Whitepaper

When ambitions meet costs: E/E architectures for the future SDV

The balance between SDV ambitions and costs is challenging. This whitepaper presents E/E architectures and products meeting that sweet spot.

Invented for life

BOSCH

Table of contents

01
The SDV honeymoon is over

As early as 2022, industry projections indicated that nearly half of all new vehicles would integrate novel vehicle-centralized architectures by the end of the decade. This widespread ambition for adopting such systems stemmed directly from automotive manufacturers' (OEMs) strategic imperative to fully support the then-emerging paradigm of the software-defined vehicle (SDV). Consequently, these designs were intrinsically built to facilitate the deployment of SDV capabilities, such as over-the-air updates and continuous functional enhancements across diverse domains to enable new revenue streams and provide a competitive edge.

Along this trajectory, however, OEMs began confronting a significant increase in the projected total-cost-of-ownership (TCO) for their novel designs. Their responses varied: Newer OEMs, keen to differentiate themselves through innovative features, largely pressed ahead, ultimately accepting reduced earnings or even initial losses. In contrast, established OEMs re-evaluated their strategies, prioritizing cost-efficiency. Consequently, some have since reverted to more traditional architectures.

Among the first group, Chinese OEMs emerged as arguably the most prominent. Unburdened by established legacy, they rapidly deployed vehicle-centralized platforms, providing a decisive advantage in the prevailing domestic price war. These platforms enabled critical differentiation and expedited time-to-market through their support for over-the-air flashable software features. Their discerning, feature-oriented clientele demonstrated widespread expectation and acceptance that vehicles would continuously improve throughout their lifecycle, even tolerating the initial absence of certain standard functionalities upon delivery. This strategy proved pivotal in establishing their market presence within a tech-savvy milieu of young digital natives. Nevertheless, the combined pressures of diminishing returns, feature saturation, and an overall stagnating market now necessitate that even these OEMs rigorously pursue cost optimization within their foundational designs.

In contrast, the remaining OEMs—those that had either paused or cancelled the development of full vehicle-centralized platforms—are now pursuing a new "SDV evo" approach. This strategy primarily involves updating existing devices to deliver a refreshed customer experience, particularly within domains such as infotainment or advanced driver-assistance systems (ADAS). While this significantly reduces total cost of ownership (TCO), it inherently entails a less holistic SDV capability. Nevertheless, these manufacturers are also developing entirely new domain-centralized systems as an interim step towards full vehicle-centralized integration, aiming to address a looming limitation in feature scalability projected to become critical in the latter half of this decade.

Both approaches, despite their divergent starting points, present comparable challenges. OEMs adopting novel vehicle-centralized platforms and pursuing vertical integration through in-house development have frequently underestimated the requisite effort, specialized knowledge, inherent complexity, and the critical importance of economies of scale. Consequently, a new trend is emerging: the formation of partnership networks to pool resources and mitigate costs. Conversely, OEMs with

OEMs start to pursue a balanced “SDV evo” approach mixing elements of domain- and vehicle-centralized architectures to reach target cost and scalability goals.

established supplier networks are actively seeking solutions that offer sufficient scalability and will help them win on features and price.

Figure 1 illustrates the cost and performance dilemma that various OEM types are confronted by across their vehicle segments. Notably, OEMs aiming to cover a broad spectrum with a "one size fits all" platform encounter significant commercial challenges in downscaling and technical limitations in upscaling. Often, while software components may be readily reusable, the chosen hardware proves inefficient for downscaling, or vice versa.

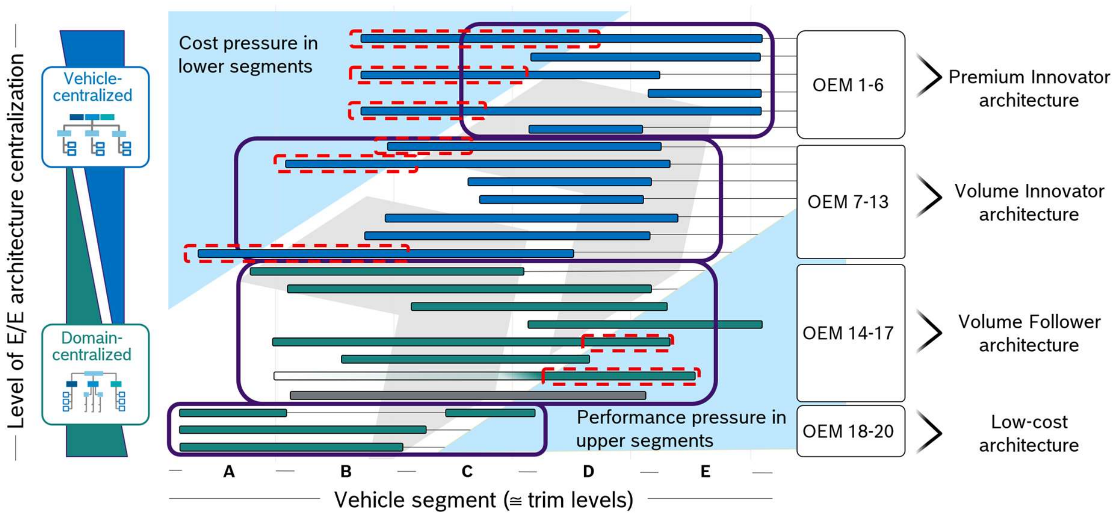

Figure 1 Architecture archetypes over vehicle segments (representing corresponding trim levels) indicating the critical ares with high cost pressure for lower vehicle segments; vehicle segment classification according to ISO 3833-1977

02
What's the way forward?

The predominant "one-size-fits-all" architectural paradigm of the beginning of this decade has proven inadequate for many large OEMs. This deficiency is rooted in an imbalance between hardware costs as the focal point for volume production and the R&D benefits that determine the time-to-market for the innovation driven premium segment. Consequently, the industry is transitioning to segment-specific vehicle architectures, typically comprising a cost-effective volume architecture and an innovation-driven premium architecture. Further architectural differentiation may be deployed for regional market penetration or brand strategic alignment. This strategic approach (depicted in Figure 2) yields considerable advantages:

(1) **Strategic Delineation:** Low-cost E/E architectures are protected from cost inflation driven by demands that are only relevant for the premium segment. Concurrently, innovation-centric architectures are unconstrained by cost optimization imperatives, ensuring uncompromised performance.
(2) **More capacity, faster time-to-market:** Commodity hardware and sub-systems for baseline architectures can be sourced off-the-shelf, significantly reducing development cycles, expenditures, and associated risks. This strategic outsourcing liberates internal resources for high-value, differentiating feature development.

The implementation of a dual E/E architecture strategy may initially appear cost-prohibitive. However, by establishing a common foundational layer and a unified governance model, significant overhead reduction and scalability gains are realized. These advantages are not exclusive to individual OEMs but present an industry-wide leverage opportunity. The strategic framework supporting this co-existence – SCALE, BRIDGE, REPLACE – is thoroughly presented in our 2025 ELIV publication[1]. Refer to this paper for an in-depth analysis.

Trend to split E/E architectures is clearly visible – a volume architecture and an innovative architecture.

Across architectures, many technical elements can be reused. Furthermore, certain elements provide scaling possibilities that can be leveraged if proper functional partitioning and technical abstraction are implemented. Required adaptations, in this context, result in lower costs compared to disconnected development. In the following we provide an application example from hardware and software within the high-performance compute area.

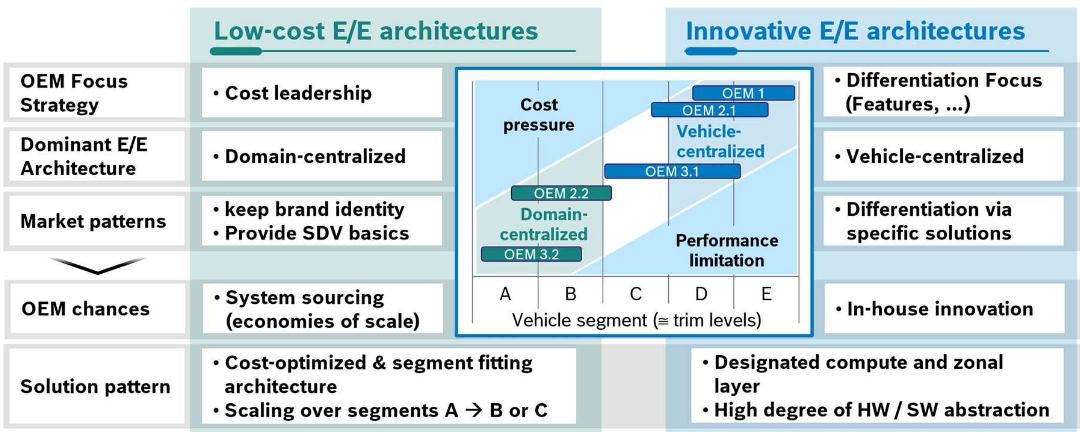

Figure 2 Market reaction on different demands across the vehicle segments shows a split of E/E architectures.

03
High-performance compute for the cost-efficient cross-segment SDV

In general, there are different ways to interweave E/E architectures on the high-performance compute level efficiently:

(1) Mechanical and software-based merge of products: Two become one.
Example: Separate infotainment and ADAS computers will be merged into one cross-domain computer hosting infotainment and ADAS. The technical solution is determined

[1] Th. Huck, A. Achtzehn: "Scale, bridge, replace: Designing E/E architectures in cross-OEM and cross-segment scenarios", 23rd Int. Congress and Exhibition on Electronics in Vehicles (ELIV), 2025

by the availability of technical fusion/separation solutions (from multi-SoC fusion to single-SoCs with hypervisors to multi-chiplet designs)

(2) Extendable products: Use of highly cost-optimized solutions with limited scalability augmented by support for extensions
Example: Additional low-cost extensions are used to increase the number of sensors/actuators in case the pin count on a device becomes prohibitive. The control remains at the base system (one brain).

How this interweaving looks in practice is depicted in Figure 3. Depending on the overall target segments the architectural split may be introduced at a different level. That why we are presenting our arguments with segment focal points.

A - B segment: Best value

This segment covers the fundamental needs - such as legislation-related features - and demands competitiveness via costs. Therefore, one central zone ECU is introduced hosting basic functions such as gateway and body controls, and all other devices will be kept separate to ensure the cheapest solutions (normally sourcing "off-the-shelf" solutions) available on the market (e.g., smart front camera for ADAS). Such systems are limited in terms of SDV capabilities, usually providing a good infotainment experience and partial updates using classical update solutions to avoid costly garage visits. For the end customer the SDV experience is mainly limited to digital content.

B - D segment: Balanced SDV-features vs. price

In this segment, the balance between providing SDV-related features on ECUs with advanced computational power and price is key.

For the central compute devices in the infotainment and ADAS domains, there are in general three options targeting different rationales:

For B-D segments the best choice is determined by technical and organizational guardrails. Each OEM will select their solution based on risk strategy, organizational transformation capability, and legacy.

Option (1) Having separate ECUs for ADAS and infotainment is the easiest way to fulfill the needs. It provides fit-to-purpose performance on one hand side, allows for re-use of legacy and, on the other side, avoids challenges in organizational collaboration. No compromises are necessary in terms of individual scaling.

Option (2) Placing several performance SoCs in one ECU (e.g., on one PCB or separate PCBs) allows for a good balance between piece-price reduction, R&D costs and need for organizational synchronization. Here, every SW domain can work separately on their SoCs while utilizing the cost benefits of sharing the ECU infrastructure (e.g., power supply, housing). In addition, combinations of different SoCs can be used to find a cost/value sweet spot for each ECU variant and costly and more challenging liquid cooling might be avoided for many variants.

Option (3) Deploying SW only on one SoC gives highest flexibility in terms of SDV ambitions. Furthermore, all cost advantages can be leveraged by sharing the SoC and ECU infrastructure across the respective SW domains. This requires superior Freedom-From-Interference (FFI) concepts to avoid unintended interference between and among the different SW domains.

These FFI concepts are available. However, they require a holistic joint design, release, and overall tighter organizational coupling.

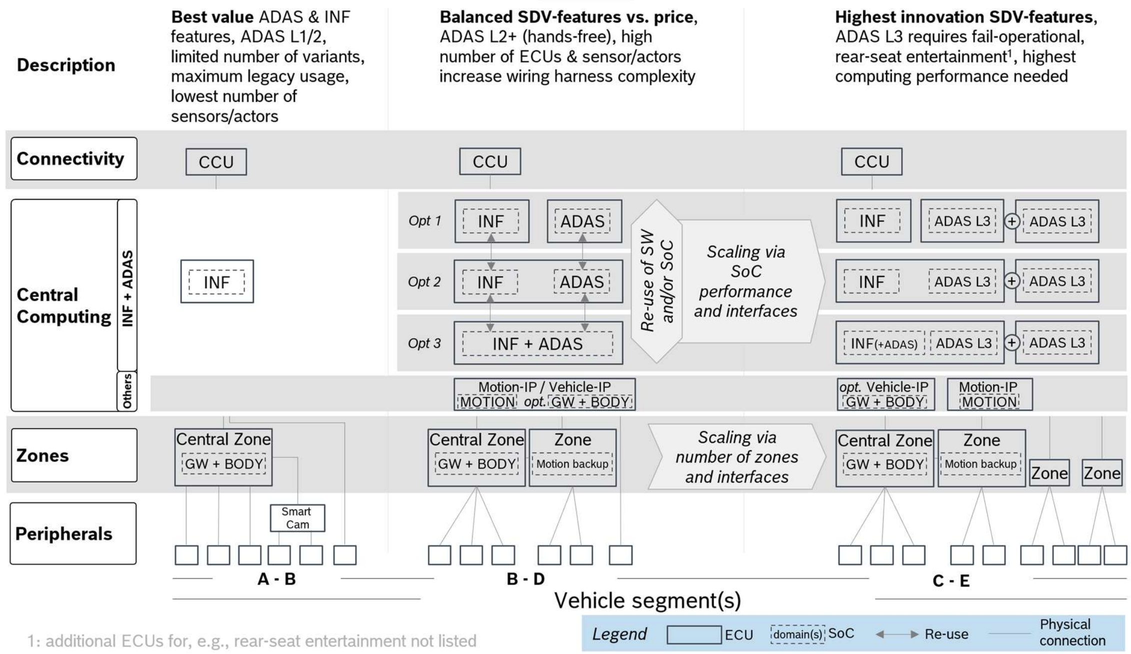

Figure 3 Intelligent scalability and re-use across different vehicle segments by combining domain-centralized and vehicle-centralized approaches
(CCU=Connectivity Control Unit; INF=Infotainment; ADAS=Advanced Driver Assistance Systems; SoC=System on Chip; GW=Gateway; IP = Integration Platform)

The decision between the different options for central compute on technical and commercial level depends on factors such as take rate, SoC strategies, SoC costs and number of variants. In some cases, best scalability can be achieved by combinations of these options: if option (3) is selected for segments B and C accounting for lower costs, options (1) or (2) can be used for segments D – E with higher performance demand while avoiding liquid cooling with the right SoC selections. The best setup for an OEM generally depends also on other driving factors such as organization and risk strategy.

C-E segment: Highest innovation SDV-features

Scaling up to ADAS SAE L3 demands redundancy with a physically separated ADAS computer (remark: one box hosting physically separated devices is possible, however, take rate and further common cause failure avoidance need to be considered). Customers can choose where the primary and secondary ADAS SW will be deployed.

In these segments, scaling via additional, separated ECUs is generally sensible, given the different ratio between hardware-related and R&D costs.

Scalability and reuse

Particularly from mid- to high segment, scaling can be realized by choosing best-fit SoCs from the same SoC-family to (1) keep the SW-architecture and development environment equal, (2) reduce segment-specific development and maintenance efforts and, thus, (3) limit adaptation costs. Decoupling in the low segment is sensible if take rates are sufficiently high and abstractions for consumer-facing functions are well established.

Figure 4 indicates that even if a mixture of options 1-3 will be taken to cover all needs, SoCs and SW can be carried over in most cases and mainly slight HW adaptations will be required. In addition, the central compute layer and setup allows the ADAS and infotainment ECUs to be used across different carlines and even architectures since in many cases only HW interfaces need to be adapted whereas the core development and SW architecture can be kept constant. By this, the OEM has the option to migrate with low risks and high re-use with their products into future E/E architectures.

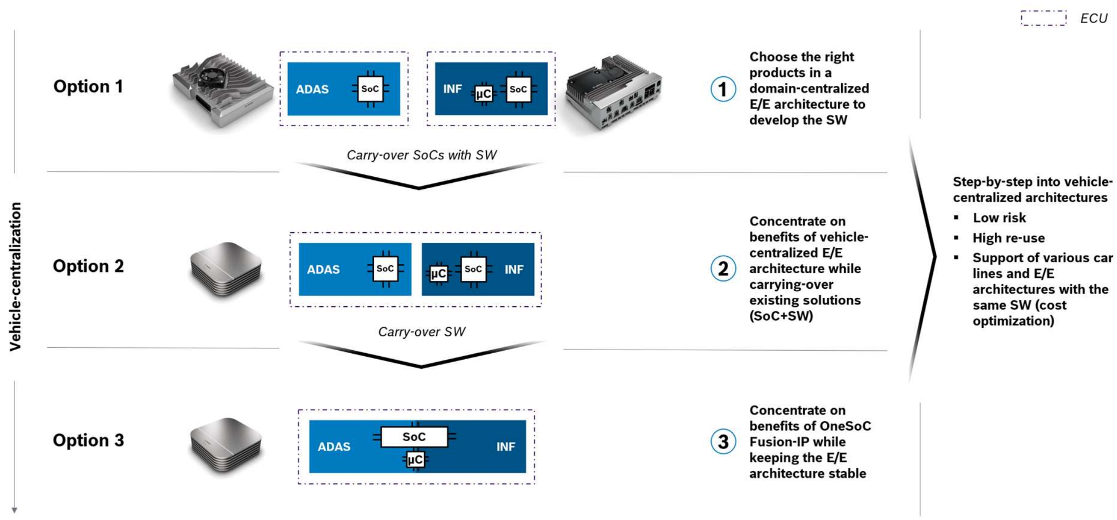

Figure 4 Migration path towards vehicle-centralized architectures with step-by-step integration of products and re-use. The options refer to Figure 3 in the central computing layer

04
Going further

This paper talks about how consumer and commercial objectives can be met when costs escalate prohibitively. It is particularly relevant for the pursuit of enabling a software-defined vehicle. Our discussion makes it evident that architectural development paradigms must be revised to concurrently address the needs of cost-sensitive and innovation-seeking consumers. This approach mandates strategic delineation through multi-architectural approaches, complemented by thorough consideration of how to scale, bridge, or replace technical solutions tailored for specific market segments.

Our analysis, exemplified by high-performance computing, substantiates the practical applicability of these principles. Their successful deployment, however, demands not only technical enablers and strategic pragmatism but also close collaboration within the automotive ecosystem. Critical partnerships with strong players like Bosch are becoming indispensable.

Bosch provides the right solutions with strong regional R&D presence to help our customers reach their SDV ambitions. We invite you to approach us to learn more about our comprehensive portfolio, e.g. in our high-performance compute area, and additional offerings.

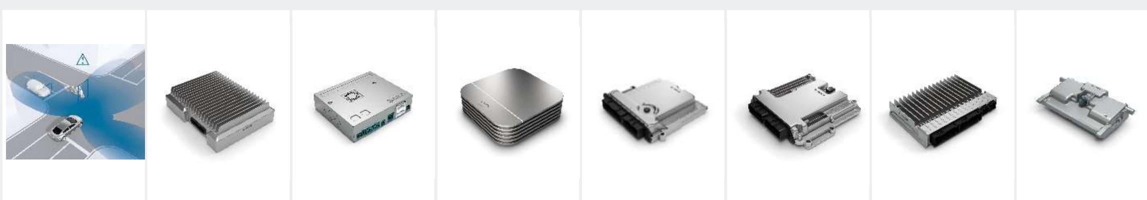

ADAS solutions	ADAS Integration Platform	Cockpit Integration Platform	Cockpit & ADAS Integration Platform	Zone ECU	Motion Integration Platform	Vehicle Integration Platform	Smart Camera

Together we can shape your current and future E/E architectures with smart solutions.

Contact us now

Get in touch with our expert team. We are looking forward to support you with our expertise to find the right solution for you!

Dr. Andreas Achtzehn

Cross-Domain Computing Solutions, Lead Architect

andreas.achtzehn@de.bosch.com
+49 (172) 378 6594

Learn more on our website
www.bosch-mobility.com

Bosch Mobility on LinkedIn
www.linkedin.com/showcase/bosch-mobility/

Contributing authors

Dr. Andreas Achtzehn, Erik Goerres, Andreas Deberling, Andreas Hoertling, Dr. David Weidenfelder

Imprint

Publisher:

Robert Bosch GmbH
Cross-Domain Computing Solutions

Poststraße 70
71229 Leonberg
Germany