

EDA vendors have recognized the adoption of RISC-V and are ensuring that there are specific tools and flows developed to support the full RISC-V automotive ecosystem. These can range from **comprehensive verification solutions** to **complementary IP** that can form part of a particular RISC-V implementation.

Tools such as Siemens EDA Ultrasight-V provide a supporting solution that can add advanced debug and trace capabilities to a RISC-V implementation. This capability, based on the E-trace standard, enables full debugging to assist during the software development phase of a project.

New **AI-enabled digital twin environments** are increasingly becoming fundamental to the development of RISC-V automotive systems. These allow the RISC-V implementation to be virtually modelled alongside many other critical components of the vehicle, such as the electrical sensors and monitors.

For automotive solution vendors, deploying a suite of custom AI-enabled tools can reduce errors, cut development costs, and accelerate total time to market. Crucially they also enable the successful verification and certification of designs to ensure the critical safety and security compliance requirements of the automotive sector.

Customer Problem	How RISC-V helped	Outcome/Benefits
To successfully adopt RISC-V based designs for Automotive applications, complementary EDA tools must be available	EDA tools and flows such as Siemens EDA Ultrasight-V developed to specifically support the full RISC-V automotive ecosystem.	• Accelerated software development and Time To Market • Digital Twin environments to virtually model whole auto systems including RISC-V • Enhanced verification and certification of critical safety and security design aspects

Regulatory Aspects

The safety-critical nature of the automotive sector dictates that systems must meet stringent regulation and rigorous certification approvals before they can be deployed in a vehicle. Furthermore, sophisticated Level 3+ and 4 ADAS systems increasingly remove the need for driver input and intervention, moving closer to the possibility of fully autonomous self-driving vehicles enabled by AI/ML.

Multiple vendors have now announced RISC-V-based automotive-specific processors that are fully compliant with Functional Safety standards ISO26262 ASIL-B and ASIL-D. Similarly, the introduction of RISC-V automotive-specific EDA tools allows developers to verify designs and model long-term in field scenarios through digital twin simulation (as referenced in the EDA tools case study above.)

Optimize Automotive AI/ML Deployments with our Recommended Best Practices

1. **"AI/ML Everywhere";** With very few exceptions, plan for some form of AI/ML capability in all systems in a vehicle. This means providing "headroom" to support the implementation of new functionality, the secure delivery of that functionality and an ongoing ability to update the capability of systems that have been deployed
2. **"Software-driven" perspective;** Develop a holistic solution enabled by a standard ISA that allows co-design of the software stack alongside the hardware. Focused on the creation of optimized architectures that are driven by the concrete needs of AI applications, different usage scenarios (e.g., training vs. inference; e.g., IoT vs. datacenter), and with an emphasis on current and evolving AI/ML algorithms.
3. **"The only constant is change";** We have learned over decades that during periods of rapid innovation, a more modular approach is required to track standards. In the area of AI, CNN models were all the rage a few years back. Now there has been a swift shift to transformer models. Yet, industry experts are predicting that those will be replaced in relatively short order. Running too much of the algorithm in software on a general-purpose CPU risks an implementation that is suboptimal in terms of performance, power and cost. Building an ASIC (application-specific integrated circuit) runs the risk of being outdated by the time the vehicle is driven away by the new owner.
4. **"Add resiliency into the supply chain";** While the technology itself is important, even more important is the associated business model that enables the technology to
   - Track the quickly evolving needs of this rapidly expanding market sector, including the delivery of new functionality
   - Enable deep collaboration across different areas of the value chain
   - Empower innovation by harnessing the best ideas from anywhere in the community—even from those outside traditional semiconductor companies.
5. **"A car is not a cloud";** The power, space, cost and reliability constraints of a vehicle are such that there are some very specific tradeoffs to be made as to where processing is performance, where models are located and how they are updated. As an example, in Tesla's earlier autonomous systems, many tasks were processed on the CPU and GPU, leading to performance and power limitations. Tesla's approach to addressing this problem was to develop a custom AI accelerator specifically for neural network processing within the FSD Computer, introduced in 2019. This dedicated chip offloads heavy AI tasks from the CPU, improving real-time performance and significantly reducing power consumption compared to previous setups. This custom accelerator allowed Tesla to reduce costs, improve the efficiency of autonomous driving functions, and make real-time processing more viable—demonstrating the benefits of using AI accelerators over CPUs for demanding AI tasks in automotive systems.

Conclusion

A successful and safe deployment of AI in the Automotive segment is not restricted to the physical processor implementation, it relies on other adjacent factors and the support of integrated partners. The RISC-V ISA's flexible and modular characteristics provide the ideal basis for AI-enabled hardware that supports advanced ADAS functionality while allowing OEMs the choice to differentiate their offering. A successful implementation requires the support of a wider ecosystem plus the ability to align with industry-standard frameworks and certifications.

The RISC-V AI proposition comprises co-designed hardware and software, complementary EDA tools, and a flexible, scalable ISA to readily address automotive-specific requirements and business challenges. Just as the AI-enabled vehicle is now a component of a connected system, the RISC-V implementation is stronger as part of its wider ecosystem.

Appendix

What is RISC-V?

Stemming from what was initially scoped as a summer project, RISC-V was invented in 2010 at the University of Berkeley in California. One of the originators of the technology was David Patterson, a co-author of what is recognized to be THE definitive book about computer architecture both at the time and for decades since.

As a new processor architecture, there was an opportunity to harness modern knowledge and zero need to be concerned about legacy. The result was a very simple base instruction set architecture (ISA). In chip terms, this means that the minimal RISC-V functionality can be implemented in silicon in a smaller size than incumbent architectures. The ISA was designed to be modular in form, enabling the extension of the base ISA to support new functions. While this provides the opportunity to drive down cost, the architecture offered two unique benefits

Open standard

Similar to standards like Ethernet or WiFi, RISC-V was established to be a standard that no one company owned/controlled. This offers the opportunity to reduce vendor lock-in and help companies be more immune and resilient to changes in the geopolitical, market and enterprise landscapes.

Customization

Everyone is empowered to define, implement and deploy processors with custom instructions, opening up a new avenue of exploration and innovation. Companies, individuals and academic institutions can use these optimized-for-purpose enhancements to address specific system challenges, including cybersecurity, performance and power.

The subject of “customization” raises the topic of fragmentation, in particular with respect to the software ecosystem.

The “so what” from this approach is that the onus is then on that company or individual to deliver the software that supports these customizations. This approach has been successfully implemented by a number of companies, especially for black box applications where the provider of the SSD delivers all of the software that runs on that platform. What this means is that customizations to accelerate specific functions are not exposed to third-party programmers and programs.

RISC-V also enables companies to build a tightly controlled architecture, featuring (by design) mandated instructions and very few options. For companies that are building a core that is going to run tens of thousands of software packages, this rigidity is needed. The way this is achieved in the RISC-V community is with Profiles. The recent RVA23 Profile mandates a number of capabilities, including vector extensions, hardware virtualization and cryptography support Fundamentally, RVA23 aligns implementations of RISC-V 64-bit application processors that will run rich operating systems (OS) stacks from standard binary OS distributions.

Profiles are planning to cover other areas of the market, including microcontrollers.

At the time of writing this paper, billions of system-on-chip (SoC) components have been deployed across a range of markets, primarily in embedded processing systems. With the knowledge of RISC-V implementations increasing and software ecosystems for a number of markets starting to form, RISC-V technology is starting to be deployed across a range of additional markets and applications, including automotive, consumer, datacenters, networking and space. Sizing the future growth is particularly challenging, but one market analyst, SHG Group, has predicted that in the year 2031, more than 25% of SoCs that ship will include RISC-V technology.

From an AI/ML perspective, companies including Google, Meta, NVIDIA have centered on RISC-V as a foundational building block for their developments in this area, driven by a combination of the business and technology benefits this architecture and supporting ecosystem provide.