The Road towards a High-Performance Automotive RISC-V Reference Platform

Report from European Working Group

14 April 2023

Annex 3 to the Report

Roadmap for Future Automotive-driven RISC-V Developments in Europe

16 February 2024

EXECUTIVE SUMMARY

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"[1]. That document formed the basis to launch two projects in the Chips-JU framework program (TRISTAN[2] and ISOLDE[3]) and another Working Group to develop a "Roadmap for Future Automotive-driven RISC-V Developments in Europe".

This resulted in a number of Roadmap Elements which are of key importance for future RISC-V (System on Chip) SoC developments. Important elements are:

- Although focus is on high-performance automotive requirements, the roadmaps also contain important (hardware) IP-elements which can be used in other vertical application domains, as e.g. industrial automation, communication, health and others.
- Within the next 8 years, 8 key milestones / demonstrators are planned, based on real silicon:
    - Q2/2027: Octo-Core AI Engine
    - Q1/2028: EURV32-RT Automotive 32-bit Controller
    - Q2/2029: EURV64-RT Automotive 64-bit Controller
    - Q2/2029: Dual octo-core AI Engine
    - Q4/2029: EURV64-AP – Application Processor
    - Q4/2030: Quad-EUR64-AP
    - Q3/2032: Heterogenous compute cluster
    - Q3/2033: 16xEURV64-AP
- Detailed roadmaps have been generated for IP-elements which can be used and re-used for different processor architectures and SoCs, both for different types of automotive systems as well as for other application domains.

[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)

0. Introduction

Context

The automotive industry has traditionally been accustomed to a multitude of interconnected *Electronic Control Units* (ECUs), each with their own specific function, architectures, software stacks and operating systems. As more functionalities and safety features were incorporated into the vehicle, some organisation of ECUs has taken place and different units communicate via a centralised gateway. With the evolution of the car and the functions now expected from the end-costumer, such as better infotainment services and autonomy, this model is becoming increasingly untenable. The presently prevalent decentralised architectures have significant drawbacks when it comes to scalability and communications performance.

The current trend is such that an increasing number of functionalities are software defined. In fact, it is projected that the number of lines of code per vehicle will go up from the current 100 million to 1 billion by 2030. From a hardware perspective, the increased autonomy being incorporated in the automotive sector necessitates that on-board computing is centralised, otherwise it would be untenable to network the increasingly interdependent ECUs together due to latency and the scale of the wiring harness, amongst other limitations.

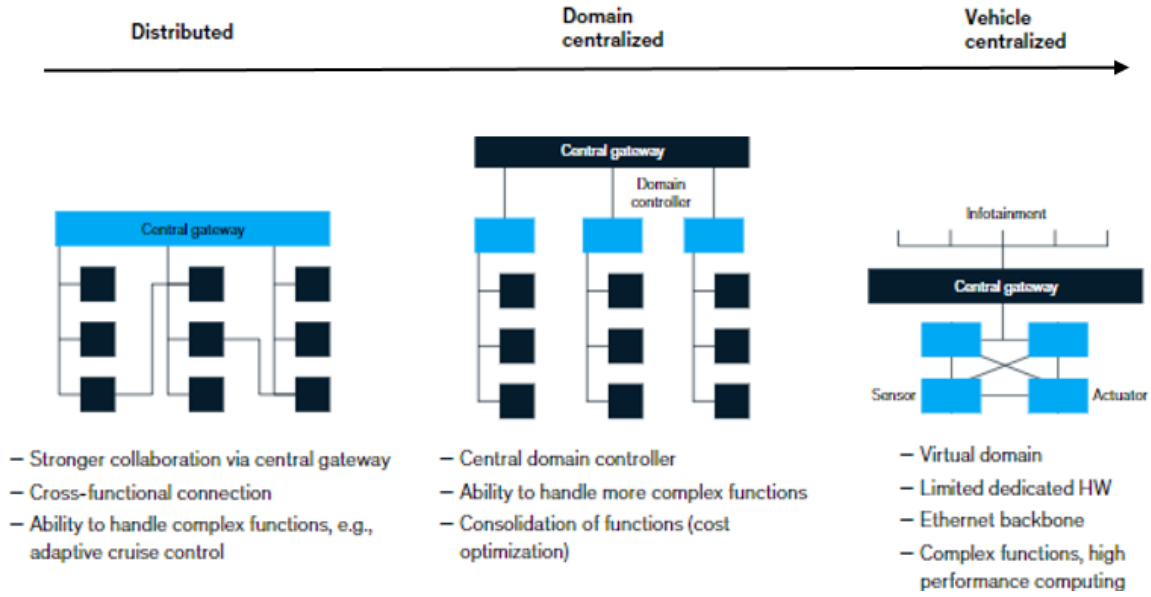

Figure 1 - Evolution of E/E architecture (Source: McKinsey & Company)

The solution to these issues is the centralisation of the E/E architecture whereby a number of ECUs are consolidated into so called *Domain Control Units* (DCUs) or *Zonal Control Units (ZCUs)* that integrate different functions into a more cost-effective solution. It is foreseen that eventually these DCUs/ZCUs will be networked to a centralised SoC that integrates and consolidates the different functions such as autonomous driving, infotainment and cabin control.

With this in mind, the RISC-V Automotive Hardware Platform initiative is designed to consolidate and strengthen Europe's leadership position in the automotive electronics sector. This initiative comes at a critical time as the automotive industry faces fundamental shifts with the centralisation of electrical/electronic (E/E) architecture in vehicles and the advent of more electrified, automated, connected, and ultimately shared modes of transportation.

Roadmap

The objective of this document is to outline a strategic framework for collaborative industrial research on a comprehensive ecosystem of processor cores and peripherals designed to support the diverse needs of a centralised electrical/electronic (E/E) architecture. The envisioned reference platform aims to facilitate the development of processor cores across all tiers of the automotive processor hierarchy prioritising both performance and energy efficiency.

Furthermore, the prospects for this market are promising. The cost share of electronic components in relation to the total value of the car is expected to grow from around 16% currently to around 35% by 2025.The automotive software and E/E hardware will grow at a rate of 7% per annum until 2030 from approx. USD 238 billion to USD 469 billion.

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.

To address this ongoing transition and guarantee that the recently launched Chips Joint Undertaking also accounts for critical industry sectors, including the European automotive industry, the Vehicle of the Future initiative, spearheaded by industry stakeholders, is championing two interlinked projects: the creation of a RISC-V based automotive hardware platform, and the formation of an open, software-defined vehicle ecosystem led by European automotive manufacturers and suppliers.

These initiatives should be seen in the context of the recently adopted Chips Act whose key aim is to build Europe’s strategic autonomy in the semiconductor industry across the value chain, from design all the way to manufacturing. These initiatives are another step forward in generating more demand for greater fab capacity in Europe and thus contributes to the sustainable fulfilment of the Chips Act’s goals.

To meet its Digital Decade goal of manufacturing 20% of the world's leading-edge semiconductor chips within Europe, it is crucial for Europe to capitalise on the demand generated by major industrial sectors, like the automotive industry. Doing so will generate additional manufacturing demand, which, in turn, will have a ripple effect—fortifying and expanding Europe's industrial foundation in semiconductors as a whole.

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.

1. Composition of the Roadmap

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

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, 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.

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.

The roadmaps including specific information on the roadmap elements, are explained in the sections below.

2. Scalable RISC-V Automotive Control Processors

2.1 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 compliance in the safety and security fields.

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)

2.2 Detailed features of the RISC-V automotive core family

It will be based on a 32-bit and 64-bit family, the second one featuring virtual memory support. Each core will have support for multi-processor configuration, integrated with the system's ability to configure itself into multiple lock-step synchronized cores.

Below are the features in details:

- Superscalar architecture
    - Higher performance than competitors in the field, in terms of Instruction Level Parallelism and microarchitecture implementation.
- Fast context switch

    - Multi-threading support
    - Fast and deterministic interrupt/exception response
- Virtualization support with Hypervisor
- Integrated vector unit
    - Efficient and scalable solution in terms of performance per unit area
    - Container of most of the custom extensions in the application fields (DSP, AI, Networking, Crypto, etc.)
    - Chained register support, multi-way and OOO execution
- Co-processor interface for external accelerators
- Safety and security (extended to memory and interconnect)
    - Spatial and temporal redundancy for temporary and permanent faults, ASIL compliance
    - Security features, SESIP compliance
- Multi-processor HW support

2.3 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.

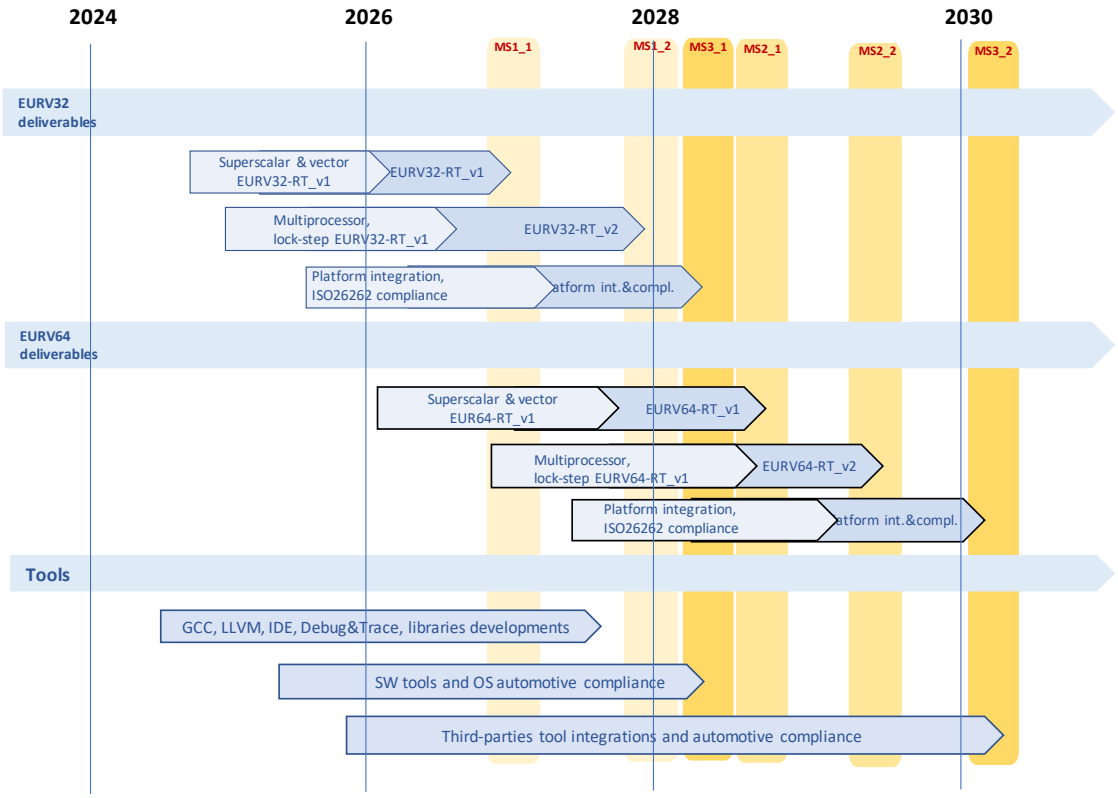

3. High-Performance RISC-V based Application Processors

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.
First needs are towards a **64-bit in-order processor** with a memory management unit (MMU) and a full hypervisor extension. A vector unit is needed here as an optional component. This in-order machine can be directly based on the Automotive RISC-V control processor line.
The computing demands of future platforms can only be satisfied through out-of-order CPUs. The performance class of an out-of-order CPU may be referred to (among others) by the issue width, i.e. how many instructions the CPU can issue for execution in one clock cycle.

3-wide out-of-order CPUs are a good starting point. This solution can be scaled by developing a quadcore cluster made up of those.

To supply compute power for continuously increasing demands, the issue width has to be increased to **4-, 6-, and eventually 8-wide machines**. Also, the number of cores in a cluster must be increased, eventually leading to **16-core clusters**.

Along with the core 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 (1) **shared cache memories** capable of supporting a large number of outstanding requests from abundant 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 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.

The developed cores can be used in **heterogenous** configurations. A development for this integration needs to happen in addition to the core development. Such configurations would combine processors with maximum power efficiency with cores that deliver maximized performance.

Instead of a real-time Operating System (OS) of the control cores in Section 2, the application processors run Linux and/or Android OSs.

Matrix Extensions will be needed and are also covered in more detail in Section 4.

For the application processors, no significant gap is foreseen in terms of SW support. This is due to these main facts

1. With the in-order 64-bit support, a complete toolchain for 64-bit automotive RISC-V is already in place
2. An application processor needs to be designed to be able to run existing applications. This is in contrast to specialized embedded cores, where differentiation plays a bigger role
3. The porting effort for rich OSs like Linux and Android is already started. This is a major effort, but expected to be complete by the time the application cores are available.

Some detailed features of the RISC-V platform system IPs based on the fact that System IPs will eventually be critical for RISC-V CPUs to get traction:
Some specific requirements in the **Industrial Automation** domain are:

- 4x~8x 64bits (comparable with Cortex A73) with virtualisation is basically the standard
- GPU
- Video processing acceleration
- Neural processing acceleration
- LPDDR4, PCIe, USB host, camera/display support

Real-time processors (Cortex M3/M4) have a stable base of aficionados but it's rather a niche.

For **robotics**, the bottleneck is not in the compute architecture, but rather in the networking stack and communication middleware.

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

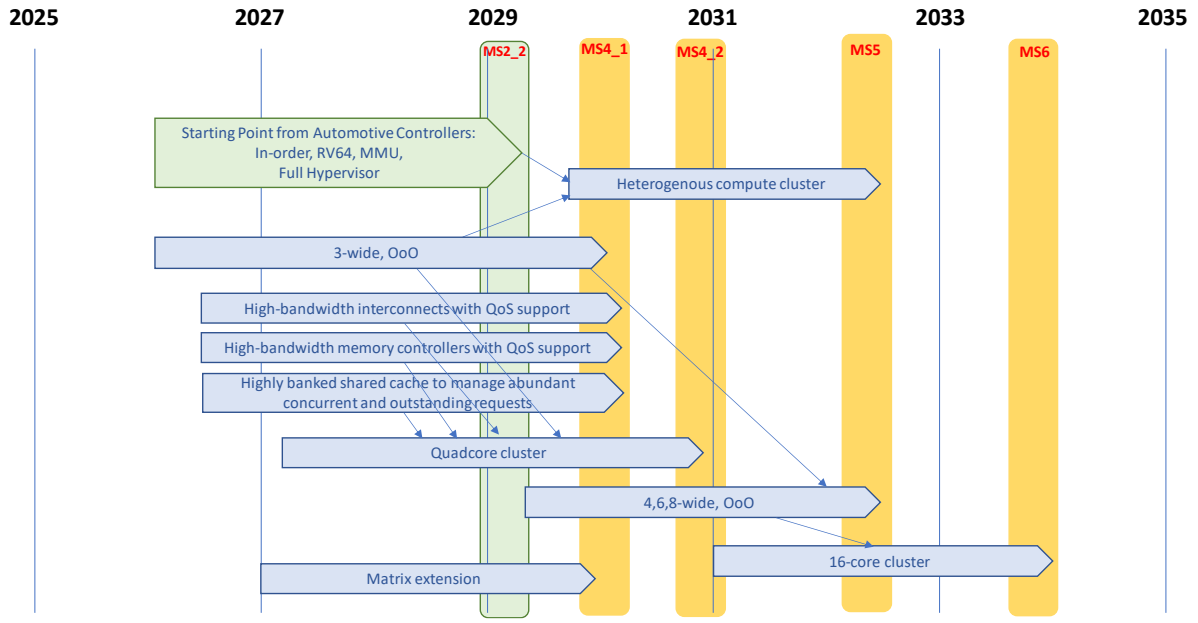

4. AI and ML Accelerators

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 ML 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.

3. 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, 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.

4. 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. 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.

5. 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.

6. A well-designed **memory hierarchy** is critical for AI and ML 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.

7. 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 straight forward (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**.

8. **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. A comprehensive software stack and support for domain-specific languages round out the requirements for a sophisticated and efficient acceleration processor.

The AI and ML Accelerators roadmap is shown below:

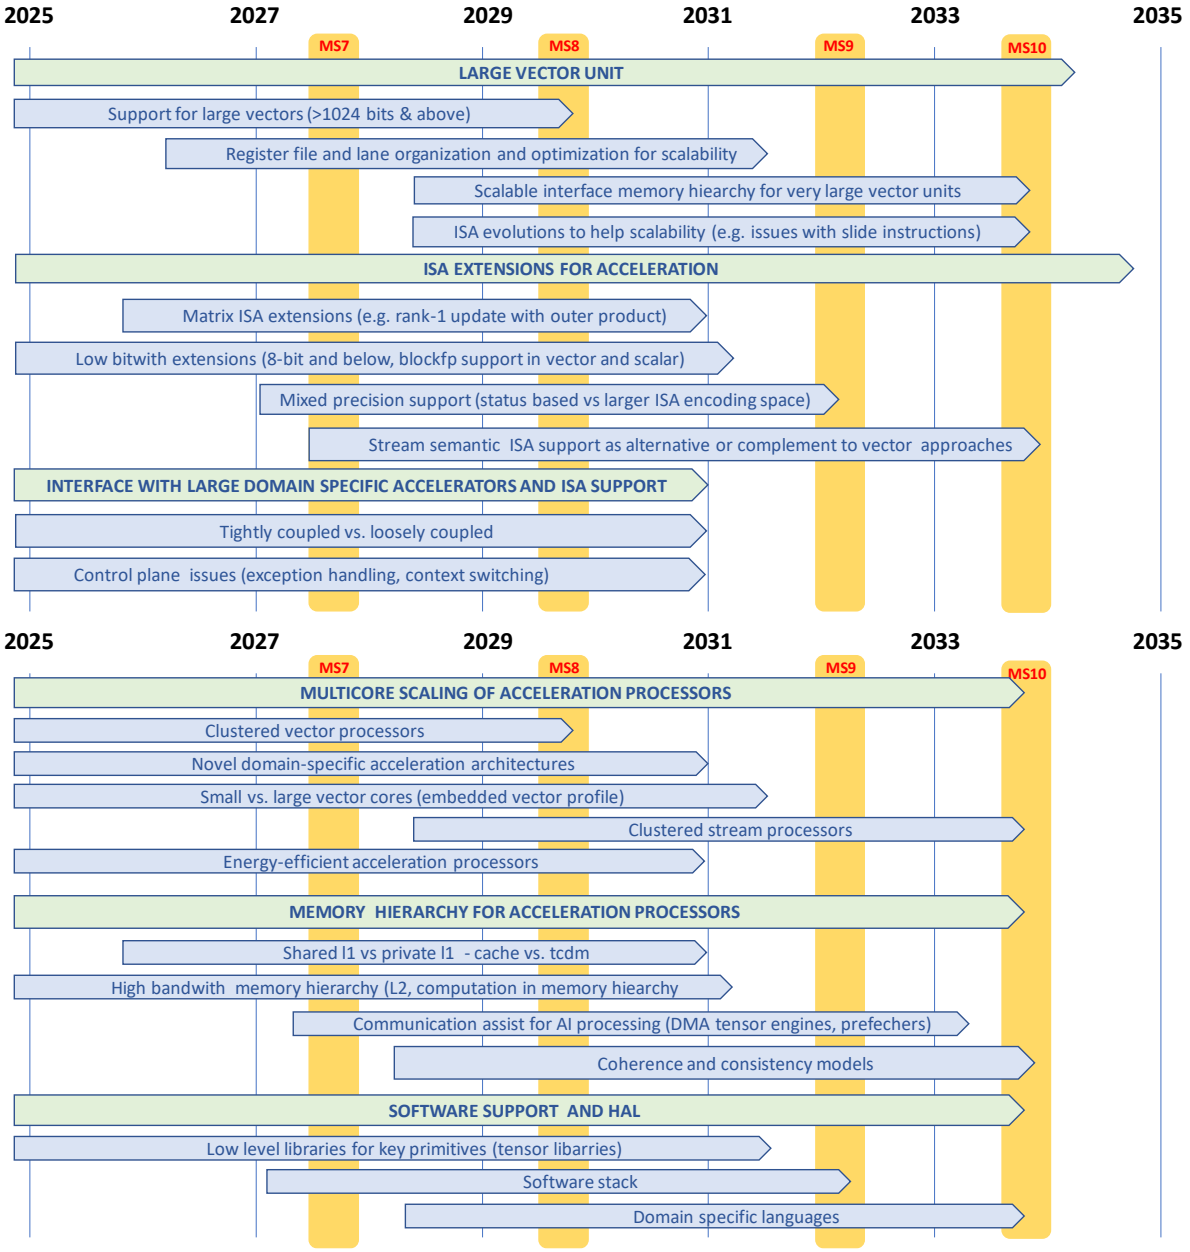

5. System Integration and Interfacing

The integration of different IPs into a product e.g., a SoC or Chiplet for an intended use-case requires considering the interaction of various IPs to build exceptional and differentiating products. Early pre-silicon integration of the relevant IP components together with an appropriate SW 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 IPs. Especially for newly designed IP or new combinations of different IPs an early verification of the overall system functionality is essential. The integration of open-source IP requires additional efforts in terms of verification and certification. While the source code is open and available to each of the 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. At 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:

- **System Peripherals.** We consider system peripherals as a subset of peripherals that are tightly integrated with the overall function of the system e.g., DMA, interrupts. In contrast to more loosely coupled peripherals which perform more isolated and independent functions. Orchestrating the system peripherals is typically critical to reach an overall system performance and requires predicting intended workloads and usage of the attached components.

- **Memory and Buses.** These central components have a significant impact on the overall performance and functionality of the system. Multicore systems and new components like NoCs or chiplet interfaces pose new challenges for predictable timing behavior under different system scenarios. Additionally, the integration of accelerators and potential memory transfers can have a severe impact on system performance.
  For shared cache memories, the following is needed to ensure data coherency & consistency:
  - High-bandwidth interconnects:
    - Define now or add a definition phase for the protocols the cores should support: CHI.E, CHI.G, and AXI for better interoperability with peripherals that must work with RISC-V and other ISAs.
    - Collect requirements for IOMMU, GIC, Security and Safety considerations as well as mechanism for debug and trace.
  - High-bandwidth memory controllers
    - DMA tensor engines, , tcdm and high be memory hierarchy, etc.
  - Power management (dynamic clock and voltage trimming and gating)
    - dynamic clock and voltage trimming and gating require some control to prevent deadlocks when off-lining some units and proper initialization when onlining ones.

- **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 the security actions inside a dedicated RISC-V microprocessor unit, based on the following prinicples:

- Physical separation between user applications and security
- Easier implementation of side-channel attack protection because not extended to the full platform
- Single Root of Trust based on the internal RISC-V design

The envisioned RISC-V controller cores belong to a class of microprocessors as defined in the TRISTAN and ISOLDE projects with the following features:

- Side channel protections
- HW and SW IP: high data throughput Vector Crypto or accelerator, Hash unit
- Random true generation, PUF
- SESIP compliance

The System Integration and Interfacing roadmap is shown below:

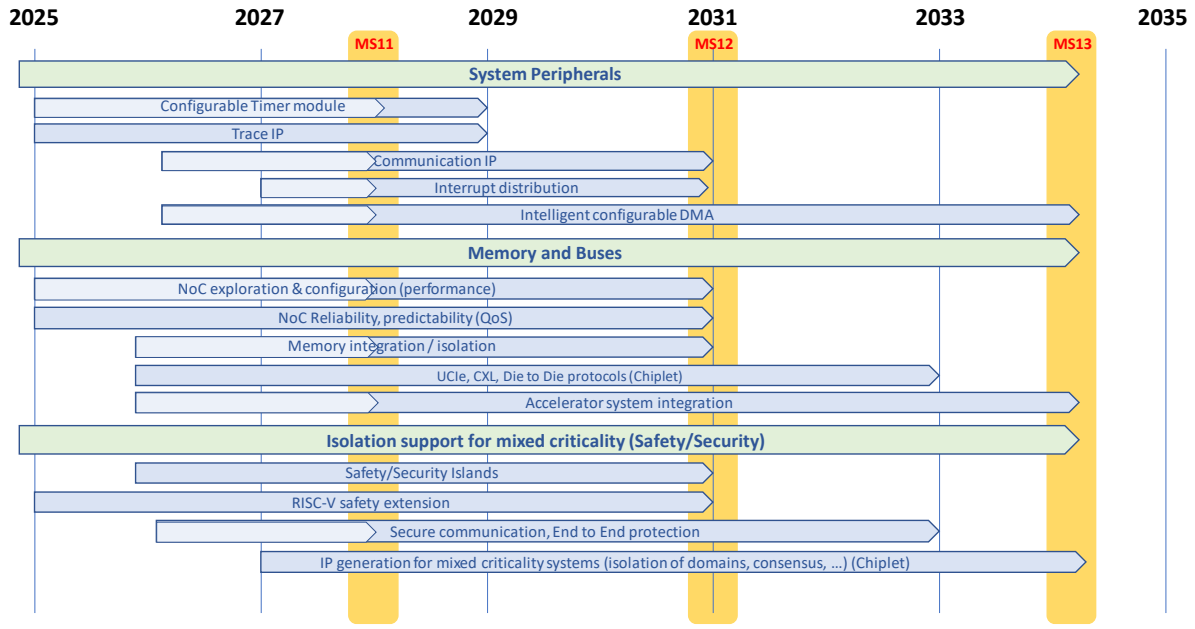

6. Milestones

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
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.

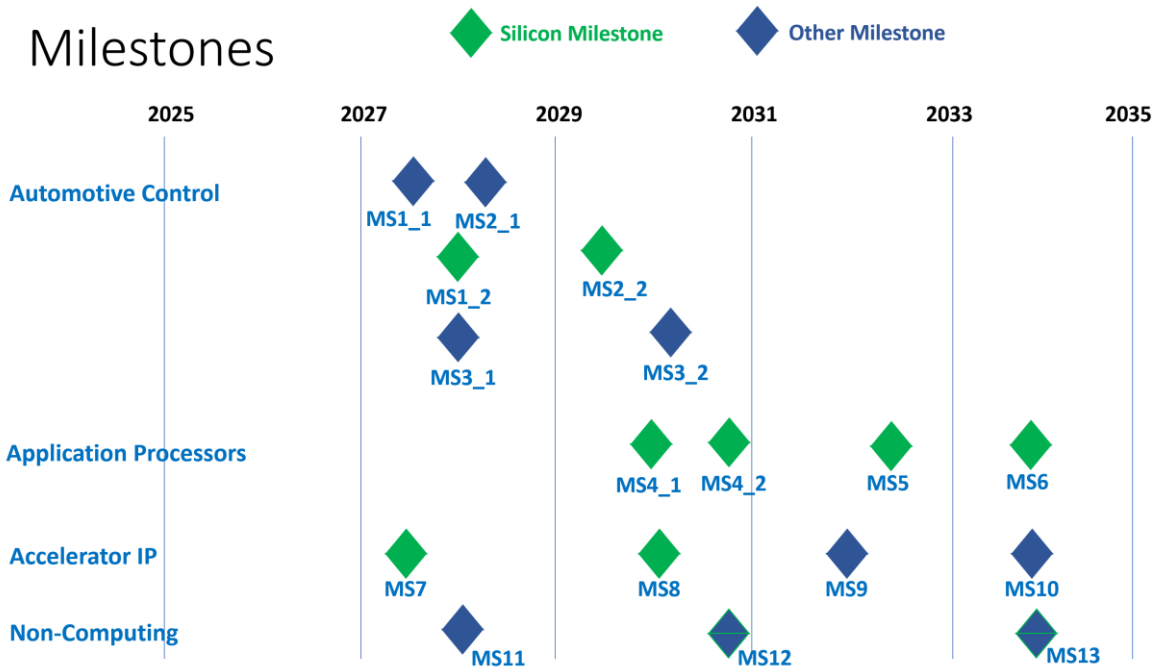

7. Estimated funding

This roadmap of the Automotive Controllers (Category 1) is based on a parallel development of the 32bit core families, 64bit core family and tools.
This global project requires a minimum investment, so organized:

- EURV32-RT development, multiprocessor/safety, compliance
    - Architecture, modeling, design, verification: average 30 persons/year
- EURV64-RT development, multiprocessor/safety, compliance
    - Architecture, modeling, design, verification: average 30 persons/year
- EURV-Tools
    - Compiler, binary utilities, IDE, libraries, compliance, third-party integration support: average 20 persons/year

The roadmap of the Applications Processors (Category 2) is based on an incremental development and enhancement of a core family. The effort to develop competitive out of order processors that are in line with the compute requirements of future Automotive applications is extremely high. To achieve an IP delivery in line with the proposed timelines, an R&D team of ~150 engineers is needed.

The roadmap of the AI and ML Accelerators (Category 3) - is based on a roadmap of two generations of in-memory compute accelerators capable of supporting advanced AI models suitable for integration in SoC and chiplet form, along with the supporting software development environment.
This platform component requires a minimum investment of 85person/year, across the following workstreams:

- Octal AI engine/chiplet development
    - Architecture, modelling, design and verification: 40 persons/year
    - Compiler, runtime, ML libraries, integration support and compliance: 30 persons/year
    - System integration, emulation, firmware, packaging and board design: 15 persons/year

The roadmap of the System Integration Interfaces (Category 4) is based on the continuous assessment of the developed HW artefacts with the available SW along the defined evaluation metrics. Building a system with APIs and DSLs that integrates well with different infrastructures and IPs is essential for avoiding repetitive steps and increased manual effort (5-10 persons/year). We assume the integration effort of SW and HW to increase towards milestones as the number of iterations and interactions increase (15 persons/year) compared to intermediate preparation and iterations (5 persons/year). For generation and specification of new peripheral IP additional capacity is needed to analyze, implement and test the required functionality (5 persons/year). These efforts don't cover efforts for SW modification, adaption or workload modeling (10 persons/year).

Overview of estimated costs:

Estimated Effort RISC-V Roadmap (pyr)										
Topic	**2025**	**2026**	**2027**	**2028**	**2029**	**2030**	**2031**	**2032**	**2033**	**2034**
EURV32-RT	30	30	30	20						
EURV64-RT		30	30	30	30					
EURV SW		20	20	20	20	20				
EURV64-AP		100	100	100	100					
Quad-EURV64-AP		50	50	50	50	50				
big.LITTLE						100	100	100		
16xEURV64-AP							50	50	100	
Accelerator IP Architecture & Design	25	40	40	40	40	40	40	40	40	25
Accelerator IP Software	20	30	30	30	30	30	30	30	15	15
Accelerator IP Integration		15	15	15	15	15	15	15	15	15
System Integration Infrastructure	5	10	10	10	10	10	5	5	5	
System IP Integration and Generation		10	10	20	10	10	20	10	10	10
SW Optimization and Workload Generation		10	10	10	10	10	10			
Total	**80**	**345**	**345**	**345**	**315**	**285**	**270**	**250**	**185**	**65**
Grand Total	**Effort:**	**2485**	**Cost:**	**~**	**522**	**Meuro**				

8. 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.

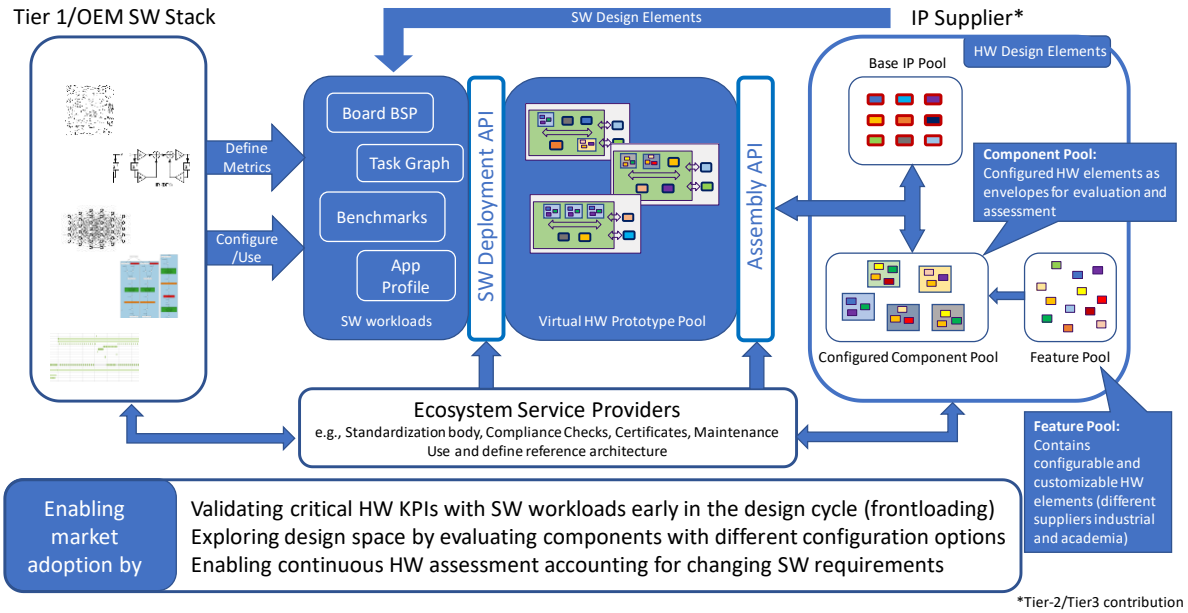

The advantage of such a distributed HW/SW Co-Design setup is that the different suppliers in the supply chain can focus on their core competencies while remaining competitive addressing a larger market.

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 scalabilty 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.

9. Feedback from OEMs and TIER1s

We received initial feedback from OEMs which did not want to have their name officially published in the report. Further meetings can and will be planned with them to further discuss. Currently discussions are on-going with their legal departments to see how and when they can launch an ''official'' view from their side.
Some of the feedback has been directly integrated within the document and roadmaps. Some elements are at a different level of granularity and will be taken into account if more detailed requirements are being derived for concrete developments in the domain.
In italics text below, we provide an answer to the original comments from OEMs and TIER1s.

Currently for next generation ECU platforms our thoughts from OEMs are in the direction of:

1. Integration real-time cores and application cores into a standardized core architecture
    a. Realtime capabilities with Application core performance (Keyword SafePosix) – *covered in EURV64-RT/AP planned activities*
    b. ASIL-D with MMU – *covered in EURV64-RT planned activities*
    c. A cycle time of 10ms is also sufficient for real-time requirements – *to be covered at platform/system-level and not at processor level*
2. SoC/Software Level – *to be covered at platform/system-level; the RISC-V roadmap will enable these features and requirements*
    a. QM and safety applications on the same SoC/MCU (Multicore, with different ASIL Levels)
    b. It must be ensured by HW/SW for example that an ASIL-D application only runs on an ASIL-D core
3. Startup times – *to be covered at platform/system-level; the RISC-V roadmap will enable these features and requirements*
    a. Always-on domains/HW support
    b. Fast bootup for quick system availability (at least partial) (<50ms)
4. HW Accelerators – *This is part of MS9 and contains standard machine learning frameworks*
    a. Simplicity of programming (preferably no special language required)

The following feedback was obtained from TIER-1 stakeholder:

Big.LITTLE concepts we have seen less in automotive, unless different types of CPU clusters. For example:

- Application Core for Linux or QNX, large applications, no real-time requirements, but possibly safety requirements
- RT Cores for Real-Time Applications, possibly (or very likely) safety requirements
- Safety cores for MACL applications, i.e., ASIL-B and ASIL-D applications
    - Big.Little, however, we could also reasonably imagine on the Application Cores domain if the applications can be moved between little < - > BIG as desired. Energy efficiency!.
    - Not suitable for ASIL-D though(!) ... feasible for ASIL-B. ASIL-B is a clear trend for application cores, ASIL-D is not.

This is agreed and will be taken into account in future developments.

Matrix extension = unclear? Does it mean "small" Matrix Multiply and Accumulate? Or more?
Matrix extension contains an attached matrix accelerator integrated into the Vector Unit of the processor core. This will follow the RISC-V standard in development.

Chiplet interconnect = UCIE!?
We acknowledge that UCIE could become a dominant chiplet interconnect standard and can be accommodated when needed.

INT8 / FP8 / FP16 or logarithmic number system is currently being discussed at Inference ML Cores.
These are discussed intensively in Special Interest Groups of RISC-V Internation and will be integrated by default in the proposed roadmaps.

If though ML/AI inference is addressed: The focus seems to be hardware development: I hope this is not the case, building the HW without having SW in mind is a medium-good idea. This becomes apparent when you see the current accelerators, especially the SDK and tools for deploying the SW to the very special hardware is decisive.
This is directly addressed in MS9.

Our issue is: The main advantage – which we are told again and again – should be the Risc-V extensions. We believe: That will be the main disadvantage because of unnecessary fragmentation of different IP cores from the get-go. The existence of a mature Eco System is very important to the Tier1s, and we see that fragmentation undermining this!
We do acknowledge this, especially for the Application Core space. In deeply embedded or specialized cores, we do see high value in differentiation and customization.

We had discussions with one single Risc-V IP and Chip(let) supplier, who took a slightly position on this:

a) Design Risc-V cores according to standardized ISAs and design Microarchitectures for the different use cases (App Core, RT Core, Safety Core), the microarchitecture is the differentiator!
b) Work with SW and other infrastructure parties to mature the Eco System (OS, IDE, SDK, Debugger, etc.)

c) In a second wave (when Risc-V is an acknowledged [superior] technology): Think about custom instructions, but also as part of defined extended ISAs.

We acknowledge this approach. In addition memory subsystems and cache play a big role in performance.

10. Counterfactual Scenario

First, it is important to note that SotA in the PC/consumer domain (in which many non-European stakeholders are active) is not the same as SotA in automotive, which is much more complex and challenging. Furthermore, RISC-V is an important option to go ahead in terms of energy efficiency, which is crucial in the automotive domain.

Therefore, if this Automotive RISC-V program would not get public funding, a momentum behind RISC-V for automotive would be lost. Europen IDM's and TIER1's would loose competitiveness on their respective markets. Without these powerful platforms, Europe's ability to innovate and advance automotive applications would be compromised. Moving towards a new generation of processor architectures, based on RISC-V - which has not been proven in the automotive market - is a key challenge and risk for European IDMs and needs public funding to achieve a successful outcome on the market.

Especially the development of the most advance Real-Time Processors and some specific Application Processor architectures would be endangered in case no European funded program is set up around these developments. It would lead to less products and less revenue for European IDMs.

11. Conclusions and Next Steps

This roadmap 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 current roadmaps 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). 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 products, 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. A first next step will be to indicated what elements in the roadmap are already partly covered in running projects, as e.g. TRISTAN and ISOLDE[4]. In addition, the goal is also to establish an Automotive RISC-V alliance to shepherd the activities around the reference platform.
The next steps 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 research topics in Europe.
In practice the Working Group recommends that these developments can be done by a series of pre-competitive funded projects, embedded in a Framework Partnership Agreement, followed up by a complementary approach via a next IPCEI program, in order to include first industrialization.

Contributors to the Roadmap:

Participants Working Group High-End Automotive RISC-V based platform		
Company	**Main Contact**	**Back-up Contact**
Moderator	Patrick Pype	
NXP Semiconductors	Christian Herber	
Infineon	Wolfgang Ecker	
STMicroelectronics	Elio Guidetti	Fabio De Ambroggi
Bosch	Ingo Feldner	Franck Bernard
Unibo/ETH	Luca Benini	Angelo Garofalo
Fraunhofer IIS	Norbert Schuhmann	Michael Konrad
PlanV	Massimiliano Giacometti	
Axelera	Paul Neil	Evangelos Eleftheriou ; Fabrizio Del Maffeo ; Giuseppe Garcea
BSC	Jaume Abella	Francisco Cazorla

[4] 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).