European Connected and Autonomous Vehicle Alliance

and

European Alliance on Processors & Semiconductor Technologies (ALLPROS.EU)

Technology Roadmap

Automotive Hardware Platform

October 12th, 2026



Commission européenne/Europese Commissie, 1049 Bruxelles/Brussel, BELGIQUE/BELGIË – Tel. +32 22991111
Office: L-51 09/DCS – Tel. direct line +32 229-62469

cnect-EU-ECAVA@ec.europa.eu

TECHNOLOGY ROADMAP
European Connected and Autonomous Vehicle Alliance (ECAVA)
Automotive Hardware Platform

Executive Summary

Strategic Political & Stakeholder Imperatives

The development and scaling of a competitive and sovereign European automotive semiconductor ecosystem demands **extreme investments, and no single player or member state in Europe is likely to succeed alone**.

Rather than focusing on isolated hardware development, Europe must drive the **fast vertical integration of technology together with the Automotive OEMs**. Tight, early-stage collaboration between semiconductor innovators and automotive OEMs is critical to accelerate the cycle from processor and chiplet design to deployment. A deep vertical alignment is needed to ensure that our hardware architectures are custom-tailored to the real-world demands of future software-defined vehicles (SDVs) and can be commercialized rapidly.

Strategic Context & the Architectural Shift

The global automotive industry is in the middle of a profound transformation from hardware-centric, decentralized architectures toward centralized, **Software-Defined Vehicles (SDVs)**. This shift is heavily driven by the intense processing demands of Advanced Driver-Assistance Systems (ADAS), autonomous driving (L4/L5), and rich, immersive in-vehicle infotainment (IVI).

At the center of this revolution sits the real-time and high-performance computing **System-on-Chip (SoC)**. However, the traditional approach of using massive, monolithic silicon dies is colliding with physical limits, thermal barriers, and soaring R&D costs. To overcome these limitations, bypass single-source non-EU dependencies, and retain control over the E/E architecture, the European ecosystem must drive **heterogeneous integration via modular chiplets**, which prioritize scalability, flexibility, system-level latency reduction and interconnect efficiency over simple geometric transistor shrinking.

Core Technological Pillars

To achieve technological sovereignty and scalability, the European roadmap establishes three concurrent strategic pillars:

- **Sovereign Processor Architectures and Software Ecosystem:** Europe should follow a maturity-based and technology-open pathway. Near-term automotive HPC platforms should leverage production-ready commercial processor IP and software ecosystems, including Arm, while Europe develops its own system-design, safety, security, AI acceleration, and advanced-node capabilities. In parallel, automotive-grade RISC-V processor IP, software, tools, and certification assets should be matured to enable progressive substitution in later platform generations.
- **High-Performance Computing (HPC) & AI Acceleration**: The central compute platform should be architecturally decomposed into three distinct components to enable multi-sourcing and avoid bundled-vendor roadmaps:
    1. *AI Accelerator (AIA)*: Run perception (CNNs, Transformers, VLMs) and planning workloads.
    2. *Application Processor (AP)*: Run the SDV middleware and general-purpose workloads on Arm application cores and in the long-run on RISC-V based application cores.
    3. *Safety Island (SI)*: Deliver independent, ASIL-D-compliant system monitoring and fail-operational safety.
- **Chiplet Integration**: Heterogeneous integration via modular chiplets empowers OEMs with scalability, strategic flexibility, and accelerated time-to-market. By scaling compute chiplets on a standardized, secure base die, OEMs can customize performance across entire vehicle fleets. This open framework eliminates vendor lock-in, enabling rapid, cost-effective hardware upgrades without complete monolithic redesigns. Open and ISA-agnostic chiplet interfaces shall allow processor technologies to evolve independently from the overall platform, enabling the progressive introduction of European RISC-V components without redesigning the complete system.

Strategic Assessment of European Control Points

Europe must leverage its engineering strengths while aggressively addressing its gaps in high-performance silicon and packaging:

Areas of Existing Strength (EU to Lead)	Key Competency Gaps (EU to Build/Source)
System Knowledge & E/E Architecture: Defining top-down requirements for E/E data flows.	**High-Performance Compute (HPC) & AI**: Cutting-edge CPUs, GPUs, AI accelerators, and memory cells.

Functional Safety & Security (ASIL): Sovereign definition of Safety Islands and real-time control.	**SoC System Design & Integration**: At-scale experience in multi-billion-transistor complexities.
In-Vehicle Network Technology: Traditional leadership in CAN, Automotive Ethernet, and I/O.	**Advanced-Node Chip Design**: Design service expertise at leading-edge nodes (5nm, 3nm, and beyond).
Standard Packaging & Reliability: Packaging for power electronics, sensors, and microcontrollers.	**Wireless Connectivity**: IP for 5G, Wi-Fi 6/7, and V2X communication.
NoC & Real-Time Cores: Routing data on-chip and open RISC-V R-cores.	**Advanced Packaging Technology**: 2.5D/3D stacking and high-density fan-out (currently concentrated in Asia).
Defining Scalability Demands: Creating cross-brand, cross-segment scaling platforms.	**EDA Tooling & Ecosystem**: High dependency on a small oligopoly of non-EU software design tools.

Ecosystem Collaboration: The "Europe Inc." Philosophy

To tackle complexity and prevent fragmented duplication, European automotive players must move away from traditional, transactional "black box" customer-supplier contracts. The industry must collaborate in a pre-competitive, antitrust-compliant partner network—mirroring the execution speed of a virtual single enterprise while strictly separating technical standardization from individual downstream commercial competition:

- **Unified Goals & Shared Success**: Agreeing on joint technical goals, such as target performance baselines, power efficiency limits, or interoperability standards across partners.
- **Radical Transparency**: Exchanging non-commercially sensitive engineering data (e.g., thermal simulation methodologies, interface protocol specs, and safety/security testing criteria), and early-stage integration challenges in real-time.
- **Integrated, Cross-Functional Teams**: Establishing multi-party engineering networks (OEMs, Tier-1s, and semiconductor providers) working off a single project backlog to solve pre-competitive technical challenges.
- **Ecosystem-Level IP Governance**: IP protection operating on FRAND (Fair, Reasonable, and Non-Discriminatory) terms, multi-party licensing to standard-essential technologies, and "clean-room" secure data-sharing environments proactively at the overarching alliance level to eliminate legal friction.
- **Maximizing Joint Undertakings**: Using established frameworks like the **Chips Joint Undertaking (CHIPS-JU)** as a foundation, while actively steering and tracking operational inter-company dependencies to maximize the return on public and private investments.

Concurrent Roadmap Phases

To accelerate deployment, the roadmap runs three parallel, interconnected phases concurrently, creating a substantial technological basis while fostering vertical integration of technology to secure the demand:

Phase	Time Horizon / SOP	TRL Focus	Strategic Objective
Build the Foundation (Gen N+3)	SOP ~2037	TRL <5	Long-term fundamental research (e.g., neuromorphic, photonics), setting standardizations and academic baselines.
Prototypes for Next Generation (Gen N+2)	SOP ~2034	TRL 3–7	Developing physical Integration Platforms and test-silicons to deliver functional proof points.
Empower & De-Risk Industrialization (Gen N+1)	SOP ~2031	TRL 8–9	Fostering qualification and certification of automotive-grade processors and chiplets from EU suppliers and establishing resilient, multi-vendor series production supply chains.

Key Message

Collaboration works where it is easiest (foundation), weakens where integration is required (prototyping), and is largely absent where systemic alignment towards industrialization is critical (empower & de-risk).

Bridging these layers is the key to industrialization towards achieving a European SDV Automotive Hardware Platform.

Contributors to the ECAVA and ALLPROS.EU Automotive Hardware WG

ORGANISATION	FIRST NAME	FAMILY NAME
Aeneas	Vincent	**LE MEAU**
Ampere	Florence	**ROWE**
AUMOVIO	Markus	**BEBENDORF**
Axelera	Manuel	**BOTIJA**
Axelera	Pavel	**ZAYKOV**
Black Semiconductor GmbH	Ali	**SAFIEI**
BMW Group AG	Holger	**VÖLKEL**
Bosch	Joern-Ole	**GODBERSEN**
CEA	Magaly	**GOUTTEBROZE**
CLEPA	David	**STORER**
EC	Panagiotis	**TSARCHOPOULOS**
EPoSS e.V.	Gereon	**MEYER**
IMEC	Serena	**FARUQUE**
Infineon	Konstantin	**KONRAD**
Infineon	Jochen	**KOSZESCHA**
Infineon	Thomas	**MARKTL**
OpenChip	Gilberto	**RODRIGUEZ**
Renault/Ampere	Cathy	**BATISSON**
Schaeffler Vitesco	Stefan Felix	**HARTL**
Stellantis	Misaele	**MARLETTI**
STMicroelectronics	Florian	**BAUMANN**
STMicroelectronics	Fabrice	**GRAIGNIC**
STMicroelectronics	Cosimo	**MUSCA**
Valeo	David	**ROINÉ**
Volkswagen	Peter	**FIEDLER**
Volkswagen	Jonas	**HENNIG**
Volvo Trucks	Vishnu Varthan	**KUMARESAN**
Volvo Trucks	Magnus	**STÅLESJÖ**
Volvo Trucks	Mauro Cesar	**ZANELLA**
ZF Friedrichshafen AG	Franz	**AUBECK**

Co-Chairs of the ECAVA and ALLPROS.EU Automotive Hardware WG

ORGANISATION	FIRST NAME	FAMILY NAME
Bosch	Thomas	**SCHAMM**
NXP	Patrick	**PYPE**

Disclaimer

This document is a preliminary scoping paper provided for discussion purposes only. It does not constitute a legally binding commitment or representation of any kind. All content is subject to change, revision, or withdrawal without notice. No stakeholder should rely on this document for legal, financial, or operational decisions. The information and views set out in this draft roadmap do not necessarily reflect the official opinion of any ECAVA member. Neither the Commission nor any ECAVA member may be held responsible for the use which may be made of the information contained therein. Comments and suggestions for improvement can be sent to CNECT-EU-ECAVA@ec.europa.eu.

All activities undertaken in the context of this roadmap, including any discussions, information exchanges, or collaborations to which it may give rise, shall be conducted in full compliance with applicable EU and national competition rules. Participants are reminded that exchanges of commercially sensitive information are prohibited, whether occurring in formal or informal settings. The involvement of the European Commission in the preparation of this document does not exonerate participants from their individual obligations under competition law. For full guidance, participants should refer to the Competition Law Compliance Guidelines set out in Annex II of the ECAVA Declaration.

Table of Contents

1. STRATEGY FOR EUROPEAN SDV HARDWARE PLATFORM – DRIVING EU-OEM VALUE CREATION 8

1.1. Challenges in Current Automotive E/E Architectures 8

1.2. Evolution Towards Software-Defined Vehicles (SDVs) 10

1.3. Evolution of Vehicle Compute Platforms 15
1.3.1. Target SoC Capability 2031/32 15
1.3.2. Vision 2035+: Future Vehicle Compute 15
1.3.3. Vision 2040+: Future++ Vehicle Compute 16

1.4. Components of a European SDV Hardware Platform 17
1.4.1. Key Deficiencies of Monolithic Automotive SoCs 18
1.4.2. Key Deficiencies of Board level Integration 18
1.4.3. Chiplet Systems – A 'Best of Both Worlds' for Automotive Compute 19

1.5. Strategic Imperatives for Chiplet Adoption, Driven by OEM Needs 21

1.6. Overview of control points needed inside Europe 22
1.6.1. Processor Architectures (e.g., RISC-V for Automotive) for European Competitiveness 23
1.6.2. Specialized Accelerator and HPC Chiplets (AI/ML, DSP) for ADAS and Autonomous Driving ... 24
1.6.3. Chiplet Integration: A Strategic Pillar for European Automotive Sovereignty 26
1.6.4. Further Semiconductor Demand 28

2. ECOSYSTEM COLLABORATION: STRENGTHENING EUROPEAN COMPETITIVENESS 29

3. ROADMAP PHASES AND KEY MILESTONES ENABLING TECHNOLOGY INTRODUCTION IN EUROPE 32

3.1. Roadmap Phases in the context of ECAVA 32
3.1.1. Build the Foundation [Target: N+3, Vehicle SOP ~2037; TRL <5] 33
3.1.2. Prototypes for next generation [Target: N+2, Vehicle SOP ~2034; TRL 3-7] 33
3.1.3. Empower & De-risk industrialization [Target: N+1; Vehicle SOP ~2031; TRL 8-9] 33

3.2. RISC-V Roadmap 34
3.2.1. Build the Foundation 34
3.2.2. Prototypes for next Generation 35
3.2.3. Empower & De-risk 37

3.3. HPC Roadmap 38
3.3.1. Build the Foundation 38
3.3.2. Prototypes for next Generation 38
3.3.3. Empower & De-risk 39

3.4. Chiplet Integration Roadmap 40
3.4.1. Build the Foundation 40
3.4.2. Prototypes for Next Generation 41
3.4.3. Empower & De-risk Industrialization 42

4. ANNEXES 43

4.1. ANNEX: RISC-V – Detailed Roadmap Document 43
4.1.1. Implementation Status of the Automotive Roadmap 45
4.1.2. Scalable RISC-V Automotive Control Processors 47
4.1.3. High-Performance RISC-V based Application Processors 49

4.1.4. AI and Compute Accelerators 50
4.1.5. System Integration and Interfacing 53
4.1.6. Milestones 55

4.2. ANNEX: Chiplet Integration – Detailed Gap Analysis 59
4.2.1. Build the Foundation 59
4.2.2. Prototypes for next Generation 64
4.2.3. Empower & De-risk 67

1. Strategy for European SDV Hardware Platform – Driving EU-OEM value creation

The Paradigm Shift in Automotive E/E Architectures

The automotive industry is in the middle of a significant transformation, moving away from traditional, decentralized electronic/electrical (E/E) architectures towards a centralized, software-defined vehicle (SDV) model. This paradigm shift is primarily driven by the increasing demand for advanced driver-assistance systems (ADAS), in-cabin infotainment, and autonomous driving capabilities. At the heart of this revolution lies the central computing System-on-Chip (SoC), high-performance processors that acts as the brain of the modern vehicle, consolidating various functions and enabling a more integrated and flexible system design.

The shift from a hardware-defined to a software-defined vehicle represents a fundamental change in where the value of a car is created. Historically, automotive Original Equipment Manufacturers (OEMs) relied on Tier 1 suppliers to provide "black box" ECUs, which were then integrated into the vehicle. Today, OEMs are aggressively moving to define their own E/E architectures from the ground up. This strategic shift is not merely a technical decision, but a core business imperative driven by several key factors.

By taking ownership of the architecture, OEMs regain control over the vehicle's core functionality and future potential, transforming it from a static piece of hardware into a dynamic, upgradable platform.

1.1. Challenges in Current Automotive E/E Architectures

Core Motivations for OEM-Defined E/E Architectures

Strategic Driver	Description
Brand Differentiation & User Experience (UX)	The E/E architecture dictates the entire digital personality of a vehicle. By controlling it, OEMs can craft a unique and consistent user experience across their model lineup. This includes the responsiveness of the infotainment system, the seamless operation of ADAS features, and the overall digital ambiance in the cabin, which are becoming primary differentiators for customers.
Enabling New Revenue Streams	A centrally managed, software-defined architecture is the foundation for future business models. It allows OEMs to offer *Features-on-Demand (FoD)*, where customers can purchase or subscribe to new functionalities (e.g., advanced autonomous features, enhanced performance modes) over the air, creating a continuous revenue stream long after the initial sale.
Faster Innovation & Agility	Defining the architecture allows OEMs to decouple hardware and software development cycles. With a stable hardware platform and standardized APIs, software teams can rapidly develop, test, and deploy new features and updates via *Over-the-Air (OTA)* updates. This dramatically reduces time-to-market and allows the vehicle to improve and adapt throughout its lifespan.
Cost Optimization & Complexity Reduction	Moving from a complex web of hundreds of individual ECUs to a more centralized or zonal architecture simplifies the vehicle's physical layout. It reduces the amount of wiring, cuts down on the number of required processors, and allows for economies of scale through the use of common, high-performance computing platforms across different vehicle models.
End-to-End Control over Safety & Cybersecurity	As vehicles become more connected and autonomous, ensuring safety and protecting against cyber threats is paramount. A fragmented architecture with numerous supplier-provided black boxes creates multiple potential points of failure and vulnerability. By defining the entire architecture, OEMs can implement a holistic, top-down security and safety strategy, controlling data flows and access points across the entire system.
Mastery of the Supply Chain	Taking ownership of the E/E architecture shifts the balance of power back to the OEM. Instead of being dependent on a supplier's proprietary solution, the OEM can specify the exact hardware and software interfaces. This allows them to source components from a wider range of suppliers, fostering competition, reducing the risk of supply chain bottlenecks, and gaining more direct control over key technologies like semiconductors.

In essence, the E/E architecture has become the strategic centerpiece for the future of mobility. For OEMs, defining it is not just about designing a better car; it's about owning the customer relationship, unlocking new value, and securing their role as the ultimate integrator in an increasingly software-driven world.

For today's automotive OEMs, the E/E architecture has become the new engine—representing the core value, brand identity, and competitive advantage that the combustion motor once solely provided—and for this reason, they must gain direct influence and ownership over the central hardware building blocks that form its foundation.

> **THE CENTRAL COMPUTING SOC: THE HEART OF THE MODERN VEHICLE**
>
> The central computing SoC is the cornerstone of the new E/E architecture, enabling the integration of multiple, previously separate, functionalities into a single, powerful chip. This consolidation of domains like ADAS, infotainment, and vehicle control into one SoC brings numerous benefits, including reduced complexity, lower cost, and a smaller physical footprint. However, this high level of integration also presents significant challenges in terms of performance, safety, and security.
> While central computing SoCs are the key enablers for the modern software-defined vehicle, the current approach to their design—dominated by large, monolithic silicon—presents significant and growing challenges for the automotive industry. These monolithic chips, which attempt to integrate every conceivable function onto a single piece of silicon, are becoming bottlenecks to innovation, cost-efficiency, and strategic flexibility. This realization is driving the industry to explore more modular and scalable alternatives.

Key Challenges and Technological Considerations for Automotive SoCs

The successful transition to a centralized E/E architecture hinges entirely on the capabilities and design of the automotive System-on-Chip (SoC). While these powerful processors are the key to unlocking new levels of integration and performance, their implementation introduces a host of complex technical challenges that must be navigated.

Strategic Rationale for a European SDV Hardware Platform

The development of a European Software-Defined Vehicle (SDV) Hardware Platform is a strategic imperative to ensure the competitiveness of the European automotive industry. Such a platform would provide a common, open, and scalable architecture that would enable European automakers and suppliers to:

- *Accelerate innovation:* By providing a standardized platform, developers can focus on creating new features and applications, rather than reinventing the underlying hardware and software infrastructure.
- *Reduce development costs:* A common platform would reduce the need for each company to develop its own proprietary solutions, leading to significant cost savings.
- *E/E Architecture Long-Term Stability:* Certain core features and functionalities within the E/E architecture remain stable over extended lifecycles; altering or removing these established baselines would trigger unnecessary redevelopment costs and significantly delay time-to-market.
- *Foster collaboration:* An open platform would encourage collaboration between different companies, leading to a more vibrant and innovative ecosystem.
- *Ensure technological sovereignty:* A European-led platform would reduce the dependence on non-European technology providers, ensuring that the European automotive industry remains in control of its own destiny.

However, realizing this vision of sovereign, high-performance central computing platforms is an incredibly resource-intensive endeavor. Designing and scaling advanced, automotive-grade silicon architectures and their accompanying software ecosystems demands extreme investments that far exceed the capabilities of individual organizations. In a global landscape dominated by highly consolidated technology ecosystems, no single player or member state in Europe is likely to succeed alone. To bridge the massive gap in capital, infrastructure, and specialized talent, European automotive stakeholders must move beyond traditional bilateral relationships and commit to a highly coordinated, cross-border development strategy. Only by pooling resources and co-investing in shared, open technical baselines can European industry partners mitigate these colossal financial risks and build a robust, globally competitive foundation for the software-defined era.

By investing in a European SDV Hardware Platform, Europe can position itself as a leader in the next generation of automotive technology, creating new economic opportunities and ensuring the long-term success of its automotive industry.

1.2. Evolution Towards Software-Defined Vehicles (SDVs)

This section delves into the critical technological considerations and hurdles that engineers and strategists face, from ensuring functional safety in highly integrated designs to managing the immense data throughput from sensors and meeting the ever-increasing demands for computational power. Each consideration represents a crucial piece of the puzzle in building the robust, reliable, and powerful brain of the future software-defined vehicle.

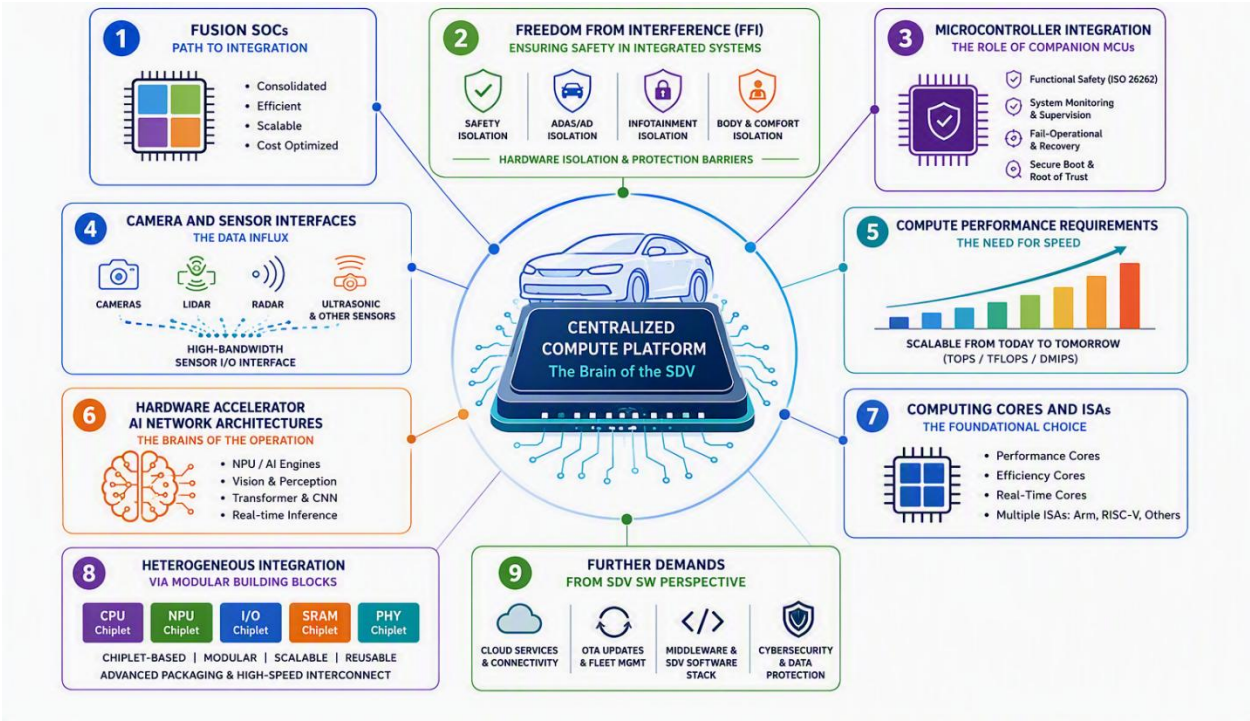


Figure 1: The main identified pillars of the Automotive Hardware Platform enabling Software-Defined Vehicles (SDVs).

1. Fusion SoCs: The Path to Integration

A "Fusion SoC" is a term for a new class of automotive processors that combine multiple functions, that were traditionally handled by separate Electronic Control Units (ECUs), into a single chip. These SoCs are designed to handle mixed-criticality workloads, where safety-critical functions (like ADAS) run alongside non-safety critical ones (like infotainment). Furthermore, the integration of third-party applications as 'black boxes' inside the infotainment operating system introduces security challenges, demanding strict separation and shielding to isolate potentially untrusted code from core vehicle functions. This integration is made possible by advanced hardware and software architectures that ensure different applications can share resources without interfering with each other.

Benefit	Description
Reduced Complexity	Fewer ECUs and less wiring simplify the vehicle's E/E architecture.
Lower Cost	Consolidation of hardware leads to a reduction in the bill of materials (BOM).
Enhanced Performance	Centralized processing allows for more efficient data sharing between different functions.

2. Freedom from Interference (FFI): Ensuring Safety in Integrated Systems

Freedom from Interference (FFI) is a critical safety concept in automotive design, ensuring that a fault in a less critical software component does not cascade and cause a failure in a more safety-critical one. In a mixed-ASIL (Automotive Safety Integrity Level) environment, where software components with different safety ratings coexist on the same SoC, FFI is paramount.

FFI is achieved through a combination of hardware and software mechanisms, including:

- *Hypervisors:* A software layer that creates isolated virtual machines, allowing multiple operating systems to run on the same hardware without interfering with each other.

- *Memory Protection Units (MPUs):* Hardware units that control access to memory, preventing unauthorized access and data corruption.
- *"Safety Islands":* Dedicated on-chip hardware, often including a separate processor core and memory, that is isolated from the rest of the SoC and is responsible for safety-critical tasks.

3. Microcontroller Integration: The Role of Companion MCUs

While powerful SoCs are designed for high-performance computing, they are not always optimized for the real-time control and safety-critical tasks required in a vehicle. This is where companion microcontrollers (MCUs) play a crucial role. These MCUs work alongside the main SoC, handling tasks such as:

- *Real-time control:* Managing time-sensitive functions like engine control and braking.
- *Safety management:* Monitoring the system for faults and taking corrective action.
- *Communication gateway:* Handling communication between the in-vehicle network and the SoC.

The key difference between an SoC and a companion MCU lies in their design focus. SoCs are optimized for performance, while MCUs are designed with safety and reliability as the top priority.

4. Camera and Sensor Interfaces: The Data Influx

Modern vehicles are equipped with an ever-increasing number of sensors, including cameras (visible and infrared), radar, and LiDAR, to provide a 360-degree view of the environment under all driving and weather conditions. This massive influx of data requires high-bandwidth, low-latency interfaces to transmit the information to the central SoC for processing.

Key interface technologies include for example:

- *MIPI CSI-2:* A widely used standard for connecting cameras to processors, offering high bandwidth and low power consumption.
- *Automotive Ethernet:* A high-speed networking standard that can handle the high data rates required for sensor fusion and other advanced applications.
- ASA-ML: ASA Motion Link is a cost efficient open standard of the Automotive SerDes Alliance (ASA) for high-speed data transmission in vehicles. It connects cameras, sensors and displays with the SoCs in ADAS or IVI applications.

5. Compute Performance Requirements: The Need for Speed

The computational demands of ADAS and autonomous driving are immense, requiring processors that can perform trillions of operations per second (TOPS). To meet these demands, automotive SoCs employ heterogeneous computing architectures that combine different types of processing units:

- *CPUs (Central Processing Units):* For general-purpose computing and task management.
- *GPUs (Graphics Processing Units):* For parallel processing and rendering of graphics.
- *AI Accelerators:* Specialized hardware for running artificial intelligence and machine learning algorithms.

6. Hardware Accelerator AI Network Architectures: The Brains of the Operation

AI and neural networks are the cognitive engine of the modern vehicle, but "AI" is not a single entity. It comprises a diverse and rapidly evolving set of network architectures, each with unique computational demands. The central challenge for SoC design is creating hardware accelerators—like Neural Processing Units (NPUs)—that can efficiently execute a mix of established and emerging models. An accelerator optimized for one type may be inefficient for another, making flexibility a critical design requirement.

The ability to process these varied network types efficiently is what allows the vehicle to perceive its environment, interact intelligently with its occupants, and make complex driving decisions.

Neural Network Type	Automotive Application	Hardware Accelerator Implications
Convolutional Neural Networks (CNNs)	*The Workhorse of Perception:* CNNs are the established standard for processing visual data. They excel at tasks like object detection (identifying pedestrians, cars, traffic signs),	Most current NPUs are highly optimized for the convolutional and pooling operations central to CNNs. The challenge is to maintain this high

	lane detection, and segmenting camera, radar, and LiDAR data.	efficiency while adding flexibility for other model types.
Transformers & Large Language Models (LLMs)	*The Heart of the In-Cabin Experience:* Transformers and LLMs are revolutionizing natural language understanding. In the car, they power sophisticated voice assistants that can understand complex, conversational commands, moving far beyond simple "play music" instructions. They can also be used for more advanced scene interpretation and driver intent prediction.	These models rely heavily on attention mechanisms and very large matrix multiplications, which have a different computational pattern than CNNs. Accelerators need sufficient on-chip memory and high memory bandwidth to handle the massive size of these models and must be designed to process the self-attention mechanism efficiently.
Vision-Language Models (VLMs) & Multi-modal AI	*The Key to Advanced Interaction and Reasoning:* VLMs are the next frontier, combining visual data with semantic language understanding. This allows for truly intelligent interaction. A driver could ask, "Is that parking spot big enough for my car?" or "What does that temporary orange sign ahead say?" The VLM fuses data from the camera with the query to provide a contextual answer, enabling a new level of human-machine partnership.	Processing VLMs requires a hybrid architectural approach. The hardware must be able to efficiently run both the CNN-based "vision backbone" and the Transformer-based "language head" simultaneously. This demands highly flexible data paths and programmable compute units that can be dynamically configured for different parts of the multi-modal network.

7. Computing Cores and Instruction Set Architectures (ISAs): The Foundational Choice

The Instruction Set Architecture (ISA) is the fundamental DNA of a processor, defining the basic commands the software can use to control the hardware. The choice of ISA has profound, long-term implications for cost, performance, flexibility, and strategic independence. While the proprietary Arm architecture has become the de facto standard in the automotive industry, its dominance creates a critical dependency. This has fueled a surge of interest in RISC-V as a strategic, open-source alternative. Arm will still remain market leader for many years, but one sees that RISC is gaining more and more momentum in the market.

The ARM Dependency: A Double-Edged Sword

ARM provides a robust, proven, and highly optimized ecosystem of processor cores (e.g., Cortex-A for performance, Cortex-R for real-time safety). This has enabled the industry to build advanced SoCs with a high degree of confidence. However, this reliance comes at a strategic cost: licensing fees, royalty payments for every chip shipped, and being locked into ARM's product roadmap and business decisions. OEMs and Tier 1s have limited ability to customize the core hardware to their specific needs.

RISC-V: The Strategic Alternative for Automotive Sovereignty

RISC-V is not a processor; it is an open-standard ISA. This means anyone can use it to design their own processor cores without paying licensing or royalty fees. This fundamental difference provides a powerful strategic rationale for its adoption in the automotive sector.

Advantage of RISC-V	Description
Strategic Independence & Freedom from Lock-In	By adopting an open standard, OEMs and suppliers can break free from dependency on a single vendor. They are no longer subject to the pricing, roadmap, or geopolitical risks associated with a single company. This fosters a more resilient and competitive supply chain, allowing companies to source cores from multiple vendors or design them in-house.
Deep Customization and Domain-Specific Acceleration	The RISC-V ISA is modular and extensible. This is its most powerful feature for automotive. Companies can add custom instructions to the core processor to accelerate specific, critical workloads—such as sensor fusion algorithms, cryptography, or specific AI operations. This allows for hardware that is precisely tuned to the vehicle's software, leading to significant gains in performance and power efficiency that are impossible with a fixed, proprietary ISA.
Cost Reduction (No Licensing or Royalties)	The open-source nature of the RISC-V ISA eliminates the substantial licensing fees and per-chip royalties associated with proprietary ISAs like ARM. Over the millions of vehicles produced, this translates into enormous cost savings, allowing for greater investment in other areas of innovation.
Unprecedented Scalability and Consistency	The RISC-V architecture is designed to be scalable from the smallest embedded microcontroller (for a window lift or sensor node) to a high-performance 64-bit multi-core processor for the central compute cluster. This allows an OEM to potentially build its entire E/E architecture—from the edge to the center—on a single, consistent ISA. This drastically

	simplifies software development, toolchains, and developer training across the entire vehicle platform.
Fostering an Open and Collaborative Ecosystem	RISC-V is managed by a global foundation with members from across the industry. This collaborative model accelerates innovation and ensures the long-term stability and evolution of the standard. It encourages the development of a broad ecosystem of tools, software, and talent, reducing the risk for any single company that adopts it.

In the table below, one can see a brief comparison[1] of Intel x86, Arm & RISC-V on a number of criteria:

ISA	Chips?	Architecture License?	Commercial Core IP?	Add Own Instructions?	Open-Source Core IP?	Automotive Maturity incl. tooling and ecosystem
x86 Established	Yes, few vendors	No	No	No	No	High
ARM Dominant ecosystem	Yes, many vendors	Yes, expensive	Yes, one vendor	No (Mostly)	No	High
RISC-V Open standard	Yes, many vendors	Yes, free	Yes, many vendors	Yes	Yes, some available	Partial

KEY TAKEAWAY **RISC-V offers the highest customization and openness, while x86 and ARM currently lead in automotive maturity.**

In summary, while Arm provides a ready-made solution, RISC-V offers a strategic toolkit. For automotive OEMs aiming to take control of their E/E architecture and differentiate themselves through software, RISC-V provides a pathway to create highly optimized, cost-effective, and sovereign hardware solutions that are tailored precisely to their vision.

Although the usefulness of RISC-V has been proven in a number of Chips-JU projects (a.o. TRISTAN & ISOLDE) by means of building the foundation for future industrialization of RISC-V based processor solutions, overall the RISC-V ecosystem is still in its early stages, requiring more foundational investment from adopters to build out the necessary tools and software libraries. In contrast, the Arm ecosystem is exceptionally mature, offering robust software support and a vast pool of engineering talent with decades of experience.

To bridge this gap, a primary strategic need is to accelerate the European RISC-V APU ecosystem to reach Arm-class maturity. This requires driving development across three critical domains: establishing production-ready software stacks, delivering safety-certified developer tooling (such as compilers, debuggers, and simulators), and securing commercialization readiness. Hardening these open-standard processor cores to meet stringent automotive-grade qualifications is the key to transforming RISC-V from a highly customizable technical alternative into a fully industrialized, sovereign hardware platform ready for high-volume series production.

8. Heterogeneous Integration via modular building blocks

At the silicon architecture level, the transition to centralized vehicle computing coincides with a profound paradigm shift in how high-performance compute scales. For decades, the semiconductor industry relied strictly on geometric transistor scaling—the traditional Moore's Law—to drive performance gains. Today, however, monolithic System-on-Chips (SoCs) are colliding with physical reticle limits, severe thermal dissipation barriers, and exponentially rising design and mask costs, making a structural shift inevitable. Future performance scaling will no longer be achieved simply by squeezing more transistors onto a single, monolithic die. Instead, it must be driven by minimizing signal transmission latency, optimizing system-level orchestration, and maximizing interconnect efficiency. This system-driven scaling paradigm makes **heterogeneous integration via chiplets** the definitive architectural path for future automotive central brains, steering the industry away from monolithic integration and toward modular, specialized excellence.

To support the processing demands of software-defined vehicles, monolithic designs must be replaced by heterogeneous systems. Rather than forcing all diverse functional blocks—from digital compute and memory to analog sensor interfaces—onto a single, massive, and cost-prohibitive silicon substrate, modular chiplet architectures decouple these functions. This approach enables the integration of diverse technologies,

[1] Extended from "What's next for RISC V?, TechSpot"

manufactured on their respective optimal process nodes, into a single, highly integrated package. The ultimate objective of this transition is to establish a dynamic, open marketplace of chiplets, where European OEMs and Tier-1 suppliers can flexibly select and combine standardized components of varying technological origins, functional scopes, and performance classes. This standardized framework is key to accelerating development timelines, improving silicon yields, and unlocking the highly scalable customization allowing to keep European software-defined vehicles globally competitive.

However, a mature and open marketplace of chiplets will not emerge overnight; establishing it requires a sustained, multi-stage industry effort. In the immediate term, the first step must focus on co-developing a highly focused, initial set of chiplets designed tightly alongside one another using open standards to validate physical and logical compatibility. Once this foundational set is proven, the next evolutionary step will extend this baseline by rolling out the validated interface standards to a wider variety of specialized dies. A fully realized, open marketplace of chiplets represents the logical culmination of this process, enabling silicon developers to monetize their designs not only within the automotive sector but also across adjacent high-performance markets, thereby securing the massive economies of scale necessary to sustain long-term European semiconductor competitiveness.

9. Further Demands from SDV SW Perspective

To support future software-defined vehicle architectures, we must coordinate technical requirements directly with the **ECAVA SDV Working Group**. High-performance workloads require massive data throughput and complex safety-critical isolation, meaning traditional microcontroller approaches are no longer sufficient, and high data bandwidth requirements are heavily driving the demand for advanced hardware platforms.

Debugging and Trace Tools

The development of modular, chiplet-based hardware requires comprehensive debugging and tracing using standard physical probes and software across all components:

Tooling/Trace Category	Specific Technical Requirement
Debugging	Transparent debugging of single dies
Programming	Transparent flashing and download mechanisms
Performance Analysis	Direct software profiling

Technical Communication Interfaces

External communication between the chiplet package and peripheral board components must be handled via low-latency, high-bandwidth interfaces supporting Direct Memory Access (DMA):

- External DMA Interface: PCIe or equivalent standard interfaces for package-external high-speed communication.
- Initial Debug Access: JTAG interface utilized for initial board bring-up and low-level diagnostic access.

Overarching System-Software Validation Concepts

Validation must be executed holistically at the system board level rather than in isolation, necessitating rigorous validation frameworks for the following critical system-software concepts:

System-Software Concept	Key Verification & Validation Focus Areas
Boot & Initialization	Verification of startup and boot phases
Cybersecurity	Implementation of secure boot and secure communication protocols
Safety & Redundancy	Safety-critical fail-operational behavior, fault tolerance, and hardware redundancy
Power Management	Integrated power management strategies and clock rate reduction
System Health	Continuous self-testing and diagnostic monitors
Physical Interfaces	Support for standardized UCIe and I3C communication lines
Operation Modes	Multi-mode scheduling and runtime configuration
Thermal Strategy	Holistic active thermal management strategies

1.3. Evolution of Vehicle Compute Platforms

The trajectory of vehicle compute is marked by a relentless pursuit of higher performance, greater integration, and increased architectural flexibility. As the industry moves from today's monolithic SoCs towards more advanced and modular solutions, the vehicle's electronic architecture will transform at an accelerated pace. This evolution can be viewed in three distinct horizons, each building upon the last to enable the fully software-defined, autonomous, and connected vehicle of the future.

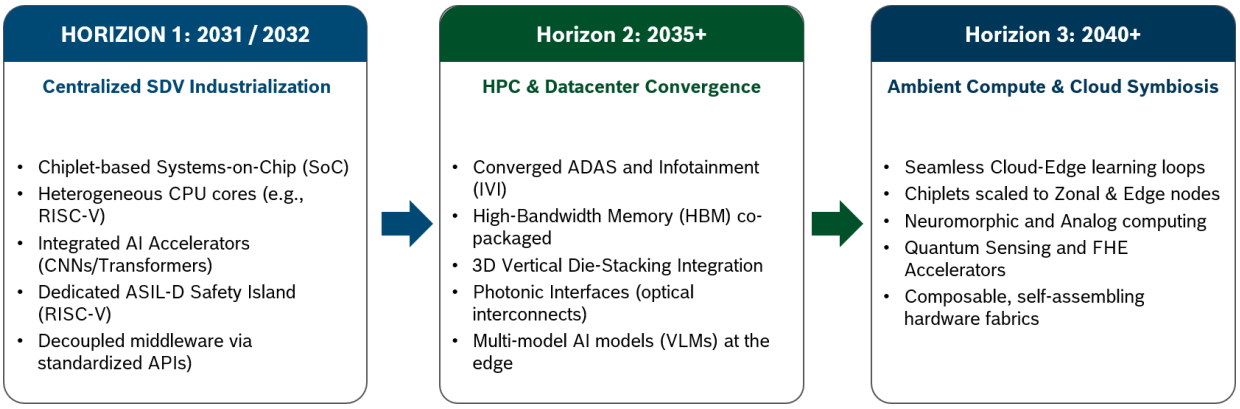


Figure 2: Evolution of Automotive SoC Capabilities and Compute Horizons

1.3.1. Target SoC Capability 2031/32

This era represents the industrialization of the first-generation centralized, software-defined vehicle platforms. The primary goal is to overcome the limitations of current monolithic SoCs by embracing chiplet-based systems and converging previously separate domains into a cohesive central compute unit. This central compute becomes the vehicle's digital brain system.

- Next-Generation Automotive SoC: The architecture will be a chiplet-based System-on-Chip, moving away from inflexible monolithic designs. This allows for a "mix-and-match" approach, combining specialized dies for CPU, AI, and I/O in a single package. This modularity reduces development costs and accelerates time-to-market for different vehicle tiers.
- High-Performance Heterogeneous Computing: The first generation should combine production-ready high-performance CPU IP and mature software toolchains with European-designed AI accelerators, safety and security components, and system integration. Arm-based application processors are expected to provide the primary near-term baseline, while automotive-grade RISC-V processors should be introduced first in functions where sufficient maturity is demonstrated, such as safety islands, real-time control, and control-plane processing.
- Integrated AI Acceleration: Hardware accelerators will be essential for processing data from cameras, radar, and LiDAR. These accelerators will be optimized for established Convolutional Neural Networks (CNNs) for perception tasks and will begin to incorporate flexibility for emerging Transformer-based models used in natural language processing for advanced voice assistants.
- Integrated Safety & Security Computing: A dedicated safety island, likely built around a real-time processor core (e.g., RISC-V with custom extensions, including defined profiles of the instruction set), will be integrated as a chiplet. This component will manage system monitoring, fault correction, and cybersecurity functions, ensuring the vehicle's core safety is sovereign and verifiable.
- Software-Defined Vehicle (SDV) Platform: Standardized APIs and a robust middleware layer will decouple software from the underlying hardware. This enables Over-the-Air (OTA) updates for bug fixes, security patches, and new Features-on-Demand (FoD), creating new revenue streams for OEMs long after the vehicle is sold.
- Reduced Vehicle Complexity: Consolidating functions into a central compute unit drastically simplifies the vehicle's E/E architecture. This leads to a reduction in the number of ECUs, a significant decrease in wiring harness weight and complexity, and improved overall system reliability.

1.3.2. Vision 2035+: Future Vehicle Compute

Marking the maturity of chiplet-based architectures, this phase will see the integration of technologies pioneered in high-performance computing (HPC) and data centers. The focus shifts from merely consolidating

functions to building distributed, high-bandwidth systems capable of handling the immense data loads required for Level 4/5 autonomy and deeply immersive in-cabin experiences.

- Converged ADAS and IVI Functionality: The system will consolidate Advanced Driver-Assistance Systems (ADAS) and In-Vehicle Infotainment (IVI) onto a single hardware platform. This is enabled by hypervisors and hardware-based "Freedom from Interference" (FFI) mechanisms, ensuring that a fault in a non-critical IVI application cannot impact safety-critical ADAS functions.
- Datacenter Technologies Entering the Vehicle: The vehicle's compute architecture will increasingly resemble a mobile data center. This includes leveraging advanced packaging and integration techniques previously confined to servers.
- ISA-Agnostic Chiplet Ecosystem: Standardized hardware and software interfaces allow Arm, RISC-V, and specialized compute chiplets to coexist or be exchanged without changing the overall platform architecture. This creates the technical foundation for progressively increasing the share of European-controlled processor IP.
- High-Bandwidth Memory (HBM): To overcome the "memory wall," processors will be co-packaged with stacks of DRAM (HBM) on a silicon interposer. This provides the massive memory bandwidth essential for processing large datasets for sensor fusion and running complex, multi-modal AI models.
- 3D Integration: Chiplets will be stacked vertically (3D stacking), not just placed side-by-side. This allows for even shorter, faster connections between logic and memory, dramatically improving performance and power efficiency while minimizing the physical footprint of the compute module.
- Photonic Interfaces: For ultra-high-bandwidth communication between different compute nodes within the vehicle (e.g., connecting a central compute unit to a zonal controller), optical interconnects (photonics) may replace traditional copper wiring. This allows for multi-terabit data rates over longer distances with near-zero electromagnetic interference.
- Distributed Compute Architectures: The architecture will evolve into a hybrid model, combining a powerful central computer with intelligent zonal controllers. These zonal ECUs will handle local real-time processing and data aggregation, communicating with the central unit via a high-speed automotive Ethernet backbone.
- Large-Scale In-Vehicle AI Capability: The hardware will be designed to efficiently run sophisticated Vision-Language Models (VLMs) and other multi-modal AI. This enables true contextual awareness, allowing a driver to ask, "Is there a coffee shop on the next block?" and have the car fuse map data with real-time camera feeds to provide an intelligent answer.
- Outcome → Highly scalable, flexible, and redundant vehicle platforms: The mature chiplet ecosystem, combined with distributed architectures, allows OEMs to create a truly scalable platform. The same foundational I/O and safety chiplets can be used across all vehicle models, while the number and type of compute chiplets can be scaled to meet performance targets, creating a cost-effective and resilient hardware family.

1.3.3. Vision 2040+: Future++ Vehicle Compute

Looking towards the N+3 generation, vehicle compute platforms will transcend the role of a networked computer and become a fully integrated, ambient component of the mobility ecosystem. The system will be defined by its ability to learn, adapt, and interact with its environment in real-time, underpinned by advances in quantum computing, neuromorphic hardware, and seamless cloud integration. Subject to competitive performance, qualified tools, automotive-grade software, safety evidence, and dependable commercial support, European RISC-V technologies may increasingly replace proprietary processor IP, including in high-performance application-processing roles.

- Seamless Cloud-Edge Symbiosis: The vehicle will operate as an intelligent edge node in a larger computational network. It will offload non-time-critical, computationally intensive tasks (e.g., training a new driving policy for a specific geographic area) to the cloud and receive personalized, updated models in real-time, creating a continuously learning system.
- Chiplet Scaling to Zonal and Edge Systems: Future vehicle platforms will extend chiplet-based architectures beyond central HPC into zonal controllers and distributed edge ECUs through lightweight, cost-optimized die-to-die interconnects (RDI). This enables modular, reusable hardware building blocks that scale efficiently into high-volume vehicle segments, supporting localized AI

processing, real-time control, and sensor aggregation close to the edge. By combining centralized HPC orchestration with distributed chiplet nodes, the architecture improves system resilience, reduces single points of failure, and allows flexible upgrades and supplier diversity across the entire E/E system

- Neuromorphic and Analog Computing: For ultimate power efficiency in perception tasks, specialized hardware will move beyond traditional digital processing. Neuromorphic accelerators, which mimic the structure of the human brain, will process sensor data using spike-based, event-driven computation, offering orders-of-magnitude improvements in performance-per-watt for tasks like object recognition and tracking.
- Quantum Sensing and Computing Integration: Early applications of quantum technology will emerge. Quantum sensors could provide unprecedented accuracy for navigation in GPS-denied environments. In the cloud, quantum computers might be used to solve complex optimization problems for fleet management, traffic flow, or material science for next-generation batteries, with the results deployed to the vehicle.
- Fully Homomorphic Encryption (FHE) Accelerators: To ensure absolute data privacy, future platforms may incorporate hardware accelerators for FHE. This would allow the vehicle or the cloud to perform computations on fully encrypted data without ever decrypting it, protecting sensitive personal and operational information from all cyber threats.
- Composable, Self-Assembling Hardware: The system architecture will become fully dynamic. Using advanced interconnects and a software-defined fabric, the vehicle could reconfigure its own hardware resources in real-time. For instance, it could allocate more computational power to the infotainment domain when parked and then dynamically re-task those same resources for sensor fusion the moment the vehicle begins to drive.

1.4. Components of a European SDV Hardware Platform

Today's automotive compute solutions expose two opposing limitations: monolithic SoCs deliver high performance and efficiency but are costly, inflexible, difficult to scale, and create supplier dependencies, while board-level integration offers modularity but suffers from latency, power, footprint, and system-complexity penalties. A European SDV hardware platform must therefore combine the best of both worlds retaining SoC-like performance while enabling modularity, reuse, scalability, and multi-vendor sourcing. Building on this approach, the platform brings together the core SDV components: application processing, AI acceleration, safety and security islands, high-bandwidth memory, standardized interconnects, sensor and vehicle interfaces, and a software-compatible base die.

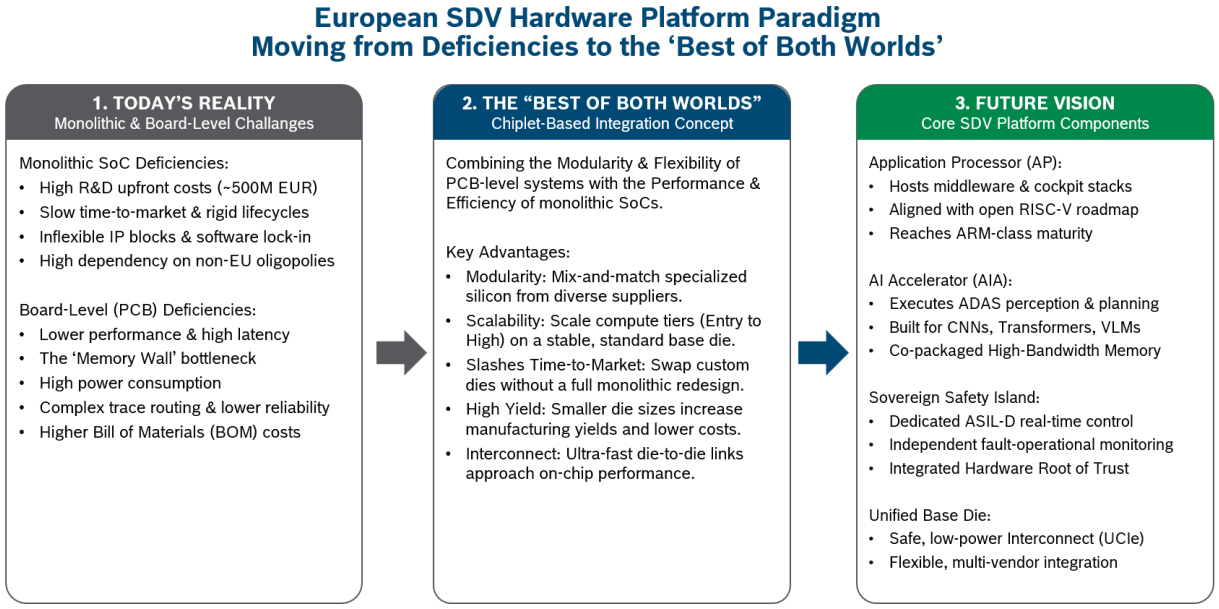


Figure 3: The European SDV Hardware Platform Paradigm

1.4.1. Key Deficiencies of Monolithic Automotive SoCs

As system complexity continues to increase, traditional monolithic architectures face growing challenges in scalability, design reuse, development cost, and time-to-market. Integrating all functionalities onto a single die can lead to lower manufacturing yields, reduced flexibility for technology upgrades, and difficulties in optimizing individual components for specific application requirements. These limitations are becoming increasingly critical as next-generation applications demand higher performance, greater customization, and faster innovation cycles. The primary deficiencies of today's monolithic SoC designs are outlined below.

Deficiency	Description
High Research & Development (R&D) Costs	Monolithic designs require enormous upfront investment in R&D. Designing, verifying, and testing a complex, large-scale SoC is a multi-hundred-million-dollar endeavor. The risk is immense, as a single flaw in the design can necessitate a full "respin" (re-manufacturing of the silicon masks), costing tens of millions and adding significant delays.
Slow Time-to-Market	The sheer complexity of monolithic SoCs leads to protracted development cycles, often spanning several years. This slow pace is fundamentally at odds with the rapid innovation in software and AI. By the time a vehicle with a new SoC reaches production, its core processing technology may already be a generation behind the consumer electronics industry, creating a persistent technology gap.
Lack of Modularity	Monolithic SoCs are an "all-or-nothing" proposition. OEMs are locked into the specific combination of IP blocks (CPUs, AI accelerators, GPUs) chosen by the chip vendor at the start of the design. It is impossible to upgrade or replace a single functional component—such as a newer, more efficient AI accelerator—without redesigning the entire chip. This inflexibility stifles innovation and prevents OEMs from integrating best-in-class technology as it becomes available.
Poor Scalability Across Vehicle Platforms	A single monolithic SoC design cannot be efficiently scaled across an OEM's entire vehicle portfolio, from budget-friendly models to luxury flagships. Creating multiple versions of a monolithic chip for different segments is cost-prohibitive. This forces OEMs into a difficult compromise: either using an overpowered (and expensive) chip in lower-end models or being unable to deliver premium features in high-end ones, undermining platform strategies.
Dependency on Single-Source Suppliers	The market for high-performance automotive SoCs is dominated by a very small number of non-European vendors. This creates significant supply chain risks and grants immense pricing power to the chipmakers. OEMs have limited leverage and are often forced to accept the technology roadmap and business terms dictated by their suppliers, ceding strategic control over their most critical component.
Inflexible Hardware for Evolving Software	The hardware is "frozen" years before the vehicle hits the road, yet the software, particularly AI models and algorithms, continues to evolve at a breakneck pace. A monolithic chip's hardware accelerators might be optimized for a specific neural network architecture, but they may perform poorly with newer, more advanced models that emerge later, rendering the hardware obsolete long before the vehicle's end of life.

1.4.2. Key Deficiencies of Board level Integration

While System-on-Chip (SoC) integration presents its own challenges, the traditional alternative—integrating multiple discrete chips on a Printed Circuit Board (PCB)—is fundamentally unsuitable for creating the powerful, efficient, and compact central compute units required by modern vehicles. This board-level approach, where individual packaged chips for processing, memory, and acceleration are connected via traces on a PCB, introduces critical bottlenecks that limit performance, increase cost, and add complexity.

The move towards SoCs is a direct response to the inherent limitations of PCB-based designs.

Deficiency	Description
Lower Performance & Higher Power Consumption	The physical distance between chips on a PCB is a fundamental performance killer. Signals must travel millimeters or centimeters between components, a journey that takes orders of magnitude longer than traveling micrometers within a single piece of silicon. This "off-chip" communication introduces significant latency and requires far more electrical power to drive the signals, directly reducing overall system performance and increasing energy consumption—a critical drawback, especially for Electric Vehicles (EVs).

The "Memory Wall" Bottleneck	High-performance computing, especially for AI, requires massive memory bandwidth between the processor and DRAM. On a PCB, this connection is limited by the number of physical pins on the chip packages and the signal integrity constraints of the board. An SoC can integrate memory controllers and even the memory itself (in advanced packaging) much more closely, enabling ultra-wide, low-latency interfaces that are impossible to achieve at the board level.
Limited Functional Integration	A PCB-based design can never achieve the tight coupling possible within an SoC. It is incredibly difficult and inefficient to create a unified memory architecture where CPUs, GPUs, and AI accelerators can seamlessly share data. Instead, data must be explicitly copied between the memory spaces of different chips, adding latency and software complexity, and preventing the use of more advanced, heterogeneous computing paradigms.
Increased System Complexity and Lower Reliability	A board with numerous high-performance chips is a far more complex system to design, manufacture, and validate. More components mean more solder joints, more potential points of failure, and more intricate thermal management challenges. The increased complexity directly translates to a higher probability of manufacturing defects and a lower overall Mean Time Between Failures (MTBF) for the finished ECU.
Higher Bill of Materials (BOM) and Manufacturing Cost	While a single monolithic SoC is expensive to design, at volume, it is often cheaper to manufacture than procuring multiple separate, high-performance chips and assembling them onto a complex, multi-layered PCB. The costs for the individual packaged chips, the larger and more sophisticated PCB required, and the more complex assembly process add up, often making the multi-chip solution more expensive for the final product.

1.4.3. Chiplet Systems – A 'Best of Both Worlds' for Automotive Compute

As the industry grapples with the opposing challenges of monolithic SoCs (inflexible, high R&D costs) and board-level integration (poor performance, large footprint), a third path has emerged: *chiplet-based systems*. This approach, also known as package-level integration, represents a paradigm shift in silicon design.

> **CHIPLETS ARE "MINI-CHIPS", THAT ACT AS MODULAR SILICON BUILDING BLOCKS FOR SEMICONDUCTORS**
>
> Instead of building one giant, monolithic chip, a chiplet system assembles multiple smaller, specialized dies ("chiplets") within a single advanced package.
>
> These chiplets are connected by an ultra-high-speed, low-latency "die-to-die" interconnect, allowing them to communicate as if they were on the same piece of silicon. This "LEGO block" approach seeks to combine the modularity and flexibility of a PCB-based system with the performance and efficiency of an SoC, creating a powerful and adaptable solution for future automotive central computers.

The Benefits of Chiplet-Based Systems:

Benefit	Description	Automotive Example
Modularity & Flexibility	This is the core advantage. OEMs and Tier 1s can "mix-and-match" chiplets from different vendors or process nodes. A single system can combine a high-performance CPU chiplet made on a 3nm process with a more mature I/O chiplet made on a 22nm process, optimizing cost and performance.	An OEM can design a base compute platform and, for a high-end variant, simply swap in a more powerful AI accelerator chiplet from an innovative startup without redesigning the entire system.
Drastically Reduced Time-to-Market for Variants	Instead of designing a new monolithic SoC for each vehicle tier, OEMs can create a family of products by simply varying the combination of chiplets. Developing a single new chiplet or a package variant is far faster and cheaper than a full SoC respin.	A "base" model might use one CPU chiplet and one AI chiplet. A "mid-range" model could add a second AI chiplet, and a "premium" model could add a more powerful GPU chiplet. A later package variant can include a new AI chiplet, which was initially not yet available. This entire portfolio can be developed in a fraction of the time required for three separate monolithic designs.

Scalability Across Vehicle Platforms	Chiplet systems are inherently scalable. The same foundational chiplets (like I/O and safety) can be used across an entire vehicle line, while the number and type of compute chiplets can be scaled up or down to meet the specific performance and cost targets of each model, from entry-level to luxury.	A small, affordable EV might use a package with a single CPU/AI compute chiplet. The flagship luxury sedan could use the same I/O and safety chiplets but connect to four compute chiplets in its package, providing immense performance while sharing the same core software architecture.
Lower Integration Hurdles for Innovation	Chiplets democratize hardware innovation. A startup with a groundbreaking AI accelerator doesn't need to build a complete SoC. They can focus on creating a world-class chiplet. OEMs can then integrate this cutting-edge solution quickly, without being locked into the slower innovation cycle of large, established semiconductor vendors.	When a new, more efficient type of neural network processor (e.g., for VLMs) becomes available, an OEM can qualify it as a chiplet and integrate it into their next product cycle in 12-18 months, rather than waiting 3-4 years for it to appear in a next-generation monolithic SoC.
Higher Performance & Efficiency than PCBs	The die-to-die interconnects within a chiplet package offer massive bandwidth at extremely low power—approaching the performance of on-chip wiring. This virtually eliminates the "off-chip" communication bottleneck of PCBs, allowing for performance and power efficiency that is far superior to board-level designs.	AI models requiring massive data throughput between compute units and memory can run far more efficiently, as the physical distance data has to travel is reduced from centimeters (on a PCB) to micrometers (within the package).
Improved Yield and Lower Manufacturing Cost	Manufacturing a single, large monolithic die is risky. A tiny defect anywhere on the die can render the entire, expensive chip useless. Manufacturing multiple smaller chiplets is far more reliable; the yield is higher because the probability of a defect on a smaller area is lower. This "disaggregation" reduces overall manufacturing cost at scale.	If the manufacturing yield for a large 500mm^2 monolithic SoC is 70%, but the yield for a 100mm^2 chiplet is 95%, it is far more cost-effective to produce five smaller chiplets and discard the few that fail than to discard 30% of the large, expensive monolithic dies.
Increased competition	Lowering entry barrier leads to more solution providers. The monolithic SoC market is an oligopoly, dominated by a few large players with the capital to fund giant chip designs. The lower R&D cost of designing a specialized chiplet opens the market to a wider range of innovative companies, fostering competition in price, performance, and features.	A European OEM is no longer limited to few major SoC vendors. They can now evaluate an AI accelerator chiplet from a European startup, a security chiplet from a trusted source, and a CPU chiplet from a US giant, forcing all suppliers to be more competitive.
Increase digital sovereignty	Chiplets give regions like Europe a viable path to regaining technological control. Instead of being entirely dependent on non-European monolithic SoCs, a European OEM can define the system architecture and strategically source key chiplets (like safety and security) from local suppliers. This reduces geopolitical supply chain risk and keeps critical IP within the region.	A European car manufacturer could mandate that the safety-critical MCU chiplet and the cryptographic engine chiplet are designed and fabricated within the EU. This ensures compliance with European safety standards (ASIL) and cybersecurity regulations, and guarantees that no foreign state has backdoor access to the vehicle's core functions.

The Deficiencies and New Challenges of Chiplets:

Despite their significant advantages, chiplet systems are not a silver bullet. They introduce new layers of complexity that must be managed.

Deficiency	Description
Advanced Packaging Complexity and Cost	Integrating multiple chiplets into a single package requires highly advanced and expensive packaging technologies (e.g., 2.5D interposers, 3D stacking). The design and manufacturing of these packages are complex, and the thermal challenges of cooling a dense collection of hot chiplets can be even more severe than for a monolithic SoC. Further research for cost-optimal solutions is imminent.
Standardization of Die-to-Die Interconnects	For a true mix-and-match ecosystem to thrive, the interfaces between chiplets must be standardized. While standards like the Universal Chiplet Interconnect Express (UCIe)

	are emerging, the industry is still in the early stages of adoption. Without a universally accepted standard, OEMs could become locked into a specific vendor's proprietary chiplet ecosystem.
Known Good Die (KGD) and System-Level Testing	In a monolithic design, the final chip is tested once. In a chiplet system, each individual chiplet must be tested and confirmed to be a "Known Good Die" *before* it is integrated into the final, expensive package. Once assembled, the entire system-in-package must be tested again. This multi-stage testing process is complex and adds to the overall cost and logistical overhead, as well as commercial risks in case the full package is tested non-functional.
Holistic System Design and Integration Responsibility	The chiplet approach shifts responsibility. Instead of buying a complete, validated SoC from a single vendor, a new role of a chiplet system integrator consolidating the responsibility will emerge. They are now responsible for the architectural definition, thermal management, power delivery, and final validation of a complex system composed of parts from multiple different suppliers, which requires significant in-house expertise.
Increased Risk of initial solution complexity	While chiplets promise future flexibility, the initial implementation is a pioneering effort. For a single company to address a first-of-a-kind chiplet solution is incredibly challenging. Adopting a multi-vendor approach, which is key to the chiplet vision, further magnifies this risk due to a current lack of mature industry standards, limited collective experience, and often unclear roles and responsibilities (R&R) between chiplet suppliers, the packaging vendor, and the system integrator.
Ecosystem and supplier coordination	Chiplets only scale if suppliers, design houses, packaging providers, OEMs, Tier-1s, and integrators align on interfaces, responsibilities, qualification, business models, and lifecycle support
Introduction of new role models	A multi-vendor chiplet integration will need a responsibility consolidation at the system integrator level, which then has the full liability for the chiplet system against the customers.

1.5. Strategic Imperatives for Chiplet Adoption, Driven by OEM Needs

Our roadmap prioritizes the development and adoption of a chiplet-based platform to build affordable, customizable, and heterogeneous Systems-on-Chip (SoCs) for next-generation vehicles. This needs strategical alignment towards a goal of an European Automotive Chiplet (Eco)System (EU-ACS) and is fundamental to reducing development risk, lowering costs, and preventing vendor lock-in.

The core of this strategy is to create a customizable reference platform that enables stagewise investment and cost-sharing with partners. By fostering an open ecosystem, we will drive competition and gain back critical control points in the vehicle computing stack. Early OEM involvement and rapid market feedback loops are built into the process, ensuring our SoC solutions are precisely aligned with crucial automotive requirements for safety, performance, and scalability. This approach will accelerate integration, de-risk our NRE exposure, and secure a competitive advantage in the future of automotive electronics.

1.6. Overview of control points needed inside Europe

Where Do We Need Control for Chiplets? A Strategic Assessment

Adopting a chiplet-based strategy is not just a technical shift; it requires a clear-eyed assessment of where to lead and where to partner. For the European automotive industry to successfully navigate this transition and achieve digital sovereignty, it must exert control over the areas that define the vehicle's unique character and safety, while strategically addressing the areas where competencies are currently lacking. This means leveraging inherent strengths in system-level thinking and safety, and building or sourcing expertise in high-performance silicon design and advanced manufacturing.

Areas of Existing Strength: The European Advantage

Europe's deep heritage in automotive engineering provides a strong foundation for defining and controlling the system architecture. These are the domains where European OEMs can and should lead.

Strength	Description
System Knowledge & E/E Architecture	The core strength of European OEMs lies in their holistic understanding of the vehicle. They are uniquely positioned to define the top-down requirements for the entire E/E architecture, dictating how data flows, where functions reside, and how different domains interact. This system-level knowledge is the crucial starting point for defining any chiplet-based compute platform.
Functional Safety & Security (ASIL)	The European automotive industry has a world-leading expertise in functional safety (ISO 26262 / ASIL) and security. This is deeply embedded in the design culture, particularly in microcontroller (µC) development. Europe can own the definition and even the design of the "Safety Island" chiplet, the security engine chiplet, and the real-time control chiplets, ensuring the vehicle's core safety functions are non-negotiable and sovereign.
In-Vehicle Network Technology	European OEMs have historically been the driving force behind the standardization of car-wired networks (e.g., CAN, FlexRay, Automotive Ethernet). This deep expertise in network architecture and traffic management is critical for designing the I/O chiplet, which acts as the gateway between the central compute unit and the rest of the vehicle.
Standard Packaging & Integration	While lacking in the most advanced packaging techniques, the European ecosystem has robust experience in standard automotive-grade packaging for power electronics, sensors, and MCUs. This knowledge is essential for ensuring the reliability, thermal management, and durability of the final system-in-package.
Network-on-Chip (NoC) & Real-Time Cores	There is growing European expertise in Network-on-Chip (NoC) technology, which is vital for routing data efficiently within a chiplet system. Furthermore, the rising adoption of the open RISC-V standard has spurred the development of European IP for real-time safety cores (R-Cores), providing a sovereign alternative to proprietary solutions.
Defining Scalability Demands	European OEMs excel at creating broad vehicle platforms that span multiple brands and segments. They understand precisely how to define the scalability requirements for a compute architecture, ensuring that a chiplet-based solution can be cost-effectively deployed from an entry-level compact car to a premium SUV.

Competency Gaps and Strategic Weaknesses

To build a complete system, Europe must be realistic about its current weaknesses in the high-performance silicon space and develop a strategy to close these gaps through investment, partnerships, or strategic sourcing.

Weakness	Description
High-Performance Compute (HPC), AI & memory	The design of cutting-edge, high-performance CPU cores, GPUs, and especially AI accelerators is currently dominated by non-European companies. Europe lacks a leading-edge supplier capable of producing the "brains" for autonomous driving and advanced infotainment at scale. This also includes Memory cells.
SoC System Design & Integration	While there are pockets of expertise (e.g., NXP's S32 family), the broader European industry lacks deep, at-scale experience in the holistic system design and integration of complex, multi-billion-transistor SoCs. The expertise to manage the intricate

	interplay of power, performance, and thermals for a complete HPC system is not widespread.
Advanced-Node Chip Design	Designing chips for the most advanced semiconductor process nodes (e.g., 5nm, 3nm) is a highly specialized skill. There are very few design service and IDM companies in Europe with a proven track record of successfully taping out complex designs at the leading edge, which is where HPC chiplets are made.
Wireless Connectivity	The IP and expertise for the latest wireless technologies (5G, Wi-Fi 6/7, Bluetooth) are largely concentrated outside of Europe. This is a critical gap, as vehicle-to-everything (V2X) communication and in-cabin connectivity are essential for future automotive functions.
Advanced Packaging Technology	The most advanced techniques for integrating chiplets (2.5D/3D stacking, high-density fan-out) are part of a supply chain that resides almost exclusively in Asia. Europe lacks the at-scale manufacturing capability for these critical enabling technologies.
EDA Tooling and Ecosystem	The Electronic Design Automation (EDA) software required to design, verify, and test complex chips and chiplet systems is supplied by a small oligopoly of non-European vendors. This represents a fundamental dependency for any silicon design activity undertaken in the region.

1.6.1. Processor Architectures (e.g., RISC-V for Automotive) for European Competitiveness

COMPUTING CORES AND INSTRUCTION SET ARCHITECTURES (ISAS): THE FOUNDATIONAL CHOICE

The Instruction Set Architecture (ISA) is the fundamental DNA of a processor, defining the basic commands the software can use to control the hardware. The choice of ISA has profound, long-term implications for cost, performance, flexibility, and strategic independence. While the proprietary ARM architecture has become the de facto standard in the automotive industry, its dominance creates a critical dependency. This has fueled a surge of interest in RISC-V as a strategic, open-source alternative.

The ARM Dependency – A Double-Edged Sword: ARM provides a robust, proven, and highly optimized ecosystem of processor cores (e.g., Cortex-A for performance, Cortex-R for real-time safety). This has enabled the industry to build advanced SoCs with a high degree of confidence. However, this reliance comes at a strategic cost: licensing fees, royalty payments for every chip shipped, and being locked into ARM's product roadmap and business decisions. OEMs and Tier 1s have limited ability to customize the core hardware to their specific needs.

RISC-V – The Strategic Alternative for Automotive Sovereignty: RISC-V is not a processor; it is an open-standard ISA. This means anyone can use it to design their own processor cores without paying licensing or royalty fees. This fundamental difference provides a powerful strategic rationale for its adoption in the automotive sector.

For automotive OEMs aiming to take control of their E/E architecture and differentiate themselves through software, RISC-V provides a pathway to create highly optimized, cost-effective, and sovereign hardware solutions that are tailored precisely to their vision.

Maturity-based adoption path. RISC-V should be developed as a strategic European option, but not treated as the exclusive architecture for near-term automotive HPC. Current European gaps relate not only to processor cores, but also to automotive-grade application-processor IP, compilers, debuggers, trace tools, simulators, operating-system support, hypervisors, middleware, safety documentation, qualification evidence, and long-term commercial support. Therefore, near-term platforms should use the most mature processor technology available, while maintaining open interfaces and software portability that enable RISC-V substitution as these gaps are closed.

A close collaboration is planned between all European projects and initiatives with respect to RISC-V based processor developments and chiplet integration use cases and demonstrators. The technical requirements, ISA agnostic architectural concepts and especially interface specifications from RISC-V based processor developments need to be appropriately considered in the targeted use cases, the overall architecture and the final system integration approach based amongst others on chiplets. This technological alignment between both major hardware research initiatives allows a later seamless integration of RISC-V products in the chiplet reference architecture. An early collaboration between both initiatives should create a clear understanding of appropriate interface specifications and standardizations, tackling technological challenges like architecture partitioning, chiplet-to-chiplet interconnect, testability, CAD tools and flows, signal integrity, thermal aspects, I/Os, advanced packaging technologies. Die-to-Die (D2D) communication, etc. from an early, integrated viewpoint. Die-to-Die (D2D) communication, in particular, is the “missing link” to leverage the chiplet-based

design ecosystem, and its development in open source would enable a wide use of the approach and could become a “de-facto” standard. This needs to target a high-bandwidth, low-latency, low-energy, ultra-short-reach link between dies. As a conclusion, the interaction is fundamental for future hardware developments for OEMs. European RISC-V developments should be designed for seamless integration into ISA-agnostic chiplet reference architectures. They should progressively complement and, where automotive maturity, competitiveness, software support, and qualification readiness are demonstrated, replace Arm-based processor components in later platform generations. The interface and integration architecture should support both technologies during the transition.

1.6.2. Specialized Accelerator and HPC Chiplets (AI/ML, DSP) for ADAS and Autonomous Driving

Central HPC silicon is the most strategic compute component in the software-defined vehicle. By the early 2030s, OEMs across all segments will rely on HPC to host autonomous driving, cockpit, and connectivity workloads on a single shared platform. Today, 100% of this silicon is supplied by non-EU vendors, and only two to three suppliers globally are capable of delivering automotive-grade HPC at the required performance, safety, and reliability levels. OEMs focused on full autonomy like Rivian or Tesla are developing their own vertically-integrated HPC solutions. Other regions are actively investing to reduce equivalent dependencies (Japan through ASRA, China through proprietary silicon at OEM and Tier 1 level). Europe currently has no comparable program.

The investment required to develop competitive HPC silicon (in the order of 500M EUR per generation) cannot be justified by European automotive volumes alone. Any European solution must therefore be designed from the outset to serve multiple markets: automotive, data center, robotics, industrial edge, and potentially defense. It must also be deployable as a scalable product family rather than a single device, covering mass-market C-segment through premium and commercial vehicles. And it must arrive in time to align with OEM platform decision points.

These three conditions, multi-market viability, product family scalability, and timing alignment, define the boundary of credible European HPC initiatives.

The Three Compute Components of Automotive HPC

A European architectural approach should treat the central HPC as three distinct compute components rather than a single bundled SoC. This separation reflects how European OEMs intend to multi-source critical IP, aligns naturally with the chiplet imperative established in Section 2.5, and contrasts with incumbent vendors who integrate all functions into monolithic devices.

Component	Role	Primary Metric	Strategic Sourcing Goal
AI Accelerator (AIA)	Executes ADAS perception, planning, generative AI, and cockpit AI workloads. The differentiating compute element.	Sustained TOPS, memory bandwidth, precision flexibility	European-designed. The most strategic IP block to control.
Application Processor (AP)	Hosts the SDV middleware, runs general-purpose workloads, manages the cockpit stack, orchestrates the AI accelerator.	kDMIPS under sustained mixed-criticality load	European-designed. Cores aligned with the European RISC-V roadmap (see Section 2.6.1).
Safety Island	Provides ASIL-D monitoring, fault detection, and fail-operational behavior, independent of the application cluster.	Worst-case interrupt latency, execution determinism, bounded memory access	European-designed. An area of existing European strength (see Section 2.4).

The architectural separation matters even when the final implementation packages all three components together. It clarifies which requirements drive which silicon decisions, opens the door to substituting individual components from different suppliers across product variants, and prevents the requirements bundling that has historically locked OEMs into single-vendor roadmaps.

Workloads Driving HPC Requirements

The central HPC must concurrently host workloads with very different characteristics on a single shared platform.

ADAS and autonomous driving represent the dominant safety-critical compute consumer. The workload decomposes into perception (multi-sensor fusion dominated by transformer backbones, current model sizes 100M to 1B parameters with sub-50 ms latency budgets), prediction (agent forecasting over 3 to 8 second

horizons), planning (transitioning from classical optimization to learned and end-to-end approaches over the roadmap horizon), control (hard real-time at 50 to 100 Hz under ASIL-D determinism), and continuous driver and occupant monitoring. The shift from modular AD pipelines to end-to-end and Vision-Language-Action (VLA) models is the single largest change to the HPC compute profile expected in the 2028 to 2032 window.

The cockpit and infotainment domain consolidates digital cluster, infotainment, multi-display rendering (premium configurations now exceeding 10K horizontal resolution), and increasingly LLM-based voice and agent interaction onto the same platform as ADAS. This consolidation is enabled by safety-certified hypervisors and hardware-enforced freedom from interference, allowing mixed-criticality workloads to coexist without compromising functional safety.

Vehicle control and connectivity workloads (chassis integration, V2X, OTA, fleet telemetry) round out the workload mix. Individually less compute-intensive, but contributing real-time determinism, functional safety, and reliability requirements.

Generative AI runs across all three domains as a cross-cutting capability rather than as a separate workload category. End-to-end and VLA driving models in ADAS, on-device LLM assistants in the cockpit, and natural-language diagnostics and agentic services in vehicle communication share the same accelerator, memory, and scheduling infrastructure. This has a direct architectural consequence: memory capacity and bandwidth become as critical as peak TOPS.

Compute and Power Targets

The European HPC family must address three product tiers across the roadmap horizon. The ranges below are indicative and to be validated against OEM platform inputs.

Tier	Vehicle SOP	AIA (sustained INT8 TOPS)	AIA precision	AIA memory bandwidth	AP (kDMIPS)	Module power envelope
Entry	~2031	300 to 1,000 or higher	INT8 / FP8 / FP16	200-600 GB/s	100 to 500	30 to 80 W
Mid	~2034	1,000 to 2,500 or higher	INT8 / FP8 / FP4 / FP16	500-800 GB/s	600 to 1,500	80 to 180 W
High	~2037	2,500 to 5,000+ or even higher	FP4 / INT4 through FP16	1 TB+/s	1,500 to 3,000	150 to 300+ W

Three design constraints cut across all tiers.

First, peak benchmark figures (TOPS, DMIPS) systematically overstate the performance available to automotive workloads. The mass-market segment in particular is thermally constrained and operates near sustained-performance limits. Sustained throughput under realistic mixed-criticality conditions, not peak burst, is the relevant evaluation metric for OEM platform selection. The roadmap should require both peak and sustained figures from any candidate platform.

Second, memory architecture rather than arithmetic throughput is the binding constraint at the upper end of the range. The Mid and High tier targets cannot be delivered against extrapolated LPDDR. They depend on advanced packaged memory, 3D-stacked memory, and near-memory compute, addressed in a dedicated section of this roadmap.

Third, the underlying technology base for the upper tiers is being developed in parallel for European data center and AI infrastructure applications, where higher per-unit silicon margins and shorter qualification cycles allow faster innovation. Coordinated transfer of these innovations from the data center to the automotive edge is the most credible path to upper-tier HPC competitiveness on the roadmap horizon. This is the multi-market viability principle stated above made concrete: technology investments amortize across markets, with data center timing typically leading automotive timing by two to three years.

Safety Architecture

A European HPC platform must achieve ISO 26262 ASIL-D at the system level. The expected architectural decomposition is: AI Accelerator at ASIL-B with diagnostic coverage, dedicated safety island providing independent ASIL-D monitoring and fail-operational behavior, and AP cluster capable of hardware-enforced mixed-criticality isolation between ASIL-D driving functions and QM-rated infotainment. The overall SoC qualification follows the Safety Element out of Context (SEooC) approach. Cybersecurity requirements include ISO 21434 compliance, post-quantum cryptography, hardware root of trust, and secure boot. Mechanical and

reliability qualification follows more than AEC-Q100 and AEC-Q200, as parts of the elements of the central compute platform need to operate 24/7 over the time span of the vehicle.

Beyond the technical requirements, the safety certification path itself is a strategic opportunity. Today's automotive HPC silicon relies on non-EU IP providers for the safety documentation that underpins certification. A European HPC family certified end-to-end through European processes removes a dependency that is currently invisible to OEM customers but operationally significant, and constitutes a tangible sovereignty deliverable distinct from compute leadership.

Competitive Position and Gap Analysis

Two incumbent platforms define the competitive frontier for the SOP 2031 generation: NVIDIA DRIVE AGX Thor (Blackwell architecture, approximately 1,000 INT8 TOPS, 14 Arm Neoverse V3AE cores, 128 GB LPDDR5X at 273 GB/s, ASIL-D, mature DriveOS and CUDA software stack) and Qualcomm Snapdragon Ride Elite (Oryon CPU with Hexagon NPU, transformer accelerators, ASIL-D, industry-leading power efficiency, Snapdragon Ride AD stack).

Current solutions meet OEM needs on high AI compute for L2+/L3, integrated cockpit and ADAS, ASIL-D capable silicon, mature AD development stacks, and emerging on-device generative AI support. The persistent gaps, where a European platform must differentiate, are:

- Single-region supply chain risk: both incumbent vendors design and source primarily through non-EU supply chains.
- High power consumption for mass-market segments: current solutions are optimized for premium tiers and scale poorly into mass-market thermal envelopes.
- Vendor lock-in on software stack and tools: CUDA dependency and proprietary AD stacks impose switching costs that grow with platform investment.
- Limited OEM control over silicon roadmap: feature inclusion, timing, and pricing are dictated by the silicon vendor.
- Cost scaling challenges into volume vehicle segments: per-vehicle silicon cost remains incompatible with C-segment economics.

These gaps define the European opportunity. A European HPC family that matches incumbents on raw compute while offering competitive power efficiency in the mass-market envelope, an open software ecosystem, and OEM influence over the silicon roadmap addresses concrete pain points that performance leadership alone cannot resolve.

Implications for the European Value Chain

Delivering this HPC family requires coordinated capability across European design houses, IP providers, foundry and packaging partners, and OEM and Tier 1 anchor customers. The detailed value-chain actions are covered in Section 3 (Ecosystem Collaboration) and in the phase-specific commitments in Section 4.3 (HPC Roadmap). Three component-level dependencies cross-link this section to the rest of the roadmap:

- High-performance CPU IP for the Application Processor, addressed under the European RISC-V roadmap (Section 2.6.1). The AP gap identified above is the highest-priority of the three dependencies.
- Advanced packaging and chiplet integration, addressed under chiplet integration (Section 2.6.3).
- Next-generation semiconductor technologies (silicon photonics, 3D-stacked and near-memory architectures, novel packaging) that the upper HPC tiers depend on for delivery, transferred from European data center and AI infrastructure programs as discussed above, and addressed in a dedicated section of this roadmap.

1.6.3. Chiplet Integration: A Strategic Pillar for European Automotive Sovereignty

For Europe's automotive industry to thrive amidst a global technology race, it must embrace a strategy that plays to its strengths while mitigating its weaknesses. Lacking the massive, single high-technology end-users that drive the consumer electronics market, Europe must champion solutions that offer scalability, flexibility, and the ability to integrate cutting-edge innovations from a diverse supplier base.

Chiplet integration emerges as the key technological and strategic pillar to achieve this, providing a powerful

answer to the unique and variable demands of the automotive sector. It offers a path to combine domain-specialized functionality, a traditional European strength, with the high-performance compute capabilities driven by innovations in cloud and mobile technology.

Bridging Two Worlds: The Best of Board-Level and SoC Integration

The strategic power of chiplet integration lies in its ability to merge two traditionally separate design philosophies. It combines the core advantages of board-level integration—modularity, scalability, and supplier flexibility—with the primary benefits of System-on-Chip (SoC) integration, namely high performance, power efficiency, and tight functional coupling.

This fusion allows for the creation of automotive-grade compute architectures where essential functions like safety and security are tightly coupled with high-performance processing units, a feat that is inefficient and costly at the board level. Simultaneously, it breaks the rigid, monolithic nature of traditional SoCs, which are often ill-suited and lacking diverse requirements of the automotive industry. This approach is critical to overcoming persistent industry challenges, including vendor lock-in, painful integration efforts with non-automotive architectures, poor scalability across vehicle lines, and inflated product prices resulting from limited competition.

Core Advantages of a Chiplet-Based Strategy

Adopting a chiplet-based model offers a clear and decisive set of advantages that directly address the needs of the European automotive market.

Advantage	Description
Scalable and Efficient by Design	Build solutions on a stable, automotive-optimized base die containing core functions (e.g., safety, security, basic I/O). Performance can then be scaled as needed by adding extension dies for high-performance compute, AI acceleration, or other specialized tasks. This "Bring Your Own Die" model allows for unprecedented customization across a vehicle platform.
Flexibility and Cost Control	Mix-and-match the most cost-efficient accelerator dies from a competitive, open market. Chiplets from different process nodes can be combined on a single, cost-effective organic substrate, ensuring that each function is manufactured with the most appropriate and economical technology. This allows for deep customization of critical components like AI accelerators and memory.
Simplified Software & Faster Development	By standardizing on a scalable SoC architecture, chiplets help reduce software complexity and long-term maintenance costs. Development cycles are significantly accelerated because workloads can be isolated in hardware, allowing parallel development and validation efforts with clear separation of concerns, which is a major advantage for complex, multi-supplier projects.

Building the Ecosystem: A Call for Collaborative Action

Technology alone is not enough. For chiplet integration to become a central pillar of European industry, the creation of a vivid, open, and collaborative ecosystem is non-negotiable. Such an ecosystem must be built upon open, specified architectures that foster genuine competition on the die-level, ensuring it remains independent of any single supplier where one company does everything. This openness is fundamental to ensuring the system is tailored to the specific demands of automotive applications rather than being a repurposed solution from another domain.

Realizing this vision requires immediate and concerted initial efforts. New business models must be pioneered, and clear roles and responsibilities for OEMs, Tier 1 suppliers, and semiconductor companies have to be developed and agreed upon. Most importantly, this ecosystem will not emerge on its own; it requires dedicated investment and close, multi-company collaboration to lift it off the ground and establish multi-vendor SoCs as a viable reality.

Adopting a chiplet-based strategy is not just a technical shift; it requires a clear-eyed assessment of where to lead and where to partner. For the European automotive industry to successfully navigate this transition and achieve digital sovereignty, it must exert control over the areas that define the vehicle's unique character and safety, while strategically addressing the areas where competencies are currently lacking. This means leveraging inherent strengths in system-level thinking and safety, and building or sourcing expertise in high-performance silicon design and advanced manufacturing. Europe's deep heritage in automotive engineering provides a strong foundation for defining and controlling the system architecture. These are the domains where European OEMs can and should lead.

1.6.4. Further Semiconductor Demand

Software-defined vehicles (SDVs) are increasing data traffic inside the car much faster than traditional architectures were designed for. More cameras (visible and infrared), radar, lidar, and smart actuators mean more data must move in real time between sensing, compute, and control. At the same time, automotive systems must stay efficient, safe, and affordable over long product lifetimes. This is why there is also a need of a low-complexity chiplet interface.

From a European industrial perspective, this approach is strategically attractive. Europe has strong capabilities in automotive semiconductors, power electronics, embedded control, and advanced packaging, including manufacturing on mature and older technology nodes. Many automotive functions do not require the latest leading-edge node; they require robustness, longevity, qualification readiness, and cost stability. A low-complexity chiplet standard enables Europe to leverage these existing local strengths: producing reliable chiplets on mature nodes, integrating them efficiently, and assembling competitive SDV hardware platforms with reduced dependency on external leading-edge capacity.

Low power is especially important in SDV systems because energy efficiency now affects the full vehicle architecture. Higher interface efficiency means less heat, smaller cooling effort, and better system-level power budgeting. This is beneficial for both electric vehicles, where efficiency links directly to range, and conventional platforms, where thermal limits and reliability remain central.

In sensor and actuator paths, low-latency communication with minimal protocol overhead also improves deterministic behavior for control loops and safety-relevant functions.

Another major advantage is flexibility across heterogeneous components. Future automotive compute platforms will combine different chiplets: AI acceleration, safety controllers, network processing, mixed-signal interfaces, and memory functions. A low-complexity, interoperable interface allows these building blocks to be combined more easily, including multi-vendor combinations, without creating a fully custom interface for each program.

In this sense two goals are addressed at once: technical optimization and strategic resilience. Technically, integration of low complexity chiplets with a low power interface enables scalable bandwidth and low-latency links with lower power and integration complexity. Strategically, it allows Europe to build high-value automotive hardware locally, using available process-node and packaging ecosystems. This combination makes low-complexity chiplet integration a realistic and sustainable path for next-generation vehicles.

We face a critical strategic decision on whether focusing on chiplets and advanced packaging is sufficient, or if Europe must also secure independent, domestic manufacturing capabilities for advanced planar silicon technologies at **2nm and beyond** to protect our supply chains from geopolitical disruptions. Although funding from initiatives like the Chips Act 2 is currently insufficient to build such massive fabrication facilities alone, a unified and highly concentrated European investment strategy is absolutely mandatory to prevent long-term reliance on external regions.

2. Ecosystem Collaboration: Strengthening European Competitiveness

The paradigm shift towards the software-defined vehicle and modular, chiplet-based hardware architectures demands an equally profound shift in how the European industry collaborates. The traditional, linear supply chain—where OEMs define requirements, Tier-1s build black-box systems, and semiconductor vendors supply components in isolation—is no longer adequate to manage the spiraling complexity and pace of innovation. This rigid model creates information silos, slows development, and fails to leverage the full innovative potential distributed across the continent.

> **CROSS-INDUSTRY RELEVANCE AND SHARED EUROPEAN SCALE**
>
> The needs identified in Chapter 1 are not unique to the automotive sector. Industrial automation, robotics, aerospace and defense, energy infrastructure, medical technology, and edge and data-center computing are undergoing comparable transitions toward software-defined, AI-enabled, connected, and increasingly autonomous systems. Across these sectors, computing platforms must combine high performance with real-time behavior, functional safety, cybersecurity, energy efficiency, long product lifecycles, and resilient supply chains. They also face the same structural limitations of monolithic silicon: rapidly rising development cost, limited scalability, long redesign cycles, and dependency on a small number of non-European technology providers.
>
> These common requirements create a strong case for a cross-industry European technology foundation. Processor architectures, AI and HPC accelerators, safety and security building blocks, memory and I/O components, chiplet interfaces, packaging technologies, verification methods, and software toolchains should therefore be developed as reusable platform capabilities rather than as isolated automotive solutions. Sector-specific differentiation can then be implemented through dedicated chiplets, software stacks, qualification profiles, and system configurations on top of shared technical baselines.
>
> A multi-industry approach increases addressable volumes, distributes non-recurring engineering cost across broader markets, and strengthens the business case for European suppliers and infrastructure investments. It also accelerates technology transfer between sectors, reduces duplicated public and private investment, and supports a larger, more competitive supplier ecosystem. The collaboration model described in this chapter should therefore align common requirements across industries wherever possible, while preserving the distinct safety, certification, lifecycle, and performance needs of each application domain.

To build the complex, secure, and scalable systems of the future, a new model is required: a dynamic, networked ecosystem where partners work together in a far more integrated and transparent fashion. This chapter explores the frameworks and principles needed to foster such an environment. It will address how Original Equipment Manufacturers (OEMs), Tier-1 suppliers, semiconductor companies, and other partners can collaborate effectively in a pre-competitive setup, defining a new blueprint for European industrial strategy that prioritizes collective strength over isolated efforts.

This overarching ecosystem model is guided by three fundamental, interconnected principles designed to maximize resource efficiency and accelerate technological deployment across the continent:

Ecosystem Principle	Strategic Objective	Operational Impact
Cross-Industry Convergence	Converging automotive, industrial, and silicon ecosystem requirements.	Unlocks massive economies of scale by allowing silicon IP and physical chiplets to be co-monetized across automotive, industrial automation, and edge-computing markets.
Cross-Functional Alignment	Co-designing software-defined vehicle (SDV), AI, HPC, and safety/security requirements directly with physical silicon architectures.	Eliminates hardware-software development silos early in the design cycle, ensuring that sovereign silicon is pre-optimized for safety, security, and high-performance software execution.
Vendor-Neutral Foundation	Building on open, non-proprietary standards to ensure seamless interoperability and prevent proprietary lock-in.	Fosters a highly competitive, modular marketplace where diverse suppliers can contribute compatible dies, securing supply chain resilience and lowering entry barriers.

A New Paradigm: Pre-Competitive Alliance on the Scale of a Single European Enterprise

Incremental improvements to the traditional, arms-length automotive supply chain are no longer sufficient. To master the immense complexity of the software-defined vehicle (SDV) and to compete on a global scale, Europe's leading industrial players must move beyond simple, siloed partnerships. They must adopt a radically new, deeply integrated collaborative model: operating in the pre-competitive space as a cohesive, virtual European enterprise.

This means abandoning the "black box" mentality and the friction of isolated development in favor of a synchronized ecosystem. It requires a profound shift in mindset where pre-competitive collaboration is not merely a side project, but the core of the strategy. Within this model, OEMs, Tier-1s, and semiconductor specialists align their pre-competitive efforts—acting with the agility and shared focus of integrated departments within a single virtual enterprise—to collectively establish and validate the open-standard hardware and software platforms that will form the foundation of next-generation mobility. Crucially, this model maintains a strict separation between joint pre-competitive development and individual, highly competitive downstream commercialization. This chapter outlines this "Europe Inc." philosophy and details the essential, legally compliant framework required to make this bold vision a reality.

The "Europe-centric" Philosophy: A Structured Model for Pre-Competitive Work

To collaborate with the speed of "one company" is not a metaphor; it is an operational mandate. It requires partners to commit to a level of integration and transparency that fundamentally redefines their relationships in the pre-competitive space while remaining strictly within the boundaries of EU competition law. This philosophy is built on three core, antitrust-compliant principles:

1. **Unified Goals and Shared Technical Success**
   In a single company, the success of the chip design department and the software application department are inextricably linked to the success of the final product. Similarly, the European automotive semiconductor ecosystem must define unified, pre-competitive technical objectives (e.g., *"Validate a scalable, certifiable, multi-vendor reference compute platform by 2029"*). The technical KPIs for all partners in the collaboration must be tied to the success of this shared standard, rather than their individual company's narrow proprietary interests. By focusing entirely on joint standardization, interoperability, and aiming high Technology Readiness Levels, success or failure is shared by all, while final product pricing, vehicle release dates, and market positioning remain strictly independent and competitive.

2. **Antitrust-Compliant Transparency and Open Communication**
   While departments within a single company share all data freely, a cross-company alliance must implement rigorous safeguards to prevent the exchange of Commercially Sensitive Information (CSI). Under this model, "radical transparency" is strictly directed at non-commercially sensitive engineering data: Jointly sharing technical roadmaps (such as open RISC-V ISA extensions), interface specifications (such as UCIe baselines), safety/security profiling, and simulated performance metrics to resolve integration bottlenecks in real-time.

   To prevent anti-competitive spillover, information-barrier guidelines and secure "clean-room" environments—managed by neutral third-party facilitators—allow to aggregate and anonymize technical data, ensuring that "fail fast, fix fast" engineering is achieved without legal risk.

3. **Integrated, Cross-Functional Reference Teams**
   To successfully execute complex projects, industry partners must transcend traditional transactional contracts and establish highly integrated, cross-functional networks of OEMs, Tier-1s, and silicon providers working as a unified team on reference architectures. These joint engineering teams work from a single project backlog focused on standard-setting, validation frameworks, and initial platform bring-up. Because this level of multi-company intimacy raises intellectual property challenges, IP governance must be managed proactively at the overarching alliance level:

   - **IP Protection & Licensing:** Establish a robust pre-competitive trust framework based on Fair, Reasonable, and Non-Discriminatory (FRAND) licensing terms, ensuring that standard-essential patents are pooled to prevent legal bottlenecks.
   - **Operational Separation:** By resolving these legal boundaries at an overarching alliance level, engineering teams are empowered to focus entirely on co-innovation, secure in the knowledge that proprietary corporate assets are protected and that downstream customization remains entirely in the hands of the individual alliance members.

Concluding Remarks: Scaling the Efficiency of Joint Undertakings

Europe does not need to build its collaborative foundation from scratch, as numerous highly successful co-innovation models already thrive within established frameworks such as the **Chips Joint Undertaking (CHIPS-JU)**. These public-private partnerships have proven that pooling resources across industry and research boundaries within structured, legally safeguarded frameworks can successfully drive pre-competitive technological breakthroughs.

To translate these achievements into rapid, market-ready industrial solutions, however, we must actively steer and monitor the operational depth of this tight inter-company collaboration. Introducing structural tracking mechanisms for cross-partner dependencies, synchronized development milestones, and shared performance indicators within consortia will prevent fragmentation and maximize the return on public and private investments. By closely monitoring these collaborative dynamics—under the strict guidance of pre-defined antitrust compliance protocols—we can optimize the execution efficiency of our joint undertakings and accelerate the path from pre-competitive research to vertical integration at the OEM level.

3. Roadmap Phases and Key Milestones enabling Technology Introduction in Europe

The successful development and deployment of a European SDV Hardware Platform require a strategic, multi-faceted approach that balances long-term research with near-term industrial needs. This roadmap is structured around **three parallel, interconnected phases**, each with a distinct focus on Technology Readiness Level (TRL), research and innovation demands, and specific outcomes.

By running these phases concurrently, we can simultaneously lay the groundwork for future breakthroughs while empowering the industry to act today. This parallel execution ensures that foundational research informs next-generation prototypes and that early industrialization efforts provide critical feedback to the entire ecosystem, creating a virtuous cycle of innovation.

EUROPEAN AUTOMOTIVE COMPUTE PATHWAY

Step 1: Establish European HPC capability with mature processor ecosystems.

Use production-ready commercial IP while building European capabilities in SoC design, AI acceleration, safety, security, software, and advanced-node implementation.

Step 2: Establish an open European chiplet ecosystem.

Create ISA-agnostic interfaces, software portability, advanced packaging, multi-vendor qualification, and clear system-integration responsibilities.

Step 3: Progressively introduce European RISC-V technologies.

Start with mature real-time, safety, and control functions and expand toward high-performance application processing once IP, tools, software, safety, and commercial maturity are competitive.

3.1. Roadmap Phases in the context of ECAVA

The three phases—Build the Foundation, Prototypes for next generation, and Empower & De-risk industrialization—are designed to mature technologies and capabilities from basic research through to series production readiness. Each phase addresses a different time horizon for vehicle Start of Production (SOP) and engages different actors within the European ecosystem.

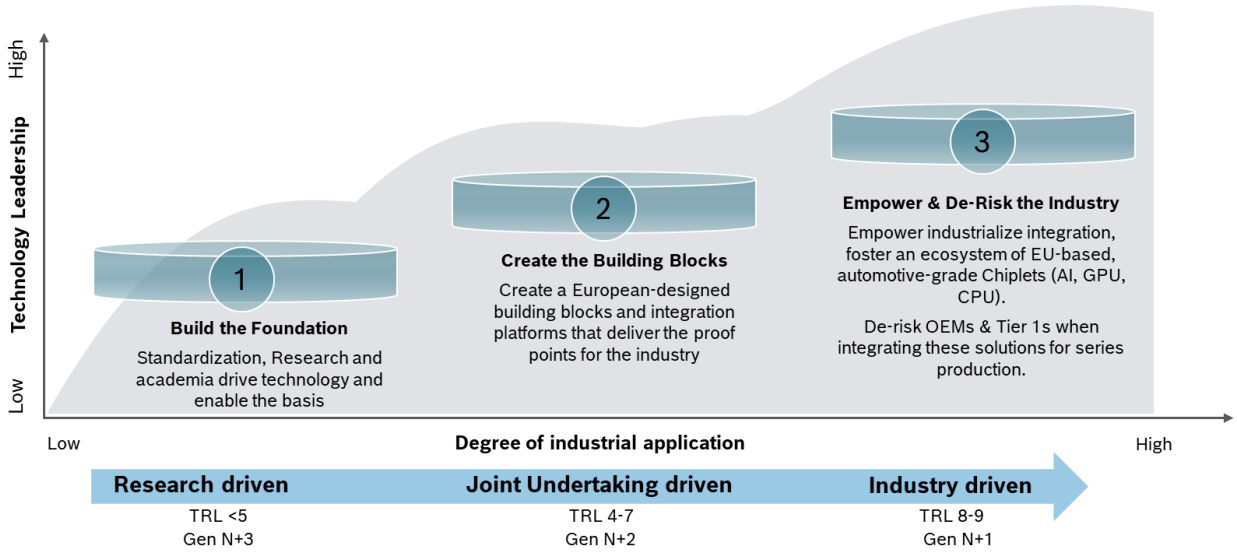


Figure 4: Research, innovation and deployment programs are categorized along three phases, continuously increasing the degree of industrial application and technology leadership.

Interdependence of the Three Technology Pillars

The three technology pillars form one integrated but maturity-sequenced system roadmap. In the near term, the HPC pillar should build European capabilities around mature processor IP and software ecosystems, while European AI acceleration, safety, security, system architecture, and advanced-node design capabilities are developed. In parallel, the chiplet pillar should establish ISA-agnostic interfaces and multi-vendor integration capabilities. The RISC-V pillar should mature automotive-grade processor IP, software, tools, safety assets,

and commercial support so that RISC-V components can be progressively introduced into later platform generations.

Progress in one pillar therefore depends directly on progress in the others. RISC-V processor IP delivers strategic value only when it can be integrated into competitive HPC systems with mature software, safety, and verification support. European HPC components become scalable product families only when standardized die-to-die interfaces, advanced packaging, thermal concepts, and multi-vendor qualification are available. Conversely, chiplet integration creates industrial value only if competitive processor, accelerator, memory, safety, and I/O building blocks exist to populate the platform.

The three parallel roadmap phases must consequently be synchronized across all pillars. Shared requirements, interfaces, demonstrators, readiness gates, and industrialization decisions should be managed jointly so that outputs converge on common European reference platforms. This integration turns three individual technology roadmaps into a coordinated path from foundational IP, through system prototypes, to qualified and commercially deployable European compute platforms.

This sequencing avoids delaying European HPC and chiplet industrialization while preserving RISC-V as an important path toward greater long-term technological control.

3.1.1. Build the Foundation [Target: N+3, Vehicle SOP ~2037; TRL <5]

This phase is the engine of long-term innovation. Focused on fundamental research with a low TRL, it brings together academia, research institutions, and corporate R&D to explore and validate future technology building blocks. The primary goals are to drive standardization efforts, investigate next-generation materials and architectures (e.g., neuromorphic computing, photonics), and build the scientific basis needed to secure Europe's technological leadership for the next decade and beyond.

3.1.2. Prototypes for next generation [Target: N+2, Vehicle SOP ~2034; TRL 3-7]

Bridging the gap from research to reality, this phase concentrates on creating tangible proof points. The central objective is to develop a European-designed Integration Platform where promising technologies from the foundation phase can be tested and validated in a representative environment. This involves building functional hardware and software prototypes that demonstrate the feasibility and benefits of the chiplet-based approach, providing the industry with the confidence to invest in these emerging solutions.

3.1.3. Empower & De-risk industrialization [Target: N+1; Vehicle SOP ~2031; TRL 8-9]

This phase is sharply focused on immediate industrial impact. Its mission is to empower the ecosystem by fostering the creation and qualification of automotive-grade chiplets (AI, GPU, CPU) from EU-based suppliers. Concurrently, it seeks to de-risk the adoption of these solutions for OEMs and Tier 1s by establishing clear integration pathways, defining robust testing and validation procedures, and ensuring a resilient supply chain for series production. This phase is critical for translating technological potential into commercial success in the near term.

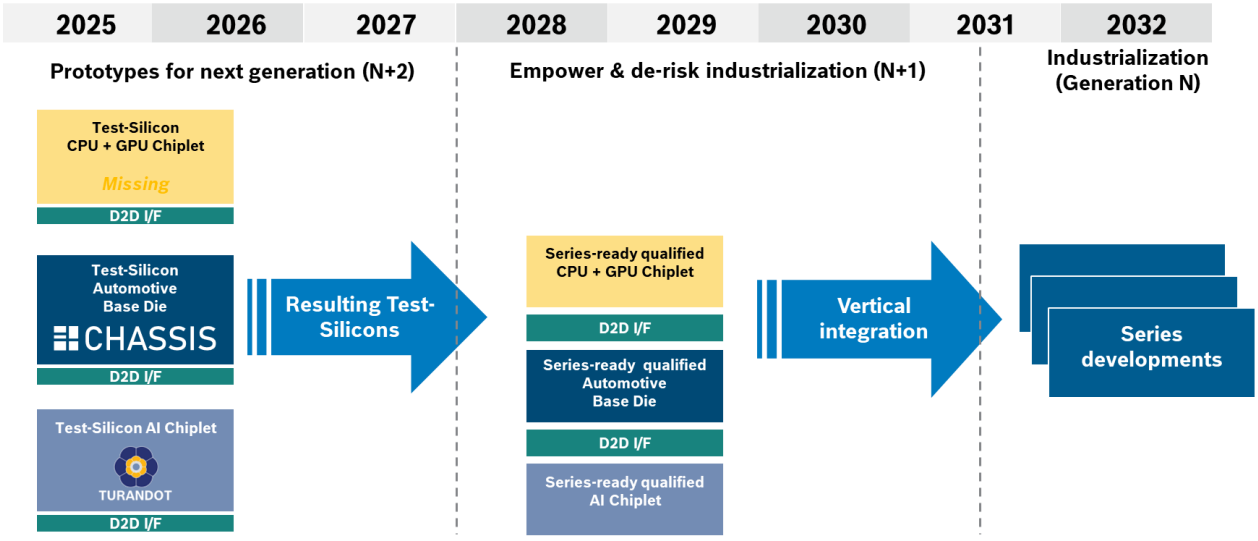


Figure 5: Exemplary transition from "Next-gen Prototypes" towards "Empower & De-risk", accelerating industrialization of technology.

3.2. RISC-V Roadmap

3.2.1. Build the Foundation

On September 8th, 2022, a European Working Group launched a report entitled: "Recommendations and roadmap for European sovereignty on open source hardware, software and RISC-V technologies"[2]. That document formed the basis to launch two projects in the Chips-JU framework program (TRISTAN and ISOLDE) and another Working Group to develop a "Roadmap for Future Automotive-driven RISC-V Developments in Europe". This latter Working Group has concluded their report and roadmap on 16 February 2024[3].

The roadmap encompasses a wide array of components including microcontrollers, domain/zonal control units, gateway controllers, and application processors. Given the advent of autonomous driving technologies and the integration of artificial intelligence within vehicles, this document also considers the development of (Artificial Intelligence) AI accelerators.

The RISC-V instruction set architecture (ISA) presents an ideal opportunity for Integrated Device Manufacturers (IDM), Tier 1s and also automotive OEMs to leverage on the flexibility provided by an open-source ISA. At the same time, a degree of standardisation and collaboration is necessary to avoid unnecessary redundant duplication of work. Through the reference platform approach proposed by this roadmap, European industry will have a springboard for further innovation in automotive electronics.

The roadmap has been split into 4 key elements or categories of the overall RISC-V based processors:

- Scalable RISC-V Automotive Control Processors
- High-Performance RISC-V based Application Processors
- AI and ML Accelerators
- System integration and interfacing

In addition, the roadmap has defined a number of milestones, which are a mixture of FPGA-based implementations and real silicon devices.

The four elements of the roadmap are linked to each other. The first contains automotive-specific real-time control processors. They are the starting point for the second category, which is more generic and applicable for other type of application domains (communications, industrial, health, etc.). One example is industrial robotics – for advanced motion control, computer vision, inspection, and health and safety applications – and drones/UAVs, where a series of grand challenges is emerging including sensor fusion and processing for detect and avoid applications in a low-cost, low-weight, low-power form factor. The third category contains accelerator elements which are specific developments for both automotive as well as other domains, and which contains starting points for next generation processors. As these accelerators are a key component of both processor families, we have opted to make a separate category to show the different IP elements to develop in the coming years. The fourth category contains specific elements for system integration, more in particular how controllers and processors fit into global system solutions.

The four elements can be considered as linked activities as part of a reference platform. By definition, this needs to include the computational elements along with the supporting infrastructure required to realise the whole platform. Therefore, the reference platform infrastructure roadmap laid out includes both SoC and chiplet interconnects.

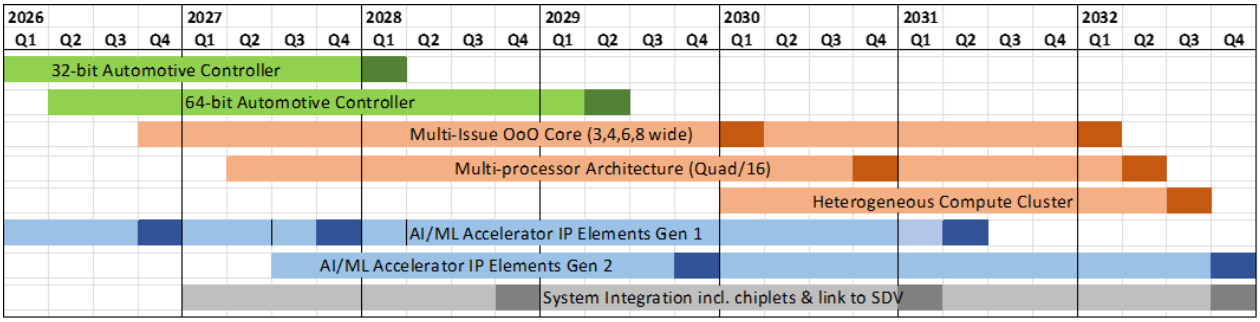
Along with the core (CPU) clusters, three key components are needed to sustain the high memory bandwidth needed by the processor, as well as controllability features towards supporting, to a sufficient extent, freedom from interference across real-time applications with integrity requirements. Those components are the following. (1) shared cache memories capable of supporting a large number of outstanding requests from a high number of out-of-order cores, both in terms of (data) bandwidth as well as in terms of managing pending requests minimizing serialization. (2) High-bandwidth interconnects with Quality of Service (QoS) minimizing contention and allowing some control on the traffic priorities through the deployment of programmable QoS support. And (3), high-bandwidth memory controllers with QoS support able to manage out-of-order requests, minimizing serialization, and allowing some controllability on the traffic priorities.

[2] Recommendations and roadmap for European sovereignty on open source hardware, software and RISC-V Technologies | Shaping Europe's digital future (europa.eu)

[3] https://ecssria.eu/Roadmap_RISC-V_v240216_Final.pdf

The goal of deploying the 4 elements together is achieving key goals for automotive safety-critical applications as well as for comparable applications in other domains. Whenever automotive control processors deliver enough performance for a given such application, they are the candidate computing elements since they provide on their own enough features (e.g., fault tolerance support, safety measures, etc.) to allow reaching the highest integrity levels. Some safety-critical applications require higher levels of performance and using high-performance application processors becomes unavoidable. However, those often lack appropriate support to reach the highest integrity levels on their own, and cooperation with the control processors is needed. Hence, the resulting system is more complex, with impact on cost and time-to-market. A third class of applications includes those massively parallel that require extraordinary levels of computing power on one hand, and extreme power efficiency on the other to provide results with tight real-time constraints and within tight power envelopes. Those applications need to resort to accelerators where programmability and maintainability on one hand, and efficiency on the other, need to be traded off. Those accelerators may not provide complete safety support and cooperation with the automotive control processors is often needed to allow realizing the safety concept on the platform. Finally, control and application processors are often multicores capable of running multiple applications potentially with mixed criticalities and heterogeneous performance needs. Hence, appropriate data storage and interface support is needed to allow computing elements execute applications at a sufficient speed to meet their requirements. Hence, caches, interconnects and memory controllers must come along with appropriate features to allow computing elements execute applications fast enough and support mixed-criticality execution.

A high-level view of the roadmap is given in the picture below.



3.2.2. Prototypes for next Generation

The objective of this phase is not to mandate RISC-V adoption, but to close the maturity gaps that currently prevent its use as a competitive automotive HPC application processor. Progression into series-oriented platforms should be based on measurable readiness gates covering processor performance and efficiency, production-quality IP, qualified software tools, operating-system and hypervisor support, workload portability, functional-safety evidence, cybersecurity, automotive qualification, and dependable lifecycle support.

The RISC-V ecosystem, while growing rapidly, still lacks the maturity, diversity, and proven industrial readiness required for large-scale deployment in the automotive domain. In particular, it falls short of providing multiple, competitive alternatives for each essential building block needed to assemble robust automotive SoCs. Industrial interest is increasing, but widespread adoption remains constrained by technical gaps, lack of strong support infrastructure, limited availability of high-quality IP, and uncertainty regarding compliance with stringent automotive requirements.

A new project, called TURANDOT, has been defined by a consortium of IDMs, TIER1s, OEMs and RTOs, and has officially started in August 2026. TURANDOT sets out to bridge this gap by building a next-generation European ecosystem for RISC-V processors that addresses the needs of future automotive electrical/electronic (E/E) architectures. The project is structured around an Application Processor-centric European automotive RISC-V platform nucleus, supported by real-time processors, AI accelerators, software tools, interfaces, benchmarking methods and integration assets. This includes the development of high-performance, real-time and application-specific microprocessors, AI accelerators and Application processors accompanied by a shared IP library, a reference hardware repository, and compatible software toolchains specifically tailored for automotive use cases.

The overarching ambition of TURANDOT is to develop an Application Processor-centric RISC-V automotive platform, combining high-performance application processor solutions with real-time processors, AI accelerators and an IP library, validated and benchmarked against state-of-the-art competition. The goal is to deploy the next-generation E/E architectures in future digital vehicles to make them most competitive in terms

of performance vs. non-RISC-V alternatives. TURANDOT is committed to ensuring that its technological developments will align with industry trends and needs while fostering scalability, interoperability, and reusability across the European Automotive sector. TURANDOT will demonstrate its results through real silicon tape-outs, industrial platform clusters and two Application Processor-centric demonstrators. These demonstrator lines act as the main system-level validation endpoints, while the other demonstrators and research satellites feed into the wider industrial platform-cluster structure.

Building on and leveraging the substantial achievements of previous projects, such as TRISTAN, ISOLDE, and RIGOLETTO – while also aligning with other relevant initiatives supported by Chips-JU – such as the SDV-related initiatives under the HAL4SDV, Federate, CHASSIS and Shift2SDV projects, TURANDOT will push the frontier of European semiconductor innovation, taking integration and silicon realization to the next level with ambitious tape-out goals and boosting TRL up to 8. The developed solutions will not only be state-of-the-art in design, but also ready for industrial deployment.

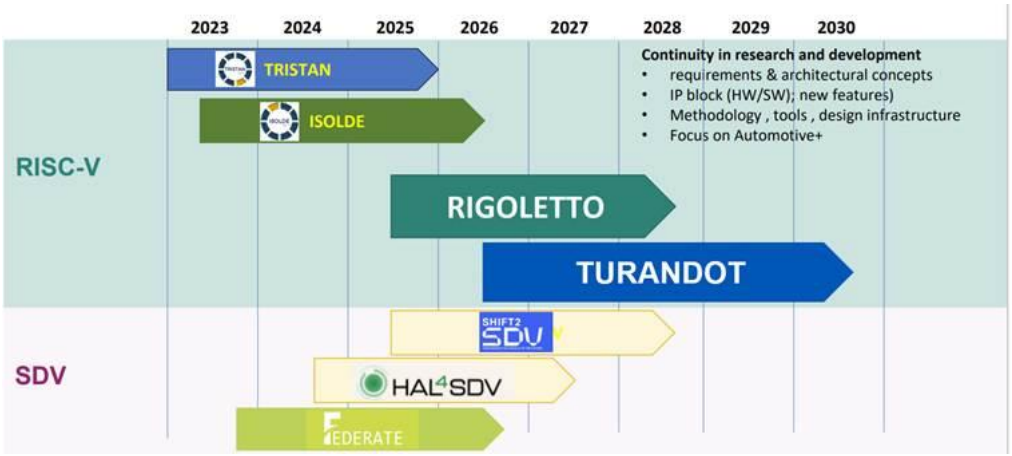
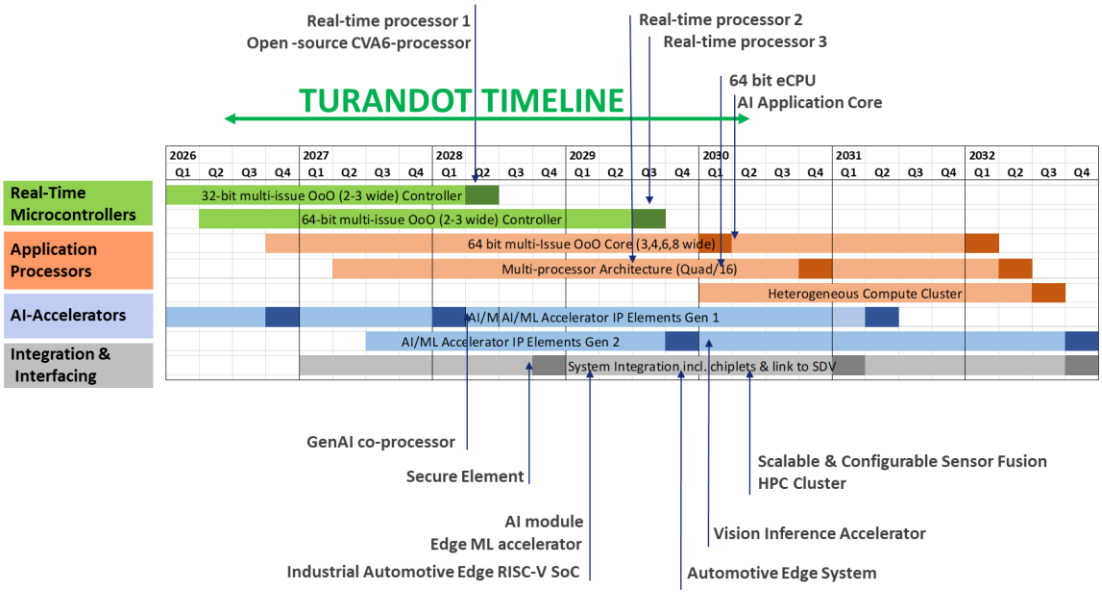


Figure 6: Positioning of TURANDOT in the EU initiative ecosystem

To ensure coherence and interoperability across the multiple industrial design tracks in TURANDOT, the project adopts a unifying platform-based approach, anchored by the Quintauris automotive-grade RISC-V reference platform. This platform-based approach will be implemented through a limited number of industrial platform clusters, which consolidate individual silicon, IP, software and integration activities into coherent automotive system contexts. Far from operating in silos, the IDMs will contribute to and benefit from this shared platform, which provides standardized interface specifications, benchmarking methodologies, and integration workflows.

Quintauris, acting as a neutral enabler, will coordinate the convergence of individual IP contributions into a Common Automotive Repository, facilitating early validation, alignment with functional safety and real-time constraints, and rapid market deployment. The repository and platform activities will focus on shareable and pre-competitive artefacts, including interface profiles, integration guidelines, benchmarking methods, IP-XACT packaging principles, safety/security assumptions and metadata for reusable automotive RISC-V components. This harmonized foundation not only supports technical comparability across demonstrators but also provides a sustainable framework for future adoption and certification of RISC-V automotive processors and accelerators, ensuring that the output of TURANDOT contributes to a common European platform and avoids fragmented outcomes. Together, these shared requirements, interface profiles, benchmarking methods, repository metadata, IP-XACT packaging principles and demonstrator mapping rules form the Common Automotive Platform nucleus of TURANDOT, enabling differentiated industrial implementations to remain interoperable, comparable and reusable.

Similar to the TRISTAN and ISOLDE projects, TURANDOT will contribute to the roadmap. The contributions of the TRISTAN & ISOLDE projects to the roadmap can be found in Section 4.1 (ANNEX: RISC-V – Detailed Roadmap Document).



In addition to monitoring progress on the roadmap implementation via collaborative projects and direct contacts with key stakeholders in the field, we are also planning to update the current roadmap as the world is not standing still and current geo-political situations enforce us to get a stronger European footprint in the RISC-V processor market.

The European RISC-V roadmap will be updated covering both R&D as well as funding & business aspects in all power and performance ranges from deeply embedded to high-end computing, incl. automotive security and safety requirements.

Thereby, the goal is to get endorsement of all relevant European automotive industrial stakeholders throughout the value chain and further develop and grow the European RISC-V hardware ecosystem from architectural exploration & design up to certification of silicon.

A high priority topic is to align the HW roadmap with other initiatives, more in particular SDV, HPC and chiplet integration developments.

3.2.3. Empower & De-risk

This phase is important to ensure that all developments have industrial impact throughout the value chain. European IDMs want to stay market leaders in the automotive semiconductor domain (3 of them are currently in the Top 5 of worldwide suppliers). In addition they want to ensure their developments are embedded in future automotive E/E architectures. The goal is to empower the ecosystem by creating different types of automotive qualified, ranging from high-performance central-compute to different types of real-time zonal processors. This phase is critical for translating technological leadership into commercial success in the near term.

The current thoughts on this empowerment have been taken in the TURANDOT project. The TURANDOT consortium will plan a number of OEM-specific workshops to get explicit feedback on processor requirements and demonstrator needs during the first 6 months of the project. OEM inputs are also reflected through formal review and acceptance milestones for the culminating demonstrator, which clarifies their gatekeeper role. these reviews are expected to focus on system relevance, robustness of the supply chain, and fit with future OEM E/E architectures, rather than on the detailed processor architecture itself. Discussions with OEMs on potential participation in TURANDOT have clearly confirmed their interest in high-performance automotive applications supported by AI-enabled real-time processors. This reflects a broader need to reduce dependency on single, centralised compute solutions and to enable a more distributed in-vehicle computing network. Furthermore, OEMs are not selecting a processor architecture in isolation; rather, they are assessing whether the overall platform strengthens resilience, flexibility, and supply-chain robustness across the vehicle architecture. This direction is fully aligned with the TURANDOT approach, which combines high-performance application processing, real-time processing and AI acceleration within a modular European RISC-V automotive platform.

Where these readiness gates have not yet been achieved, the platform should retain mature commercial processor IP while preserving the interfaces and software portability required for later substitution.

3.3. HPC Roadmap

This section translates the European HPC family targets defined in Section 1.6.2 into phase-specific commitments. The three phases run concurrently, in line with the parallel-execution principle established in Section 4.1: foundational research informs next-generation prototypes, and early industrialization provides feedback to the broader ecosystem.

Each phase delivers one of the three HPC product tiers defined in Section 1.6.2. Empower & De-risk delivers the Entry / Mass Market tier for SOP ~2031. Prototypes for next Generation delivers the Mid / Premium tier for SOP ~2034. Build the Foundation enables the High / L4 tier for SOP ~2037. The phases share the same architectural principles (AP, AIA, and safety island as distinct components; chiplet-based integration where the technology supports it; multi-market viability; OEM-aligned timing) and differ in the technology base they draw on, the TRL of their deliverables, and the ecosystem actors engaged.

The milestones below are indicative and to be aligned with OEM platform decision points and with the parallel commitments in the RISC-V, Chiplet Integration, and Further Semiconductor Demands roadmaps.

3.3.1. Build the Foundation

This phase is the research backbone for the upper tier of the European HPC family. The silicon products it enables (2,500 to 5,000+ sustained 8-bit TOPS/TFLOPS, 800+ GB/s memory bandwidth driven by 3D-stacked memory, FP4 through FP16 precision, 1,500 to 3,000 kDMIPS application processor performance, 150 to 300+ W module envelope) cannot be reached by extrapolating today's silicon. They depend on technology directions currently at low TRL that require coordinated EU investment to mature in time for SOP 2037.

The technology directions below are not automotive-specific. Each is being developed in parallel for European data center and AI infrastructure applications, where higher per-unit silicon margins and shorter qualification cycles allow faster innovation. The strategic role of this phase is to ensure that as those technologies mature in the data center, they are transferred to automotive-grade qualification in time for SOP 2037. The multi-market viability principle stated in Section 1.6.2 makes this co-development not just possible but a precondition for cost-justifying the investment.

Technology directions to fund in this phase:

- Silicon photonics and on-package optical I/O, to break the bandwidth and energy-per-bit limits of electrical interconnect for chip-to-chip and chip-to-memory links. Automotive adoption follows once thermal and reliability envelopes are validated.
- Advanced memory integration: 3D-stacked memory, near-memory compute, and compute-in-memory architectures, transferred from data center HPC applications where they are maturing first. The binding constraint for on-device VLA and LLM inference at automotive cost and power points, and the single most important technology direction for the upper tier.
- Novel packaging and heterogeneous integration: hybrid bonding, wafer-level integration, and 2.5D / 3D stacks qualified to automotive reliability and lifecycle requirements.
- High-performance European Application Processor cores at the kDMIPS class required for the upper tier, with the architectural trajectory defined by the RISC-V Roadmap (Section 4.2). This closes the AP gap identified in Section 1.6.2 over the longest time horizon; the medium-term steps are in the Prototypes phase below.

Ecosystem engagement. European research institutions, university programs, and corporate R&D, anchored in the IAG Industry Policy Blueprint Focus Areas 3, 5, 6, and 7. OEMs and Tier 1s participate as requirement validators rather than direct customers in this phase. Coordination with European data center HPC programs is the principal cross-domain dependency.

Readiness gates:

- 2028: technology down-select on photonic interconnect topologies, advanced memory architectures, packaging approaches, and AP core direction.
- 2031: technology demonstration in representative test vehicles or surrogate platforms.
- 2034: handoff to the Prototypes phase for product integration.

3.3.2. Prototypes for next Generation

This phase delivers the mid-tier European HPC family (1,000 to 2,500 sustained INT8 TOPS, 300 to 500 GB/s memory bandwidth, INT8 / FP8 / FP4 / FP16 precision, 600 to 1,500 kDMIPS, 80 to 180 W module envelope) and

the shared European integration platform on which it is demonstrated. The focus is on translating research outcomes from the Build the Foundation phase into functional silicon, on proving chiplet-based mid-tier HPC at vehicle-relevant performance and power, and on initiating the European Application Processor development effort that the Empower phase does not address.

Silicon and integration deliverables:

- AIA chiplet at TRL 5-6 by 2030, integrated in a working multi-die package on a UCIe-mature interposer, with the AIA architecture inheriting the 3D-stacked memory direction from the Build the Foundation phase.
- Initiation of a European high-performance Application Processor development effort. The AP gap identified in Section 1.6.2 is the most strategically significant unaddressed component in European HPC silicon planning, and the Prototypes phase is the window to close it. Functional European AP silicon at TRL 5-6 by 2032, aligned with the RISC-V auto-grade AP track in Section 3.2.
- Safety island chiplet co-designed with the safety architecture defined in Section 2.6.2.
- Advanced packaged memory subsystem, drawing on the 3D-stacked memory work maturing in parallel for European data center HPC.
- Reference software stack reaching functional parity with incumbent platforms for AD workload portability.

Technology dependencies. Chiplet packaging maturity (linked to Section 3.4 Chiplet Integration Roadmap), advanced memory adoption, and the European RISC-V auto-grade AP and safety cores reaching production readiness (linked to Section 4.2). The prototypes phase is the convergence point where the four ECAVA technology streams (HPC, Chiplet, RISC-V, Further Semiconductor Demands) must be co-funded; misalignment in any one of them blocks the phase.

Ecosystem engagement. At least two European OEM platform commitments by 2030 and at least one Tier 1 integration partner. The European integration platform is shared across the ECAVA HPC, Chiplet, and RISC-V streams to compound investment and to give multiple silicon candidates a common validation environment.

Readiness gates:

- 2030: chiplet prototypes (AIA, safety island) available for OEM evaluation.
- 2031: integration platform demonstration on representative test vehicles.
- 2032: European AP silicon at TRL 5-6; silicon validation complete; transition to production-ready industrialization.

3.3.3. Empower & De-risk

This phase delivers the entry / mass-market European HPC family (300 to 1,000 sustained 8-bit TOPS/TFLOPS, 100 to 200 GB/s memory bandwidth driven by LPDDR6-class memory, INT8 / FP8 / FP16 precision, 100 to 500 kDMIPS, 30 to 80 W module envelope) on technology nodes and packaging approaches available today. The foundational silicon for this tier is being developed under the EU-funded Turandot project, which covers the European AI Accelerator chiplet targeting tape-out in 2029 and the RISC-V microcontroller building blocks for the safety island and control plane. The Empower & De-risk phase funds what Turandot does not: the industrialization runway from first silicon to OEM-qualified production, including ASIL-D system-level certification, Tier 1 integration support, OEM anchor coordination, and the EU-sovereign safety pipeline. It is the phase that translates European HPC ambition into vehicles on the road by 2031. The platform decision window for SOP 2031 vehicles closes in 2027 to 2028, which means silicon definition must be locked by 2027 and engineering samples in OEM hands by 2030.

Silicon products:

- Entry HPC AI Accelerator, with the compute, memory, and power characteristics above.
- Chiplet implementation using mature UCIe-class die-to-die interconnects, enabling product family scaling through AIA / AP / memory chiplet combinations.
- AIA at ASIL-B with diagnostic coverage, combined with a dedicated ASIL-D safety island and a mixed-criticality-capable AP cluster, delivering ASIL-D at the system level. Cybersecurity per in accordance with ISO 21434 for the engineering process, and with J3101 and EVITA recommendations for the hardware and software architecture, including a Hardware Security Module (HSM)-based Hardware Root of Trust.

- Tape-out target 2029; engineering samples 2030; production qualification 2030 to 2031.

Technology base. Established process nodes (5 nm to 3 nm), LPDDR5X/6-class memory, conventional advanced packaging where applicable. AP cluster baseline uses licensed high-performance commercial CPU IP (Arm AE-class) in the SOP 2031 window: European RISC-V cores developed under Turandot are integrated where they are production-ready, principally on the safety island and in control-plane functions.

Ecosystem engagement. The Empower phase complements Turandot, which funds silicon design and first silicon, by funding the industrialization layer Turandot does not cover: ASIL-D system-level safety certification, Tier 1 integration validation, EU-sovereign safety certification pipeline, and OEM platform anchor coordination. Concrete commitments needed in this phase: anchor customer commitments from at least two European OEMs by 2027; Tier 1 integration validation by 2029; EU-sovereign safety certification pipeline operational by 2030, removing the current dependency on non-EU IP providers for the safety documentation underpinning ASIL-D system-level certification.

Funding rationale. The Empower phase represents the industrialization step that converts Turandot silicon outputs into deployable automotive HPC products. Without coordinated EU funding for this phase, the existing investment in Turandot delivers working silicon that European OEMs cannot adopt at the SOP 2031 cycle: the gap is not technology, it is the ASIL-D certification dossier, the Tier 1 integration runway, the OEM anchor coordination, and the EU-sovereign safety pipeline. Each of these is fundable at modest scale relative to the silicon design effort already committed, and each is the binding constraint for converting Turandot outputs into vehicles by 2031.

Readiness gates:

- 2027: OEM platform decision and silicon definition lock.
- 2028: silicon design freeze; customer evaluation kits (FPGA / emulation); Tier 1 integration architecture defined.
- 2029: tape-out; first silicon back.
- 2030: engineering samples in OEM and Tier 1 hands; ISO 26262 ASIL-D system-level certification dossier complete; integration validation in OEM development vehicles.
- 2030 to 2031: production qualification.
- 2031: vehicle Start of Production.

3.4. Chiplet Integration Roadmap

Chiplet integration is a strategic enabler for future European automotive hardware platforms. It allows complex SoCs to be decomposed into reusable, scalable, and interoperable building blocks. Instead of developing every future vehicle computing platform as a monolithic full-custom chip, chiplet systems enable a modular architecture where compute, AI, safety, security, memory, I/O, and communication functions can be combined according to the needs of different vehicle classes, performance levels, and market segments.

The roadmap is structured into three parallel phases. These phases are intended to be executed in parallel. They address different time horizons, maturity levels, and ecosystem needs in parallel: long-term technology foundations, mid-term ecosystem enablement, and near-term industrial deployment.

3.4.1. Build the Foundation

Target horizon: 2037 to 2040+ | Technology readiness: TRL below 5

This phase establishes the scientific, architectural, and manufacturing foundations for long-term European sovereignty in chiplet-based computing. Its purpose is not to create immediate products, but to mature the technologies that will define future vehicle, industrial, robotics, aerospace, and critical-infrastructure platforms. Europe must strengthen its position in advanced packaging, die-to-die communication, interoperable Networks-on-Chip, reusable safety and security components, and future compute architectures.

A central priority is advanced packaging and heterogeneous integration. Research should address glass substrates, silicon and organic interposers, 2.5D and 3D integration, co-packaged memory, integrated photonics, and automotive-grade thermal management. These capabilities are needed to reduce dependence on non-European packaging ecosystems and to make large heterogeneous systems reliable and economically viable over long product lifecycles.

Open and efficient communication technologies form the second foundation. Future platforms require low-cost and ultra-low-power die-to-die interfaces, coherent inter-chip communication, package-scale links, and eventually optical connections where electrical interconnects reach limits in bandwidth, reach, or energy efficiency. These interfaces must be complemented by vendor-independent coherence models, common transaction protocols, standardized NoC abstraction layers, and interoperable resource management. The goal is that chiplets from different organizations can operate within one system fabric without dependence on a single NoC or interface supplier.

Europe should build on its strengths in functional safety and cybersecurity by developing reusable safety controllers, security islands, trust anchors, hardware security modules, and system supervisors. These components should be designed for use beyond automotive markets, improving their economic viability while supporting sovereign control of critical functions. Research must also prepare for active, self-protecting architectures that detect abnormal behavior, isolate compromised chiplets, and recover safely during operation. Long-lifecycle platforms additionally require cryptographic agility and efficient post-quantum security building blocks, including quantum-safe boot, attestation, key management, and update mechanisms.

The architecture itself should become more adaptive and distributed. Research should investigate fine-grained compute and memory disaggregation, heterogeneous accelerator integration, runtime workload allocation, adaptive power and thermal management, fault-aware orchestration, and self-healing behavior. Sensor-in-package concepts can move preprocessing and AI inference closer to radar, lidar, imaging, and other data sources, reducing raw data movement, latency, and energy consumption. Optical die-to-die and package-to-package communication should be explored as a longer-term option for massively interconnected AI-native systems.

By the end of this phase, Europe should have a portfolio of sovereign technology building blocks: open interfaces, interoperable NoC concepts, advanced and cost-competitive packaging, reusable safety and security chiplets, quantum-safe hardware components, intelligent sensor packages, optical communication technologies, and self-adaptive system architectures. These results should be specified from the outset for interoperability and multi-vendor integration.

3.4.2. Prototypes for Next Generation

Target horizon: 2034 to 2037 | Technology readiness: TRL 3 to 7

The second phase transforms research results into an interoperable European chiplet ecosystem. The main challenge is no longer the availability of individual technologies, but the ability of different organizations to develop, exchange, integrate, verify, and qualify compatible components. Europe therefore needs shared governance, common engineering methods, and demonstrator platforms that prove multi-vendor integration under realistic conditions.

A European governance framework should define technical decision-making, compliance, qualification, certification, lifecycle management, and participation rules. Joint European programs should establish several interoperable integration platforms for automotive HPC, AI acceleration, edge computing, industrial control, robotics, and aerospace. The objective is not one universal product, but a portfolio of platforms built on a common technological foundation and shared interface principles.

Genuine mix-and-match integration requires standardized chiplet descriptors, machine-readable metadata, common packaging requirements, verification methods, and system-assembly frameworks. Existing and emerging standards should be extended where necessary to cover package-level and system-level integration. This enables independent chiplet development while ensuring that power, thermal, mechanical, safety, security, and communication assumptions remain compatible at system level.

EDA toolchains must evolve accordingly. Future tools should support cross-company design flows, package-level optimization, multi-vendor verification, reusable integration models, system composition, and marketplace-ready component descriptions. Europe also needs to expand advanced-node design and tape-out competence and establish automotive-grade qualification flows for leading-edge components. These capabilities remain essential even when final products combine advanced-node compute chiplets with mature-node I/O, safety, or mixed-signal components.

Collaboration also requires a trusted intellectual-property framework. Shared interoperability principles, cross-licensing approaches, standard-essential IP rules, and defensive patent mechanisms can protect participants while preserving downstream competition and differentiated products. The ecosystem must support monolithic, hybrid, and fully chiplet-based implementations because market adoption will progress at different speeds across segments.

The expected outcome is a functioning European chiplet ecosystem with established governance, interoperable development environments, multiple demonstrator platforms, repeatable cross-vendor integration, and shared IP structures. These assets provide the bridge from technology exploration to product-oriented industrialization.

3.4.3. Empower & De-risk Industrialization

Target horizon: 2031 to 2034 | Technology readiness: TRL 8 to 9

The third phase focuses on industrial deployment and volume adoption. It translates mature ecosystem results into production-ready semiconductor platforms that can be designed, integrated, qualified, manufactured, and supported by European value chains. Initial products may use a reusable sovereign base die or SoC foundation with open die-to-die interfaces, allowing later extension through AI, compute, memory, safety, or I/O chiplets. This staged approach lowers market-entry risk and supports product-family scaling.

Industrialization must cover the complete platform rather than silicon alone. Priorities include automotive-grade qualification, package and system validation, safety and cybersecurity certification, software drivers and middleware, AI runtimes, board-support packages, thermal and power integration, manufacturing ramp-up, supply assurance, lifecycle management, and long-term software maintenance. The system-integrator role and liability model must be explicitly defined, particularly where components originate from multiple suppliers.

A package-based product family can reuse validated chiplet libraries while varying compute performance, AI acceleration, memory, and interfaces for different vehicles or industries. This allows common base architectures to serve entry, mid-range, and premium products without creating a new monolithic design for every variant. Close vertical collaboration among semiconductor suppliers, packaging providers, Tier 1s, OEMs, and industrial customers is essential to align platform decisions with real deployment schedules and qualification requirements.

Sustainable governance is equally important. Open standards need long-term stewardship, qualification rules must remain stable, and ecosystem operators must coordinate updates, security maintenance, component obsolescence, and supplier transitions over the full vehicle lifetime. Funding and business models should reward both shared platform contributions and differentiated commercial products.

By 2031 to 2034, this phase should deliver production-grade European chiplet platforms, repeatable multi-vendor qualification processes, industrialized packaging and integration capabilities, OEM and industrial deployments, and scalable product families based on reusable chiplet libraries. Together with the first two phases, it creates a continuous European value chain from architecture and silicon through packaging, software, ECU integration, vehicle application, and industrial deployment.

4. Annexes

4.1. ANNEX: RISC-V – Detailed Roadmap Document

The main goals of creating and continuously updating a Roadmap[4], are the following:

- Cover both R&D as well as funding & business aspects in all power & performance ranges from deeply embedded to high-end computing, incl. automotive security & safety requirements.
- Get endorsement of all relevant European automotive industrial stakeholders throughout the value chain and further develop and grow the European RISC-V hardware **ecosystem** from architectural exploration & design up to automotive certification of silicon.
- Document the outcomes of **European projects** on the European RISC-V roadmap and make a gap analysis of missing IPs.
- Align the RISC-V HW roadmap with other initiatives, more in particular **SDV, HPC and chiplet** developments.

The roadmap remains split into 4 key elements or categories of the overall RISC-V based processors:

- Scalable RISC-V Automotive Control Processors
- High-Performance RISC-V based Application Processors
- AI and Compute Accelerators
- System integration and interfacing

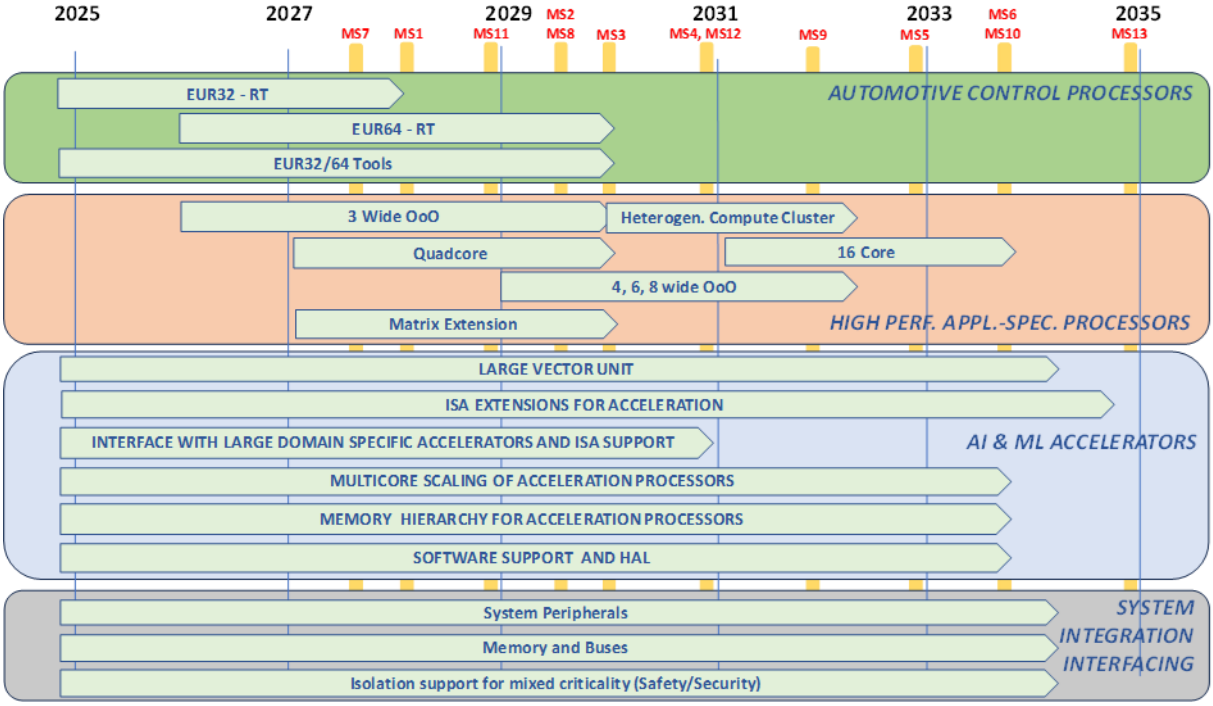
The four elements are linked to each other. The first contains automotive-specific real-time control processors. They are the starting point for the second category, which is more generic and applicable for other type of application domains (communications, industrial, robotics, autonomous systems, health, etc.). The third category contains AI and compute accelerator elements which are specific developments for both automotive as well as other domains, and which contains starting points for next generation high-end processors. As these accelerators are a key component of both processor families, we have opted to make a separate category to show the different IP elements to develop in the coming years. The fourth category contains specific elements for system integration, more in particular how controllers and processors fit into global system solutions.

The four elements can be considered as linked activities as part of a reference platform. By definition, this needs to include the computational elements along with the supporting infrastructure required to realise the whole platform. Therefore, the reference platform infrastructure roadmap laid out includes both SoC and chiplet interconnects.

The goal of deploying the 4 elements together is achieving key goals for automotive safety-critical applications as well as for comparable applications in other domains. Whenever automotive control processors deliver enough performance for a given such application, they are the candidate computing elements since they provide on their own enough features (e.g., fault tolerance support, safety measures, etc.) to allow reaching the highest integrity levels. Some safety-critical applications require higher levels of performance and using high-performance application processors becomes unavoidable. However, those often lack appropriate support to reach the highest integrity levels on their own, and cooperation with the control processors is needed. Hence, the resulting system is more complex, with impact on cost and time-to-market. A third class of applications includes those massively parallel that require extraordinary levels of computing power on one hand, and extreme power efficiency on the other to provide results with tight real-time constraints and within tight power envelopes. Those applications need to resort to accelerators where programmability and maintainability on one hand, and efficiency on the other, need to be traded off. Those accelerators may not provide complete safety support and cooperation with the automotive control processors is often needed to allow realizing the safety concept on the platform. Finally, control and application processors are often multicores capable of running multiple applications potentially with mixed criticalities and heterogeneous performance needs. Hence, appropriate data storage and interface support is needed to allow computing elements execute applications at a sufficient speed to meet their requirements. Hence, caches, interconnects

and memory controllers must come along with appropriate features to allow computing elements execute applications fast enough and support mixed-criticality execution.

Below you can find a summary overview of the Roadmap and the Milestones.



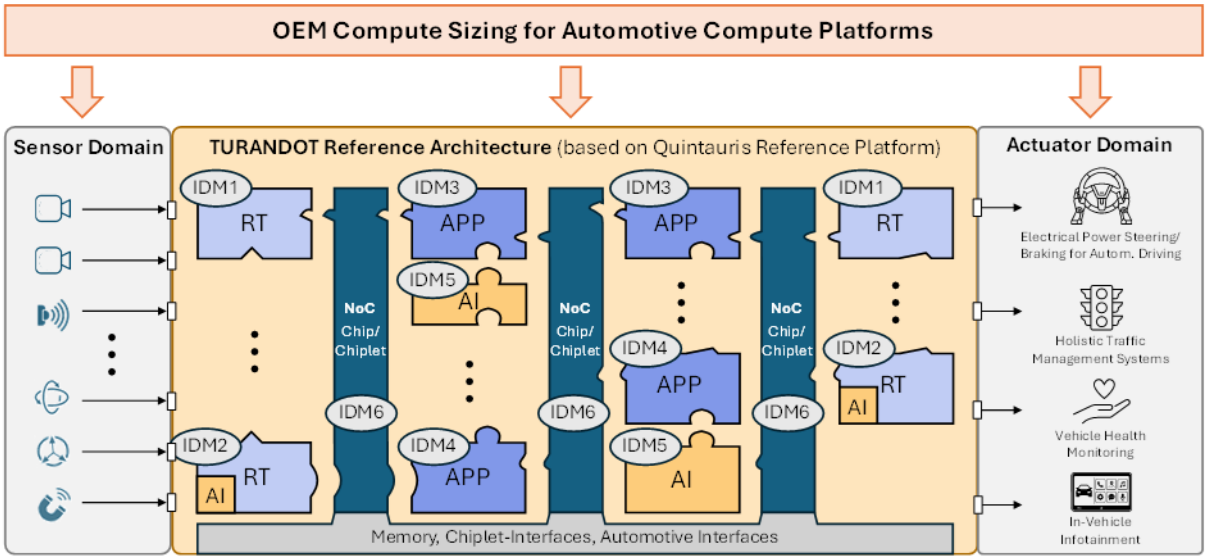


- Q1/2028 - MS1: Full **EURV32-RT processor**, divided in 2 releases MS1_1 and MS1_2
  - Features MS1_1: Contains all key automotive features with limited performance (Q2/2027)
  - Features MS1_2: : Full performance and ISA customization (Q1/2028)
- Q2/2029 - MS2: Full **EURV64-RT processor**, divided in 2 releases MS2_1 and MS2_2
  - Features MS2_1: Contains all key automotive features with limited performance (Q2/2028)
  - Features MS2_2 Full performance and ISA customization (Q2/2029)
- Q1/2028 - MS3_1: **Software** release for **EURV32-RT** controller, containing the following elements:
  - Compilers, binary utilities, libraries, IDE
- Q1/2030 - MS3_2: **Software** release for **EURV64-RT** controller, containing the following elements:
  - Compilers, binary utilities, libraries, IDE
- Q4/2029 - MS4_1: Release of first **out-of-order** processor **EURV64-AP** Application processor
- Q4/2030 - MS4_2: Release of **Quadcore** cluster processor **Quad-EURV64-AP**
- Q3/2032 - MS5: Release of **heterogeneous compute cluster**
- Q3/2033 – MS6: Release of 16-cluster processor **16x-EURV64-AP**
- Q2/2027 - MS7: AI and ML Accelerators: **Octal-core AI engine**
- Q2/2029 - MS8: AI and ML Accelerators: **Dual octal-core AI engine**
- Q4/2031 - MS9: AI and ML accelerators : **Full platform** toolchain integration
- Q3/2033 - MS10: AI and ML accelerators : **Dual octal-core AI chiplet**
- Q4/2028 – MS11: **Decision on IP component configuration** to test complete. SW tests and metrics defined and in alignment with SDV. Definition of Safety and Security measures complete.
- Q4/2030 - MS12: **Comparison** and profile of workloads and systems defined in MS11
- Q4/2034 – MS13: **Integration** and **test of chiplet use-case complete**. Accelerator – Memory – DMA performance verified. Safety and Security measures implemented.

The roadmaps including specific information on the roadmap elements, are explained in the sections below. Before going into this, we have added a specific section of the current implementation status of the Automotive Roadmap, based on the two initial projects TRISTAN & ISOLDE.

The roadmap elements have been defined based on a future generic automotive E/E compute architecture (see picture below), which is representative of the move towards centralized and scalable vehicle compute platforms that support multiple vehicle domains, including in-vehicle infotainment, ADAS, safety monitoring, vehicle control, connectivity and software-defined vehicle functions. From a system perspective, such an architecture requires the ability to size and configure compute resources according to sensor configurations, actuator interfaces, application workloads, data rates, latency constraints, safety requirements and available power envelopes. As can be seen in the picture below, this future automotive E/E compute architecture contains different types of processor families (RT= Real-Time processors, APP=Application Processors & AI=AI-

Accelerators) and system integration requirements, esp. HW/SW integration (SDV) and advanced heterogeneous packaging (a.o. NoC (Network-on_Chip) and Chiplets).

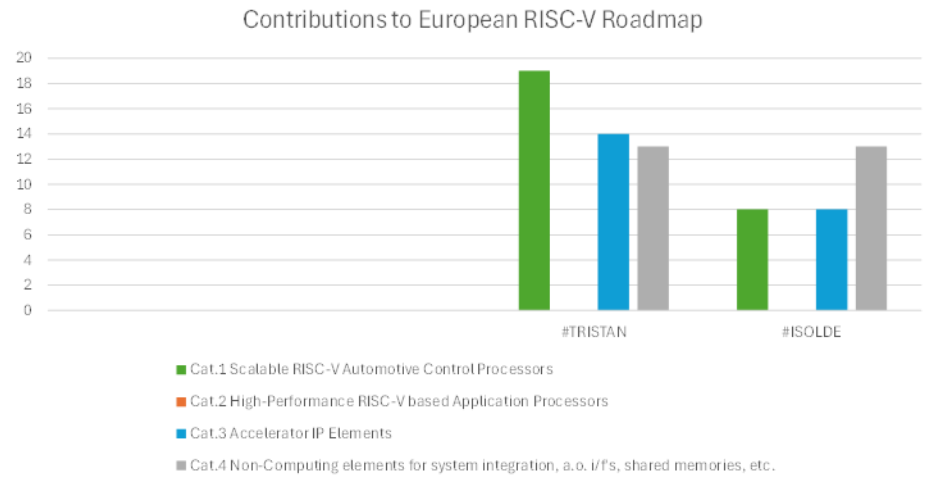


4.1.1. Implementation Status of the Automotive Roadmap

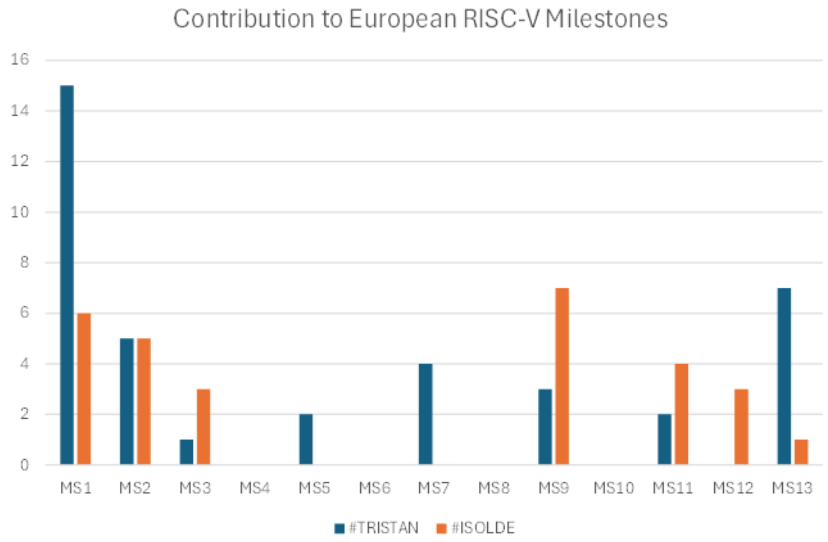
The TRISTAN & ISOLDE projects are the first two projects which have started to contribute IP Building blocks for the Automotive RISC-V Roadmap. The IP blocks generated in this project have been documented and a reference has been made to the Roadmap and to the Milestones of this Roadmap.
TRISTAN is a 54MEuro project with 46 partners and was running from December 2022 to June 2026. ISOLDE is a 42MEuro project with 38 partners and was running from May 2023 to October 2026.
Both projects have made an overview of their IP developments and have documented this in an Excel-sheet, of which you can see an excerpt in the picture below.

Developments in TRISTAN & ISOLDE towards the RISC-V EU Roadmap						
TRISTAN Partner	Brief discription of the development	Type	Status	Implementation	Roadmap Element	Contribution to Milestone
TECHNISCHE UNIVERSITAT DARMSTADT	Scalable&portable core extension interface (SCAIE-V)	*HW-IP*	*2nd Prototype*	*Silicon*	*3.2 ISA extensions for Acceleration*	*MS2: Full EURV64-RT processor, divided in 2 releases MS2_1 and MS2_2*
POLITECNICO DI TORINO	Open-source AI kernel library	*SW-IP*	*Ready for*	*Simulation*	*3.6 Software support and HAL*	*MS9: AI and ML accelerators : Full platform*
POLITECNICO DI TORINO	Tools for full-stack deployment: from high-level AI to deployment	*SW-IP*	*Ready for customer testing*	*Simulation*	*3.6 Software support and HAL*	*MS9: AI and ML accelerators : Full platform toolchain integration*
POLITECNICO DI TORINO	Development of a System-C simulator to co-simulate cycle-accurate routines on TRISTAN HW and system-level I/Os	*SW-IP*	*Ready for customer testing*	*Simulation*	*3.6 Software support and HAL*	*MS9: AI and ML accelerators : Full platform toolchain integration*
ABSINT ANGEWANDTE INFORMATIK GMBH	Hybrid measurement-based timing analysis for RISC-V-based systems	*SW-IP*	*Ready for customer testing*	*other*	*1.3 Tools*	*MS3_1: Software release for EURV32-RT controller*
semify e.U.	Accelerator implementing ISA extensions for signal compression	*HW-IP*	*2nd Prototype*	*FPGA-Implementation*	*3.2 ISA extensions for Acceleration*	
AICAS GMBH	Securly supervise and manage RISC-V systems from cloud	other	2nd Prototype	other	4.3 Isolation support for mixed criticality (Safety / Security)	MS13: Integration and test of chiplet use-case complete.
AICAS GMBH	Securly supervise and manage RISC-V systems from cloud	other	2nd Prototype	other	4.3 Isolation support for mixed criticality (Safety / Security)	MS13: Integration and test of chiplet use-case complete.
ACCEMIC TECHNOLOGIES GMBH	Trace IP with extended support for safety-critical systems	*HW-IP*	*1st Prototype*	FPGA-Implementation	4.1 System Pheripherals	MS11: Decision on IP component configuration to test complete..
ACCEMIC TECHNOLOGIES GMBH	Implementation of runtime observation and analysis tool chain	*SW-IP*	*1st Prototype*	other	1.3 Tools	MS11: Decision on IP component configuration to test complete..
Thales	CVA6 open-source RISC-V 32/64 bit application core	Processor Core	2nd Prototype	other	1.1 EURV32	MS1: Full EURV32-RT processor, divided in 2 releases MS1_1 and MS1_2
Thales	CVA6 open-source RISC-V 32/64 bit	Processor	other	other	1.2 EURV64	MS2: Full EURV64-RT processor, divided in 2
ROBERT BOSCH GMBH	Instruction set extensions for automotive applications	HW-IP	other	FPGA-Implementation	1.1 EURV32	MS1: Full EURV32-RT processor, divided in 2 releases MS1_1 and MS1_2
ROBERT BOSCH GMBH	Accelerator building blocks for automotive applications	HW-IP	other	FPGA-Implementation	1.1 EURV32	MS1: Full EURV32-RT processor, divided in 2 releases MS1_1 and MS1_2

Based on this overview, we have generated pictures with the number of elements contributing to the different elements of the Roadmap and to the Milestones, as can be seen in the figures below. In the future, this will also be done with the other RISC-V related projects in the Chips-JU program.



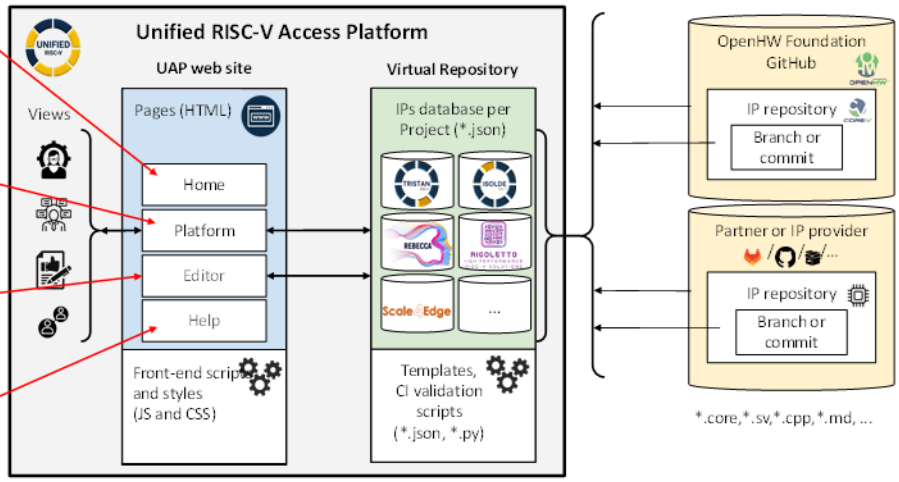
As can be seen in the picture above, the key focus of TRISTAN is on the 1st category of scalable real-time automotive control processors. In ISOLDE, the key focus is on non-computing elements, peripherals and system integration elements. No activity is yet planned on the high-performance RISC-V based application processors, as these are longer term oriented and will be planned in the new and upcoming RISC-V projects.



In terms of milestones, the focus in both projects has been on MS1, MS2 (real-time control processors), MS9 (AI accelerators) and MS13 (system integration).

Important is also the fact that all IP is available via a Unified RISC-V IP Access Platform, which is managed by the Eclipse Foundation / RISC-V Open HW Working Group, as shown in the picture below.

- **Welcome page** – explains the role of the UAP as reference for RISC-V IPs developed under EU projects
- **Platform page** – dynamic table with the latest IP information under the Virtual Repository.
- **Editor page** – to edit or update IP and project entries.
- **Help page** – guidance on how to contribute to the UAP or its Virtual Repository catalogue of IPs, *currently under construction*.



4.1.2. Scalable RISC-V Automotive Control Processors

Global requirements for automotive real-time RISC-V microprocessors

The proposed program and roadmap addresses the development of a family of real-time RISC-V microprocessors to be used mainly in the automotive market with certifications in the safety and security fields. Automotive platforms remain the primary focus. We anticipate an increase in computing capabilities driven by new algorithms using artificial intelligence techniques, with workloads partitioned between CPUs and AI accelerators. These same platforms can also be deployed in emerging markets, often classified under "Physical AI", which largely share requirements similar to those of the automotive sector.
The following requirements are mandatory:

- RISC-V ISA compliance
  - With extensions to increase performance and features in automotive applications (to be proposed to the RISC-V International Association for ratification)
- Competitive performance and special features compared to future expected automotive microprocessor solutions on key benchmarks:
  - Instruction and data level parallelism (scalar and vector instructions)
  - Tuned micro-architecture for high speed in the set of the chosen technologies
  - Fast response from asynchronous events and fast context switch with deterministic computing decisions (microprocessor and associated memory hierarchy)
  - ISO26262 safety compliance
  - Hardware (HW) support for virtualization enabling isolated compartments
  - Better security features to avoid malicious access and interference (side channel attacks, etc.), and compliance
  - More energy-efficient solutions
  - Cost effective solutions (performance per unit area)

Next generation features of the RISC-V automotive real-time core family

In the evolution towards scalable RISC-V high-performance real-time processors, following new additional features are needed:

- Next generation Edge Computing System for Automotive and Physical AI applications
  - A new class of high performance real-time microprocessors for:
    i. Automotive realtime tasks previously defined in the Turandot project requirements

ii. New class of Physical AI algorithms, from the pre-processing of the data extracted from sensors to the AI processing alongside with Generative AI units

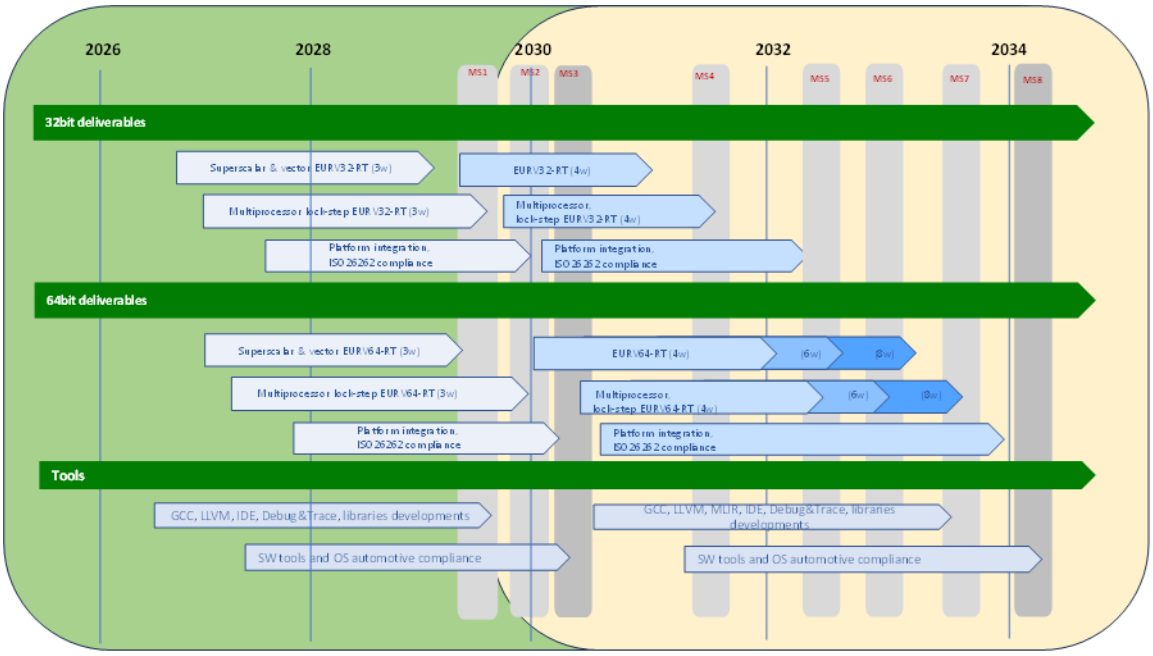
- New memory hierarchy structures to support low latency with the big data to be elaborated, both in instructions and data domain, with specific requirement for multi-processor clusters
- Matrix interconnect IP to enable high thruput and low latency data transfer among microprocessor, AI accelerators and memories
- Augmented SW tools to manage processing allocation between CPU and accelerators

- **Scalable High Performance Real-Time Family: 32/64bit superscalar OutofOrder, 4-8 issues, with integrated vector/matrix unit**
  - Domain Specific Architecture for high KPI in terms of performance over cost
  - Simplified architecture in respect to the Application processors (no requirement to support high-end OS)
  - Ability to make instruction customization
  - SW ecosystem under supplier/customer control

Scalable Automotive RISC-V Roadmap

The roadmap covers three main parallel activities:
- The development of the 32bit core family
- The development of the 64bit core family
- The design of the software (SW) tools

For each family we target two microprocessor releases: the first with the key features, trying to make a fast delivery, and the second with new advanced features (a.o. specific customizations, finetuning of the micro-architecture) and working in parallel, producing a step ahead among the competitors.
Thereby it is important to note that the released versions are initial prototypes for initial usage, validation and testing. These are not yet commercial products on the market. For full commercial products, more effort is needed in terms of validation, full compliance, full software integration and release. Multiple silicon runs are typically needed. On average, it still takes a few years before a full commercial release, taking into account the complexity of the processors and the advanced technology used.



4.1.3. High-Performance RISC-V based Application Processors

The roadmap of High-Performance Application Processors has been derived from the basic automotive profile processors, which contain an embedded hypervisor, multi-threading and embedded vector. So the starting point is the same as what has been outlined in the section before.
Overall, we see a clear need for industrial-grade and automotive-ready application processors because of the following reasons:

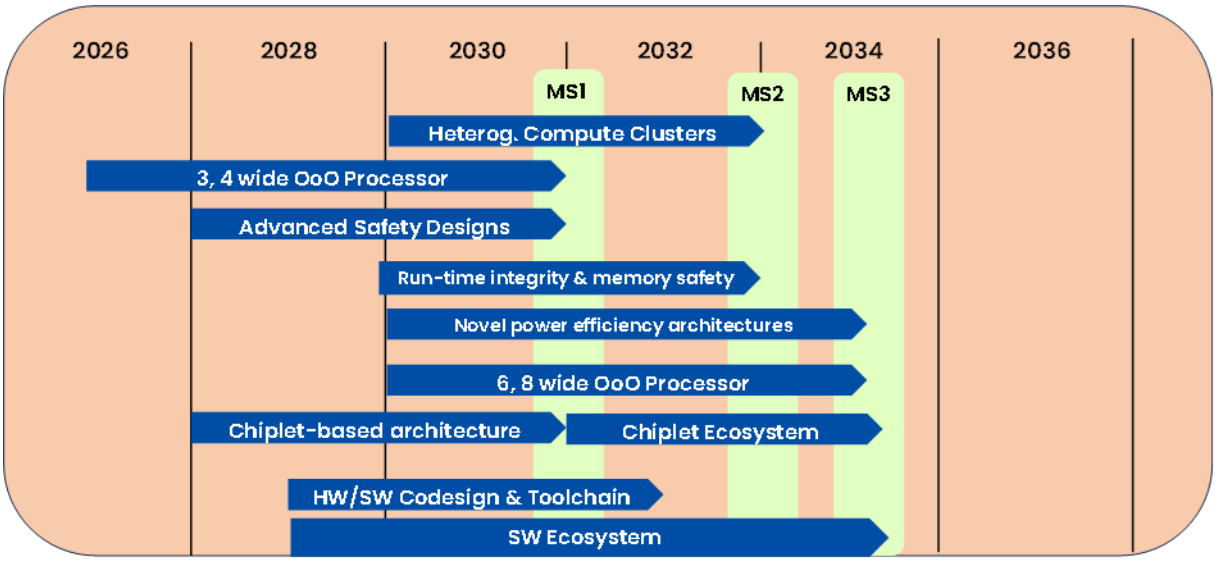
- Run-time integrity & memory safety: protect embedded software from unauthorized modification, unexpected crashes, and malicious exploits like buffer overflows, ensuring system stability and security;
- Advanced safety designs: (1) Efficient solutions for integration of such cores in safety-critical systems required à HW-SW Co-Engineering; (2) Tradeoffs in terms of time-bound execution in hybrid application scenarios à Hybrid architectures MCU/MPU: trade-off between RT and performance
- Novel power efficiency concept/architectures required to address thermal challenges
- Chiplet-based architectures supporting efficient integration of automotive compute elements on system-level, e.g. application-specific accelerators
- Enabling “fast track” for 4,6,8 wide OoO cores: not necessarily going through a 3-wide OoO step

Some elements of the current AP-roadmap seem essential for all high-performance processor types and are not Application-Specific (e.g. high-bandwidth interconnects & memory controllers, shared cache memories; matrix extensions..).

Key elements of the Application-Specific processors are:

- Functional safety & security
    - Mixed-criticality system architecture supporting safe and non-safety domains on a shared platform
    - ASIL-B/-D system alignment with integrated safety-island architecture (heartbeat, fault containment, fail-operational)
    - Lockstep / dual-execution and ECC across compute and memory paths, with RERI-compliant fault reporting
    - Cybersecurity compliance per ISO 21434 (HW root of trust, secure boot)
    - Speculative vulnerabilities mitigation (Spectre, Meltdown variants)
    - Memory tagging/ virtual memory protection
    - Appropriate testing (fuzzing)
- SW Ecosystem & Standardization
    - Compliance with ratified mainstream RISC-V ISAs to prevent ecosystem fragmentation
    - Active engagement with and compliance to emerging automotive SW and HW standards
    - Automotive RISC-V application-class profile definition and alignment (e.g., Quintauris)
    - EU-specific encryption standards addendum (23.1 profile)

The High-Performance RISC-V based Application Processors roadmap is shown below:



4.1.4. AI and Compute Accelerators

RISC-V processors will play a crucial role in AI, (Machine Learning) ML and data-intensive application acceleration due to the advantages offered by ISA extensions and the flexibility inherent in an open ISA. This is particularly important in AI and ML, where workloads evolve at a fast pace as they are tuned key application domains such as advanced Automotive, Industrial automation. Designers can tailor RISC-V processors to specific application and domain requirements, enabling optimization for diverse compute-intensive tasks, from deep learning to computer vision. RISC-V makes it possible to integrate specialized instructions tailored for AI, ML and other data-intensive tasks directly into the processor. This facilitates the development of dedicated in-core and near-core hardware accelerators for specific operations commonly found in these workloads, such as matrix multiplication or tensor operations. The following platform elements compose the AI and Compute accelerators envisaged in the roadmap:

1. **Vector ISA (RVV) specification**. Large vector machines usually run into the problem of scalability since you can not efficiently provide an all-to-all lane swap. Hence, the design considerations for a processor with a large vector unit are multifaceted. First and foremost, research is needed to tune processor architecture for very large vectors, surpassing 1024 bits. This requires attention to the register file and lane organization, optimizing them for scalability. Achieving a scalable interface memory hierarchy is also imperative, as it supports the efficient functioning of these expansive vector units. The evolution of the ISA is pivotal for scalability, addressing issues such as those arising from slide instructions.

2. **ISA extensions beyond basic vector instructions**. Matrix operations of all sizes are prevalent in neural networks. Matrix ISA extensions, exemplified by rank-1 updates with outer product operations, can contribute significantly to the versatility of the processor. A matrix ISA extension should be carefully defined to be able to leverage the data structures and interface memory hierarchy optimized for a large vector unit, as the vector and matrix extensions are likely to be implemented jointly. At the same time, such an ISA extension should be kept as lightweight as possible to allow efficient and scalable implementations. A tight integration of vector processing and matrix processing elements, with potentially shared structures and interfaces, is essential for energy efficiency and scalability. A well-defined set of matrix ISA extensions can be leveraged into the necessary software infrastructure work (frameworks and libraries) supporting reuse between hardware architectures. Extensions focusing on highly quantized low-bitwidth number representations, spanning 8-bit and below, as well as block floating-point support in both vector and scalar contexts, are critical for energy efficiency.

Support for mixed precision and new data formats, whether through status-based mechanisms or a broader ISA encoding space, is a key direction for enhancing adaptability. In the longer term, introducing stream semantic ISA support offers an alternative or complementary approach to vector methodologies, providing flexibility in handling data streams.

3. Ensuring the processor's ability to **interface seamlessly with large domain-specific accelerators** necessitates a design space exploration effort, covering tightly coupled versus loosely coupled architectures. Control plane issues, encompassing exception handling and context switching, must be addressed to ensure efficient operation of acceleration engines.

4. Ultimately, efficient **multicore scaling of acceleration processors** is paramount, with clustered vector processors and small versus large vector cores playing pivotal roles in defining the profile of high-performance and scalable processor. In addition, a mixture of multiple enabling technologies (e.g., RVA, compute-in-memory, custom compute) is required. Memory hierarchy with multiple interconnectivities between accelerator cores and accelerator-side QoS handling will be needed too. As automotive processors will be required to deliver an order of magnitude jump in performance operating with high resolution images in order to support ADAS level 5 functionality, system energy efficiency is a key design consideration. Hence, the ability of domain-specific accelerators to deliver both very high compute performance as well as high levels of energy efficiency will be a key success factor for this reference platform.

5. A well-designed **memory hierarchy** is critical for AI and Compute acceleration, as most applications are bound by the available memory bandwidth. By providing a more sophisticated memory hierachy that exploits temporal re-use of data items the memory bandwidth limit can be mitigated and more applications made compute bound. Memory hierarchy design involves choices between shared L1 and private L1 caches, as well as the interplay between cache and Tightly Coupled Data Memory (TCDM). High-bandwidth memory hierarchy considerations, including L2 caches and computation in the memory hierarchy, further contribute to the overall efficiency of the system. Communication assistance for AI processing, such as DMA tensor engines and prefetchers, is also a key research area is it enhances data flow and processing speed.

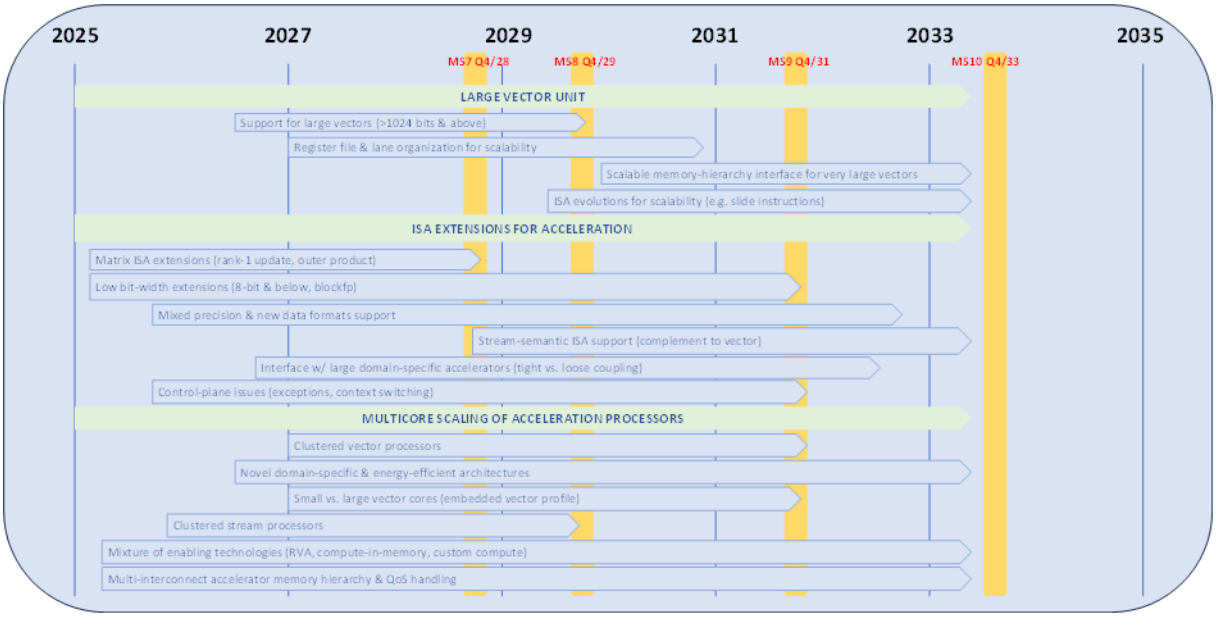
6. As scaling the vector architecure is limited by the above mentioned lane swap effects as well as the fact that designs need partitioned and grouped for physical implementation, scaling over the amount of cores is a straightforward (from a hardware perspective) way of achieving more compute power over a given amount of silicon area. The disadvantage of scaling via the number of cores becomes apparent when we look at sharing data between cores: As soon as cores have private copies of data items in their caches that aren't up-to-date in main memory, it becomes a necessity that copies are exchanged between caches. Hence, **coherence and consistency models are essential to ensure the correctness of data-parallel computations**.

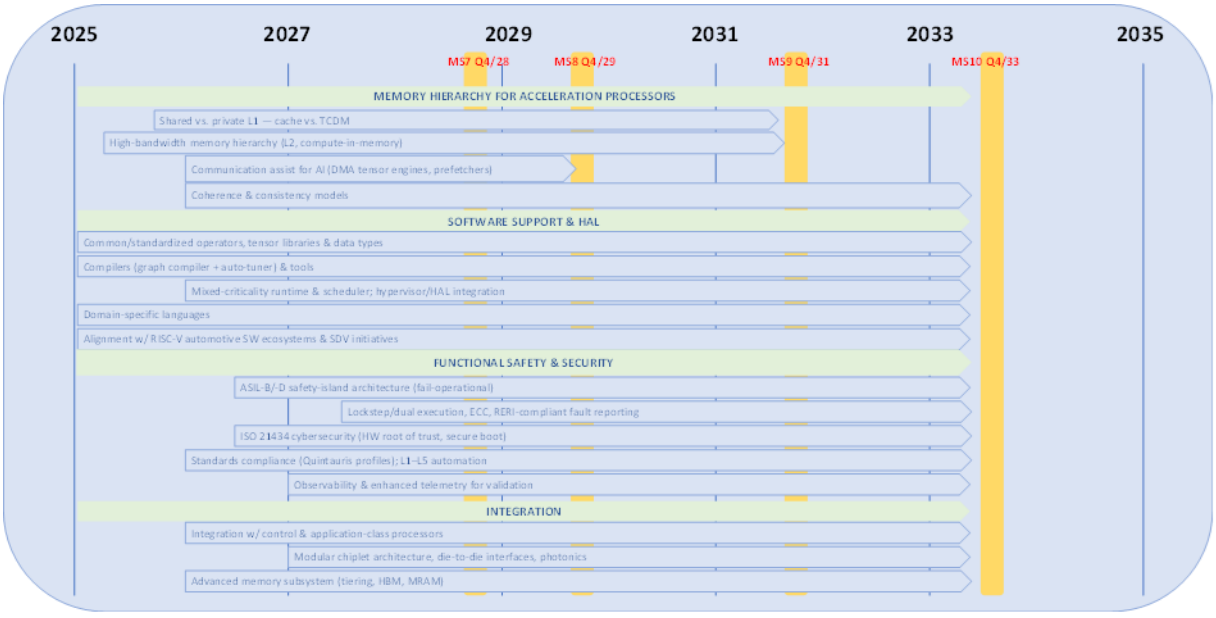
7. **Strong software support and Hardware Abstraction Layer (HAL) implementations** are necessary for seamless integration into existing application ecosystems. The availability of low-level libraries for key primitives, such as tensor libraries, empowers developers to harness the full potential of the processor. This should include the definition of a common standardized operators, libraries and data types. Also compilers need to be included (Graph compiler + auto-tuner). A comprehensive software stack and support for domain-specific languages round out the requirements for a sophisticated and efficient acceleration processor, including mixed-criticality-aware runtime & scheduler, hypervisor and HAL integration. An important must-do is the alignment and contributions to the RISC-V SW Automotive ecosystems, especially the alignment with automotive SDVs.

8. **Functional Safety & Security** are key requirements for automotive applications.. This should contain the following features:

- ASIL-B/-D system alignment with integrated safety-island architecture (heartbeat, fault containment, fail-operational)
- Lockstep / dual-execution, ECC across compute and memory paths, and RERI-compliant fault reporting
- Cybersecurity compliance per ISO 21434 (HW root of trust, secure boot)
- Active engagement with and certification compliance to emerging automotive SW and HW standards (e.g. Quintauris profiles)
- Observability capabilities and enhanced telemetry for validation
- Compliance w/ the needs of L1 – L5 automation

The AI and Compute Accelerators roadmap is shown below:





4.1.5. System Integration and Interfacing

The integration of different IP blocks into a product e.g., an SoC or a Chiplet for an intended use case, requires careful consideration of the interaction between IP blocks in order to build robust and differentiated products. Early pre-silicon integration of the relevant IP components together with an appropriate software configuration and test application creates confidence and helps to minimize the risk of late failures for the partners collaborating in the supply chain.
The system integration task includes the generation of new IP and the assessment of existing IP blocks. Particularly, for newly designed IP blocks or new combinations of different IP blocks, an early verification of the overall system functionality is essential. The integration of open-source IP requires additional efforts in terms of verification and qualification. While the source code is open and available to all involved parties, the integrator must understand all implications and risks involved in the design process.

As the components mentioned in this roadmap only function in a system context, there is a close relationship to the other roadmap elements, especially the CPU cores. Therefore, the peripheral integration will start with a default system that must be capable of integrating new components upon availability. During project definition, integration platforms for the intended core classes and integrated peripheral components must be aligned with verification and SW bring-up tasks for providing a meaningful system assessment.

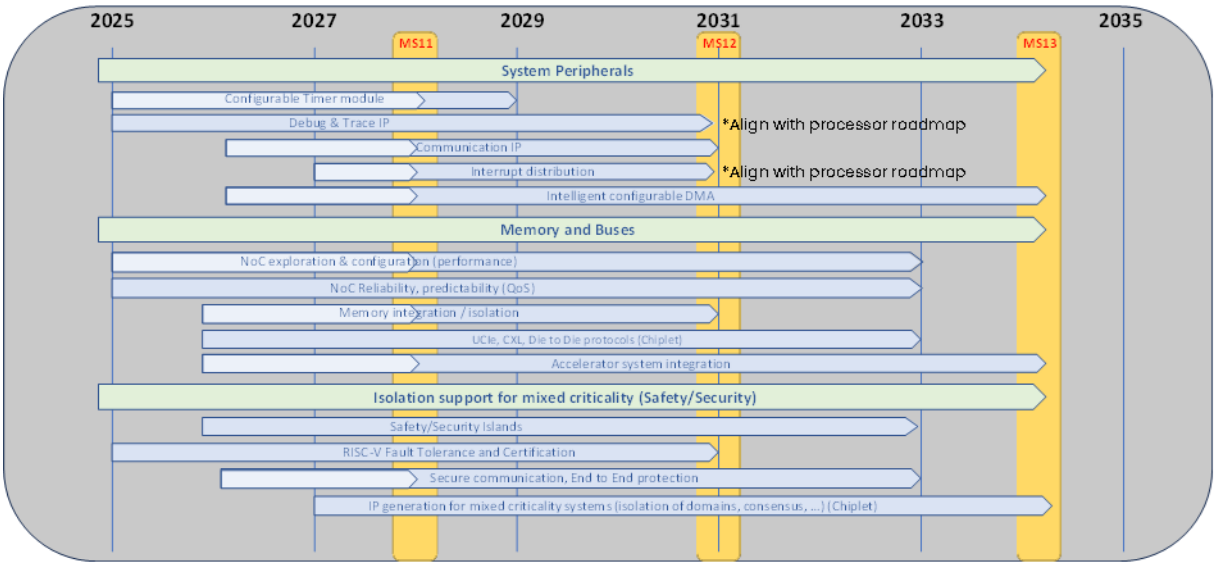
The following section briefly explains the elements of the System Integration Roadmap, with a particular focus on the components and interfaces needed to support the updated processor roadmap.

- **System Peripherals:** System peripherals are a subset of peripherals that are tightly coupled with the overall system function, such as DMA engines, interrupt controllers, debug and trace infrastructure, and other components that directly influence system observability, responsiveness and data movement. In contrast to more loosely coupled peripherals, which typically perform isolated functions, system peripherals must be orchestrated together with processors, accelerators, memories and interconnects. This orchestration is critical to achieve predictable performance and to support safety, security and validation requirements in complex automotive platforms. **Trace and debug IP** will need to become significantly more capable for future European RISC-V processors and heterogeneous SoC platforms. The increasing number of processing elements, accelerators and clock domains, together with higher overall processing performance, will lead to substantially higher trace data rates and more complex debug scenarios. Therefore, trace and debug IP must scale in terms of throughput, configurability and integration depth, while remaining suitable for automotive validation, safety analysis and in-field diagnostics. The roadmap must address updated trace and debug requirements resulting from new European processor IP, including support for multiple processing components, coherent and/or synchronized execution across processing components, multiple clock domains, higher processor throughput and more complex heterogeneous execution scenarios. Compatibility with emerging and ratified trace and debug standards is essential. RHTI is the ratified standard for the processor trace interface, but it is not yet implemented by all processors. Broader processor support for RHTI is therefore an important roadmap topic. In addition, two complementary standards exist for the data stream after the trace encoder: N-Trace and E-Trace. Further alignment is needed because N-Trace is still incomplete without full data trace support. For debug access, an efficient two-wire debug standard is needed that fulfils automotive robustness, reliability and throughput requirements. Such an interface should support scalable bring-up, production test, failure analysis and service diagnostics without creating excessive pin-count or integration overhead. **Interrupt distribution solutions** are tightly coupled to the processor architecture and are expected to be influenced by the RISC-V Advanced Interrupt Architecture. Future platforms must support scalable interrupt routing, prioritization and virtualization across multiple cores, accelerators and safety islands. The interrupt infrastructure must therefore provide deterministic behavior, low latency and sufficient flexibility to support mixed-criticality workloads and heterogeneous compute configurations. Future automotive RISC-V platforms also require **intelligent and configurable DMA solutions** that can serve as complex alternatives to established DMA architectures. Such DMA engines must support efficient data movement between processors, accelerators, memories and peripherals, while providing configurability, isolation and predictable behavior. Compatibility with coherent interconnect protocols, including CHI-based environments, is important to enable efficient integration into heterogeneous and chiplet-based systems.

- **Memory and buses:** Memory and bus architectures are central elements of the overall system and have a significant impact on both performance and functionality. Multicore systems, network-on-chip architectures and chiplet interfaces introduce new challenges for achieving predictable timing behavior across different system scenarios while preserving high performance. In addition, the integration of accelerators and the resulting memory transfers can strongly affect system performance and must therefore be considered early in the architecture design. Scalable methods for **architecture exploration, NoC configuration and processor performance assessment** are required, particularly for new high-performance and multicore processor developments. **Reliability and predictability** are equally important, as increasingly distributed and complex automotive systems require end-to-end Quality-of-Service mechanisms, protocol interoperability across CHI.E, CHI.G and AXI, and robust cache coherency concepts. **Memory integration and isolation** must also address advanced automotive requirements, including cache and local memory safety mechanisms, address error detection, end-to-end safety, and appropriate MMU and MPU support.

- **Isolation support for mixed criticality.** Ensuring safe and secure operation of a platform must be considered as early as possible in the design process on system level. Furthermore, the integration of applications with different levels of criticality poses additional challenges for ensuring the correct operation and prioritization. The impact of these measures for safety and security on the overall system is critical for ensuring correct system behavior. RISC-V with its customizable ISA provides new possibilities for efficiently supporting these features. In a platform with multiple Computing Islands, it makes sense to centralize safety and/or security actions inside a dedicated RISC-V microprocessor unit.
  For such **safety and security islands**, requirements are increasing due to the growing complexity of SoCs, and the hardware-based security extension CHERI will require adaptations to the architectures of safety and security islands. The envisioned RISC-V controller cores used in safety and security islands belong to a class of microprocessors as defined in the TRISTAN and ISOLDE projects with features such as side channel protection, vector crypto or accelerator, hash unit, random true generation, PUF or SESIP compliance.
  With respect to **RISC-V fault tolerance and qualification**, further input from the Functional Safety SIG and the Automotive SIG will be embedded in future integrated system developments. A key topic is how RISC-V extensions can be leveraged for performance, safety, security and domain-specific functionality while preserving compliance with automotive safety requirements.

The System Integration and Interfacing roadmap is shown below:

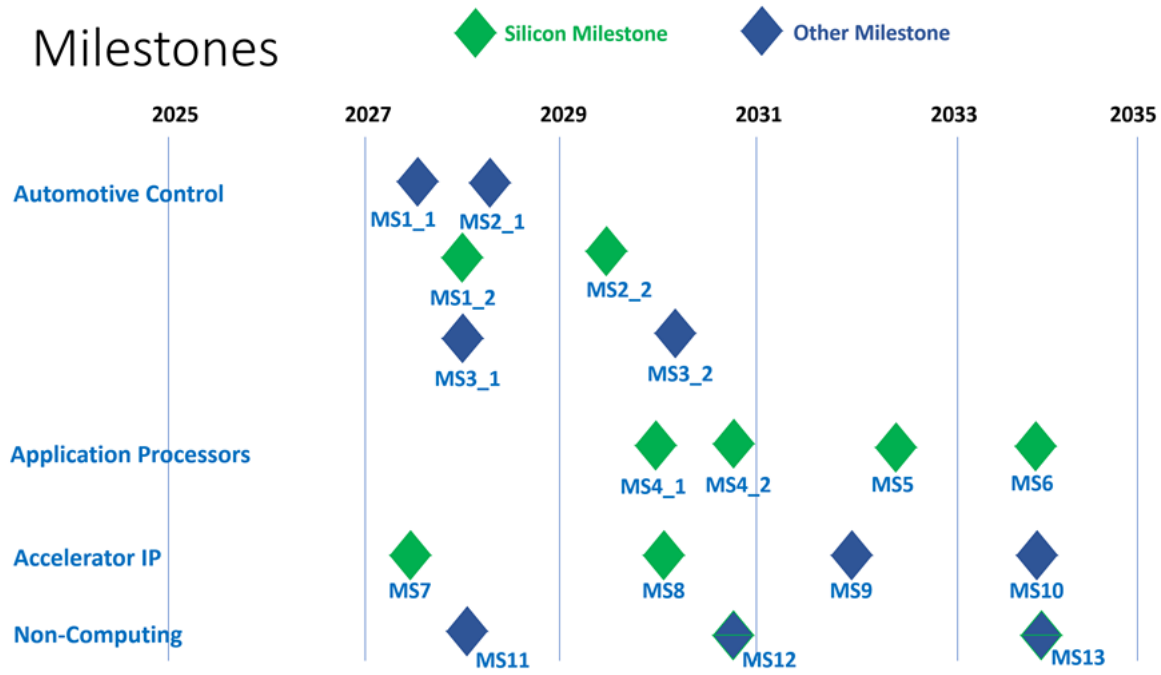


4.1.6. Milestones

The milestones typically consist of an integration of multiple deliverables from the different categories of elements of the roadmaps, which are used to build a demonstrator that can be used by first customers to further refine and integrate the results into their final products.

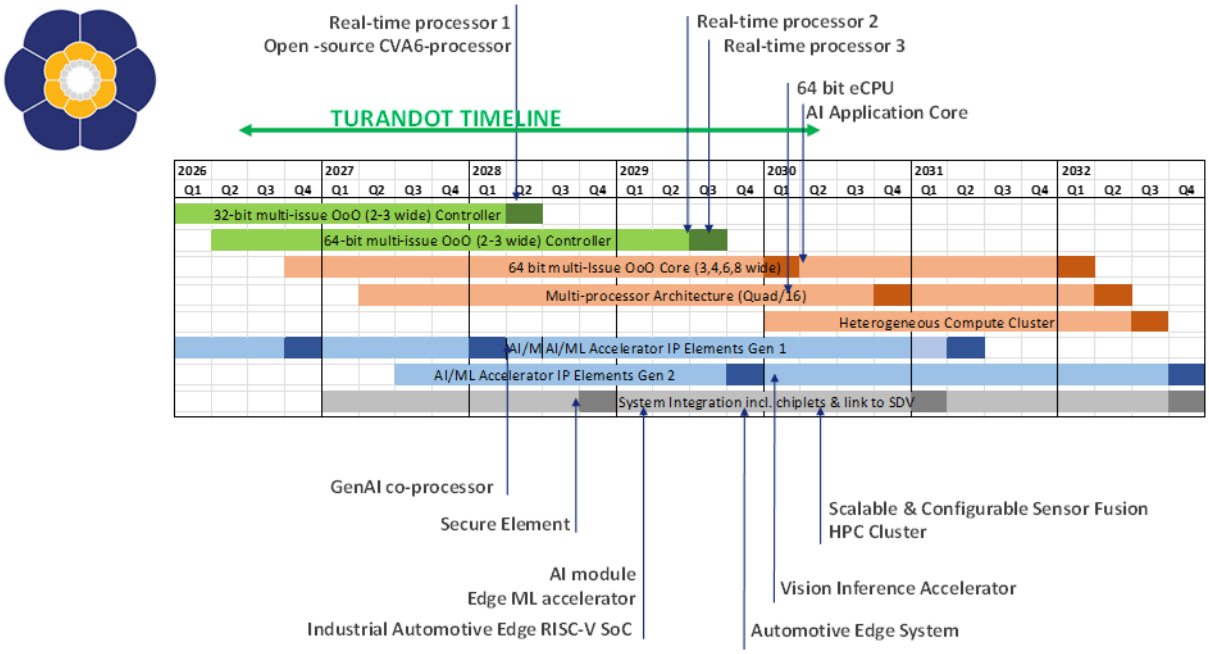
List of Milestones:

Q1/2028 - MS1: Full **EURV32-RT processor**, divided in 2 releases MS1_1 and MS1_2
Features MS1_1: Contains all key automotive features with limited performance (Q2/2027)
Features MS1_2: : Full performance and ISA customization (Q1/2028)
Q2/2029 - MS2: Full **EURV64-RT processo**r, divided in 2 releases MS2_1 and MS2_2
Features MS2_1: Contains all key automotive features with limited performance (Q2/2028)
Features MS2_2 Full performance and ISA customization (Q2/2029)
Q1/2028 - MS3_1: **Software** release for **EURV32-RT** controller, containing the following elements: Compilers, binary utilities, libraries, IDE
Q1/2030 - MS3_2: **Software** release for **EURV64-RT** controller, containing the following elements: Compilers, binary utilities, libraries, IDE
Q1/2031 MS4_1: First industrial-grade 3-4 wide OoO processor architecture available
Q1/2031 MS4_2: First chiplet-based architectures in automotive RISC-V with embedded AI&toolchain
Q1/2033 MS5: tbd
Q1/2034 MS6: European modular composable Chiplet-architecture for high-end SoCs and Edge Computing with embedded AI
Q2/2027 - MS7: AI and Compute Accelerators: **Octal-core AI engine**
Q2/2029 - MS8: AI and Compute Accelerators: **Dual octal-core AI engine**
Q4/2031 - MS9: AI and Compute accelerators : **Full platform** toolchain integration
Q3/2033 - MS10: AI and Compute accelerators : **Dual octal-core AI chiplet**
Q4/2028 – MS11: **Decision on IP component configuration** to test complete. SW tests and metrics defined and in alignment with SDV. Definition of Safety and Security measures complete.
Q4/2030 - MS12: **Comparison** and profile of workloads and systems defined in MS11
Q4/2034 – MS13: **Integration** and **test of chiplet use-case complete**. Accelerator – Memory – DMA performance verified. Safety and Security measures implemented.



TURANDOT and the RISC-V Automotive Roadmap

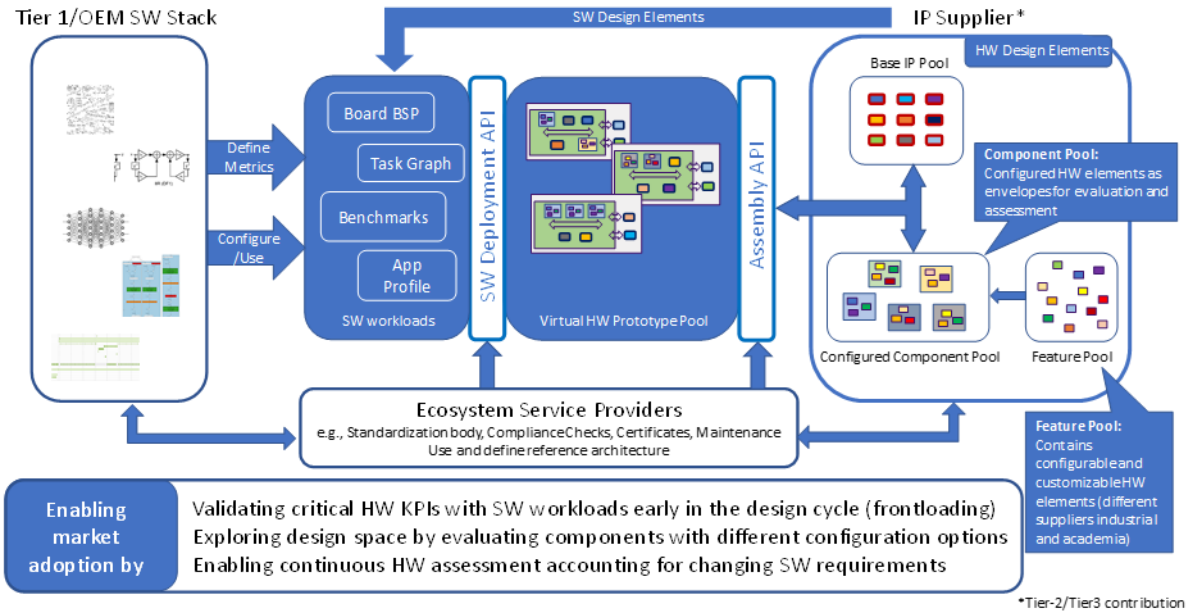
The TURANDOT project is a 285 MEuro project with 43 partners from 13 countries in Europe, running from 1 August 2026 till 31 July 2030. The main goal is to implement largely the roadmap as described in this document by building real silicon to be tested by TIER1's and OEM's for their future E/E architecture. This is the current view on how TURANDOT will contribute to the roadmap:



HW/SW Integration

The HW-developments will be carried out in cooperation with the future requirements on Software Defined Vehicles. These requirements are typically on a much higher level of abstraction than needed for deriving HW KPIs and the process of mapping these high level requirements to predict HW performance or select HW IP are manifold. The input from the SW side to this process can range from high-level SW workloads to compiled binaries for dedicated HW architecture. It is important to note that the process of providing an optimized implementation for a dedictated HW archicture is typically a lengthy and complex process. Having a defined set of APIs and Digital Subscriber Lines (DSLs) for assembling, building, running and analyzing the respective artefacts on SW and HW side is essential for enabling continuous SW/HW assessments. While some companies have very successfully established this flow by a dedicated SW stack with a dedicated HW-IP, the challenge is to promote this flow in a distributed automotive supply chain setup. In addition to that, concepts for decoupling SW from dedicated HW components like middleware or special runtimes must be considered by Tier-1 and Tier-2 suppliers.

As a starting point for future interaction, we propose the following high-level system architecture for interaction in the automotive supply chain. This starting point will create and re-use open APIs, DSLs and state-of the-art SW methods in order to establish a scalable and continuous architecture assessment.



The advantage of such a distributed HW/SW Co-Design setup is that the different suppliers in the supply chain can focus on their core competencies while remaining competitive addressing a larger market.
The OEM/SW stack provides the basic requirements and workloads needed to select a certain IP. These requirements need to be defined or transformed into an executable profile which can then be deployed different realizations of HW prototypes. It is important that this HW envelope is generic enough to embrace multiple HW realizations to avoid manual interactions when re-mapping the SW components. The metrics from the SW side can be seen as a contract between the SW, which is continuously evolving while development advances, and the HW which also adds features or adapts its implementation. Therefore, we see the collaboration with the Software Defined Vehicle as a key part to derive this interface and ensure full consistency with the SW requirements. A core aspect of this cooperation needs to define the underlying data formats that facilitate the proposed scalability and automation.
On the IP supplier side, there are different HW design elements which can be assembled according to a defined system description. The pre-configured HW envelopes contain configured IPs for a certain use-case e.g. DMA, IRQ controllers, busses, memories where the configuration can automatically be changed for different HW assessments. This setup is complemented by the Base IP pool which contains standard peripherals e.g. for communication or simple traffic generators.
The ability to orchestrate the HW for assessing different configurations is essential for an efficient mapping of the SDV requirements.

Conclusions and Next Steps

This roadmap and current and future updates can be considered as a framework for future RISC-V based developments in Europe with initial focus on the high-end automotive market. A main focus area thereby is the domain of high-performance cores in real-time applications, that are usually in-order cores with deep pipelines (8-10 integer stages) and low latency response times on critical events like interrupts, exceptions and context switches. The focus of the roadmap is on the hardware developments, but some initial ideas on software, libraries and standardized interfaces are also included, but need to be further worked out. This will be done in close cooperation with the different initiatives and projects in Europe on a SW Defined Vehicle (SDV). In addition, close cooperation is envisioned with European initiatives and working groups on chiplets, as a new packaging technology which can serve as a solution of closely networked processors in one package instead of monolithic single high-performance solutions. An important contribution will also be on defining and performing benchmarking to assess the platform throughout all the phases of its development. This implies being able to run them in environments with varying degrees of maturity, ranging from simulation environments to FPGA prototypes, test chips, and finally, fully-fledged silicon, because the platform is targeted for commercial use and should comply with industry standards with respect to quality.

The Automotive RISC-V reference platform is geared towards maximizing the benefits of RISC-V through healthy competition, while streamlining the ecosystem to facilitate hassle-free integration of RISC-V cores and SoC into automotive systems. Competition should happen on the levels of HW implementation and reference platform compatible innovative vendor extensions.
Finally, this is expected to be a continuous activity. Further iterations and refinements will be needed. Since the creation of the initial roadmap, elements of this roadmap has been developed in projects, as e.g. TRISTAN and ISOLDE[7]. Also a new company has been set up, Quintauris, with shareholders NXP, Infineon, Bosch, STMicroelectronics, Qualcomm and Nordict, in order to work on an automotive reference platform and certification of RISC-V based solutions for automotive applications. The intention is to further follow up developments in Quintauris and in other European projects, as a.o. RIGOLETTO and TURANDOT to further implement the roadmap as described in this document. Finally, an important next step will be to get further and more detailed feedback and buy-in on the proposed roadmaps by OEMs, TIER1s and other stakeholders in order to continuously monitor and align on priorities and new pre-competitive automotive processor research topics in Europe.

[1] Recommendations and roadmap for European sovereignty on open source hardware, software and RISC-V Technologies | Shaping Europe's digital future (europa.eu)
[2] TRISTAN – Expand, mature and industrialize the European RISC-V ecosystem (tristan-project.eu)
[3] ISOLDE | (isolde-project.eu)
[4] Roadmap_RISC-V_v240216_Final.pdf
[5] RIGOLETTO Project | Automotive RISC-V Hardware Platform for Software-Defined Vehicles - Rigoletto-Project
[6] Home | Rebecca
[7] Both projects are running under the KDT-Framework with funding from the European Commission and different Member States. For more details, see: TRISTAN – Expand, mature and industrialize the European RISC-V ecosystem (tristan-project.eu) and ISOLDE | (isolde-project.eu).

4.2. ANNEX: Chiplet Integration – Detailed Gap Analysis

The current workstream material already describes the three roadmap categories as: Build the Foundation, Prototypes for next generation, and Empower & De-risk industrialization. It also explicitly notes that these phases run in parallel from a SRIDA perspective and have different research and innovation demands.

The three phases form a coherent roadmap from research to ecosystem to industrial deployment.

Phase 1 creates sovereign technologies.

It addresses foundational research, advanced packaging, interfaces, NoC interoperability, safety-critical chiplets, optical interconnects, and future architecture disaggregation.

Phase 2 creates a trusted and interoperable European chiplet ecosystem.

It establishes governance, Joint Undertaking platforms, metadata layers, EDA flows, defensive IP structures, and multi-vendor interoperability.

Phase 3 creates scalable European compute platforms for automotive and adjacent industries.

It industrializes platforms, qualifies chiplet systems, enables package-based product families, aligns with OEM SDV roadmaps, and creates commercially viable European automotive SoC platforms.

Together, these phases provide a European partnership roadmap for closing current gaps in chiplet integration. The overall ambition is to move from fragmented research and isolated demonstrators toward a scalable, interoperable, and industrially deployable European chiplet ecosystem.

4.2.1. Build the Foundation

Target Horizon: 2037–2040+

Technology Readiness: TRL < 5

Objective: Foundations & Technology Building Blocks

Establish the scientific, technological, and architectural foundations required for long-term European sovereignty in chiplet-based computing systems.

The focus is not on immediate products but on developing the building blocks that will define future generations of vehicle computing platforms, industrial control systems, robotics, aerospace systems, sovereign infrastructure, and critical digital services.

Current Gaps

Europe currently lacks leadership positions in several fundamental technologies required for future chiplet systems:

Advanced Packaging Technologies

- Limited European industrial capability in advanced 2.5D and 3D integration
- Dependency on non-European packaging ecosystems
- Limited automotive-grade packaging expertise for large heterogeneous systems

Chiplet Interconnect Technologies

- Missing next-generation coherent die-to-die standards
- Limited research into long-reach chiplet interconnects
- No common architecture for heterogeneous inter-chip communication

Network-on-Chip Ecosystem

- Today's NoC solutions remain largely proprietary
- Lack of interoperability between NoC vendors
- No common abstraction layer for heterogeneous chiplet fabrics

Safety-Critical Compute Building Blocks

- Lack of reusable safety and security chiplets
- Limited availability of automotive-grade sovereign IP
- Absence of standardized chiplets usable across markets

Emerging Compute Architectures

- Need for future AI-native architectures
- Memory-centric computing approaches
- Optical and photonic integration technologies

Adaptive System Management

- Today's chiplet systems are largely statically configured
- Limited capability for runtime workload optimization
- Lack of self-healing and fault-aware architectures
- No common frameworks for dynamic resource orchestration
- Increasing system complexity requires autonomous management mechanisms

Autonomous Security and Self-Protection

- Security mechanisms in today’s embedded computing platforms remain largely reactive.
- Hardware-based threat detection and containment capabilities are still limited.
- Increasing connectivity, AI integration, and software-defined vehicle architectures expand the attack surface.
- Current systems do not provide sufficient autonomous cyber-resilience capabilities at chiplet or package level.
- Future platforms require security architectures that can detect, isolate, and mitigate threats in real time.

Long-Term Post-Quantum Cryptography

- Existing cryptographic algorithms may become vulnerable to future quantum computing capabilities.
- Automotive and industrial platforms require security lifetimes of 15 to 20 years or more.
- There is limited availability of hardware-accelerated post-quantum cryptographic solutions for embedded and automotive use cases.
- Current architectures lack sufficient cryptographic agility to migrate securely across algorithm generations.
- Europe needs stronger sovereign security building blocks for long-lifecycle critical infrastructure and vehicle platforms.

Distributed Intelligent Sensing

- Sensor processing remains highly centralized in many current architectures.
- Large volumes of raw sensor data increase communication overhead, latency, and energy consumption.
- Intelligence is not yet sufficiently integrated close to the source of data generation.
- Existing perception architectures often depend on centralized compute pipelines rather than distributed intelligence.
- There is limited integration of sensing, preprocessing, AI acceleration, and communication within a single package.

Optical Chiplet Communication

- Electrical interconnects face increasing scalability limitations in bandwidth, latency, reach, and energy efficiency.

- Future AI-native and disaggregated systems will require communication performance beyond conventional electrical package interconnects.
- Europe has limited industrial maturity in optical chiplet integration and photonic package technologies.
- There is no established optical communication standard for package-level chiplet systems.
- Current photonic integration capabilities are not yet sufficiently aligned with automotive-grade reliability, safety, and cost requirements.

Strategic Research Areas

Advanced Packaging and Integration

Research shall establish European competencies in:

- Glass substrates
- Silicon and organic interposers
- 2.5D integration
- 3D stacking
- Thermal management technologies
- Advanced cooling concepts
- Co-packaged memory
- Integrated photonics

The long-term goal is to create European manufacturing capabilities that reduce dependency on external advanced packaging ecosystems.

Open Chiplet Interface Technologies

Future chiplet systems require interoperable communication mechanisms.

Research activities should focus on:

- Low-cost chiplet interfaces
- Ultra-low-power interfaces
- Long-reach package-scale communication
- Optical chiplet interconnects to achieve high-bandwidths over long reach
- Coherent inter-chip communication

Particular attention should be given to optical interfaces capable of maintaining extremely high bandwidth over distances exceeding the practical limits of electrical interconnects.

Coherent Network-on-Chip Architectures

One critical gap is the absence of interoperable NoC ecosystems.

Europe should establish:

- Vendor-independent coherence models
- Common transaction protocols
- Standardized NoC abstraction layers
- Interoperable topology models
- Cross-vendor resource management

The vision is that future chiplets developed by different organizations can operate within a common system fabric regardless of the underlying NoC supplier.

Reusable Safety and Security Chiplets

Automotive expertise represents a uniquely European strength.

A research program should create reusable:

- Safety controller chiplets
- Security islands
- Trust anchors
- Hardware security modules
- Functional safety supervisors

These building blocks should be designed for deployment beyond automotive markets, enabling reuse across:

- Industrial automation
- Defense
- Aerospace
- Critical infrastructure
- Medical devices

This cross-market approach improves economic viability while strengthening strategic autonomy.

Disaggregated Future Computing Architectures

The roadmap should actively encourage further SoC disaggregation.

Research should investigate:

- Fine-grained compute partitioning
- Memory disaggregation
- Distributed AI architectures
- Heterogeneous accelerator integration
- Chiplet libraries enabling architecture composition

The objective is to move beyond large fixed SoCs toward dynamically composable systems.

Self-Adaptive Chiplet Systems

Future computing platforms should be capable of dynamically adapting to changing workloads, operating conditions, and system health. Self-adaptive chiplet systems enable runtime orchestration of compute, memory, communication, safety, and accelerator resources across heterogeneous chiplets.

Research activities should focus on:

- Dynamic workload distribution across heterogeneous chiplets
- Runtime resource allocation and optimization
- Adaptive power and thermal management
- Fault-tolerant and self-healing architectures
- Reconfigurable performance scaling
- AI-assisted system orchestration
- Health-aware operation of chiplet-based platforms
- Runtime balancing between performance, safety, energy efficiency, and reliability

The objective is to create chiplet-based computing platforms that continuously optimize themselves throughout their operational lifetime.

Self-Protecting Architectures

Future computing systems must evolve from passive protection toward active self-defense capabilities. As systems become more connected, autonomous, and software-defined, security must become an intrinsic and adaptive property of the hardware platform.

Research activities should focus on:

- Autonomous threat detection
- Hardware-based attack containment
- Self-monitoring security subsystems
- AI-driven anomaly detection
- Adaptive security policies
- Autonomous recovery mechanisms
- Runtime isolation of compromised chiplets or subsystems
- Hardware-rooted trust and attestation across chiplet systems

The objective is to establish computing platforms capable of detecting, isolating, and mitigating cyber threats in real time while maintaining safe and reliable system operation.

Long-Term Post-Quantum Cryptography

Future vehicle and infrastructure platforms require security mechanisms that remain trustworthy throughout their operational lifetime. Since future quantum computing capabilities may compromise currently deployed cryptographic algorithms, long-lifecycle platforms must be prepared for quantum-resilient security.

Research activities should focus on:

- Post-quantum cryptographic accelerators
- Quantum-resistant key management
- Hardware trust anchors for quantum-safe security
- Secure update mechanisms
- Cryptographic agility frameworks
- Long-term identity protection concepts
- Quantum-safe secure boot and attestation
- Efficient implementation of post-quantum algorithms in embedded platforms

The objective is to establish security foundations resilient against future quantum computing threats.

Sensor-in-the-Package Intelligence

Future computing platforms should move intelligence closer to the source of data generation. Sensor-in-the-package concepts integrate sensing, preprocessing, AI acceleration, and communication capabilities at package level to reduce latency, bandwidth demand, and energy consumption.

Research activities should focus on:

- Integrated radar, lidar, imaging, and sensing solutions
- Edge AI processing within sensor packages
- Sensor fusion at package level
- Low-latency perception architectures
- Distributed intelligence concepts
- Energy-efficient sensor processing
- Package-level preprocessing and event detection
- Close coupling between sensor data acquisition and AI inference

The objective is to reduce raw data movement, improve system responsiveness, and enable highly scalable distributed sensing architectures.

Optical Chiplet Interconnects

Future high-performance computing systems will require communication technologies beyond the limits of electrical interconnects. Optical chiplet interconnects offer a path toward ultra-high-bandwidth, low-latency, and energy-efficient communication between chiplets, packages, and potentially distributed compute nodes.

Research activities should focus on:

- Optical die-to-die communication
- Photonic interposer technologies
- Optical Network-on-Chip architectures
- Co-packaged optics
- Package-to-package optical links
- Ultra-low-power optical communication fabrics
- Automotive-grade photonic integration
- Optical communication concepts for AI-native and disaggregated computing systems

The long-term goal is to establish optical communication infrastructures capable of supporting massively interconnected chiplet systems.

Expected Outcomes

By 2040 Europe should possess:

- Sovereign chiplet building blocks
- Open and standardized interface technologies
- Interoperable coherent NoC architectures
- Advanced, cost-competitive packaging capabilities
- Safety-certified chiplet components & self-protecting hardware architectures
- Photonic interconnect technologies
- Quantum-safe cryptographic hardware building blocks
- Sensor packages integrating sensing, AI processing, and communication
- Optical chiplet communication technologies
- Self-adaptive chiplet-based computing systems

These building blocks are expected to deliver along the open ecosystem requirements, meaning open specifications for interoperability and multi-vendor integration are considered.

4.2.2. Prototypes for next Generation

Target Horizon: 2034–2037

Technology Readiness: TRL 3–7

Objective: Collaborative Innovation & Ecosystem Enablement

Transform technological building blocks into an interoperable European chiplet ecosystem capable of supporting multi-vendor platform development and providing the governing framework.

Current Gaps

Today the industry lacks:

- Industrial governance models
- Multi-vendor qualification processes
- EDA tool supported exchange of common chiplet metadata
- Cross-vendor testing concepts
- Advanced node design capabilities
- System-level interoperability standards
- Shared ecosystem IP frameworks

Without these elements, chiplet innovation remains fragmented.

Strategic Research Areas

European Ecosystem Governance

A central priority is the creation of a governance framework.

The ecosystem must define:

- Technical governance
- Compliance processes
- Qualification mechanisms
- Certification procedures
- Lifecycle management
- Participation models

This governance becomes the long-term foundation enabling ecosystem scalability.

Joint Undertaking Platform Development

Joint European programs, similar like CHASSIS[4], should establish multiple chiplet integration platforms.

These platforms may target:

- Automotive HPC
- AI acceleration
- Edge computing
- Industrial control
- Robotics
- Aerospace

The objective is not a single platform but a portfolio of interoperable platforms sharing a common technological foundation.

Multi-Vendor Integration Framework

Europe must create the technical conditions enabling genuine mix-and-match integration.

This includes:

- Standardized chiplet descriptors
- Integration metadata

[4] CHASSIS Project | Chiplet-based Hardware Architectures for Software-Defined Vehicles - CHASSIS-Project

- System assembly frameworks
- Verification methodologies
- Common packaging requirements
- Contribution to and extending standards like “IEEE-P3537 3Dblox” / “JEP30 CDX-ML” on system level

The goal is enabling independent chiplet development while preserving system-level compatibility.

EDA Toolchain Evolution

Perhaps the largest ecosystem gap resides in EDA infrastructure.

Future tooling must support:

- Chiplet marketplaces
- Cross-company design flows
- Standardized integration models
- Package-level optimization
- Multi-vendor verification
- Automated system composition

The roadmap should identify which EDA abstractions are required to make heterogeneous chiplet development practical.

Advanced Node Design Capabilities

Europe has only a small number of companies with proven tape-out experience at advanced nodes:

- Scarce advanced-node tape-out competence in Europe (≤5nm).
- Lack of automotive-grade qualification and validation flows for leading-edge SoCs.
- Dependence on non-European EDA ecosystems, IP portfolios, and expert talent.

Defensive Patent Ecosystem

An additional strategic element is the creation of a European defensive intellectual property framework to safeguard the ecosystem.

The objective is not protectionism.

Rather, participating organizations establish:

- Shared interoperability principles
- Cross-licensing approaches
- Standard-essential IP frameworks
- Defensive patent mechanisms

This creates a protected innovation space where European organizations can collaborate while reducing exposure to external patent aggression.

Hybrid Architectures

The future will not be exclusively chiplet-based.

Phase 2 therefore includes:

- Monolithic SoCs
- Hybrid chiplet systems

- Package-level scaling
- Modular upgrades

Scalability must remain seamless across all implementation technologies.

Expected Outcomes

By the end of Phase 2 Europe should possess:

- A functioning chiplet ecosystem
- Established governance
- Interoperable development environments
- Multiple demonstrator platforms
- Cross-vendor integration capabilities
- Ecosystem-level IP structures

4.2.3. Empower & De-risk

Target Horizon: 2031–2034

Technology Readiness: TRL 8–9

Objective: Industrial Deployment & Volume Adoption

Establish a European capability to design, integrate, qualify, industrialize, and deploy scalable AI-enabled semiconductor platforms based on open chiplet architectures, spanning from silicon IP to OEM applications.

Translate ecosystem innovations into production-ready platforms deployed by European OEMs and industrial customers. This can lead to development and qualifications of SoC computing platforms which can integrate computing extensions via chiplets in the subsequent steps.

Results will contribute to open platforms, allowing integration into a multi-partner chiplet system, multi-industry ecosystem to maximize market impact.

Key Industrial Gaps

- Automotive-grade qualification and certification
- Multi-vendor supply chain management
- Commercial risk and responsibility models
- Long-term support and lifecycle management
- Volume manufacturing readiness
- Vertical integration between semiconductor and OEM ecosystems

Strategic Research Areas

1. Base Die Platforms
   Reusable sovereign SoC foundations with open D2D interfaces.
2. AI Chiplet Ecosystem
   Multi-vendor AI accelerators enabling scalable, cross-industry computing platforms.
3. Chiplet Industrialization
   Packaging, integration, supply-chain readiness, and ecosystem business models.
4. Platform IP
   Reusable safety, security, middleware, and virtualization building blocks.

5. Qualification Framework
   Silicon, software, package, and multi-vendor certification methodologies.
6. Platform Software
   Drivers, middleware, AI runtimes, toolchains, and long-term support.
7. System Integration
   ECU, thermal, power, and vehicle-level integration.
8. OEM Verticalization
   Joint development, deployment, and evolution of production platforms.

Industrialization Priorities

Platform Industrialization

- Qualification and certification
- Lifecycle and cost management
- Production ramp-up
- Supply assurance

OEM-Centric Platform Development

- SDV platform architectures
- Vehicle integration
- Platform evolution
- Feature-on-demand enablement

Package-Based Product Families

- Reusable chiplet libraries
- Configurable SoC variants
- Exchangeable AI accelerators
- Market-specific platform differentiation

European Automotive SoC Platforms

- AI-native computing architectures
- Centralized vehicle computing
- Integrated ADAS and IVI
- Safety and cybersecurity by design

Ecosystem Sustainability

- Governance and funding models
- Open standards maintenance
- Long-term interoperability
- Ecosystem stewardship

Long-Term Support

- Platform maintenance over vehicle lifetime
- Security and software updates
- BSP and base die support
- Responsible ecosystem operators

Expected Outcomes

By 2031–2034 Europe should have:

- Production-grade European chiplet platforms
- Multi-vendor qualification processes
- Industrialized packaging and integration ecosystems
- Automotive volume deployments at OEMs
- Scalable product families based on reusable chiplet libraries
- Sustainable European semiconductor ecosystem and governance model