McKinsey & Company

GSA

Getting ready for next-generation E/E architecture with zonal compute

The next generation of electrical/electronic architecture will play a vital role in the software-defined vehicles of the future—OEMs and chip makers need to be strategic in their choices.

This article is a collaborative effort between the Global Semiconductor Alliance, and Ondrej Burkacky, Johannes Deichmann, Martin Kellner, Fabian Steiner, and Julia Werra, representing view from McKinsey's Semiconductors Practice.

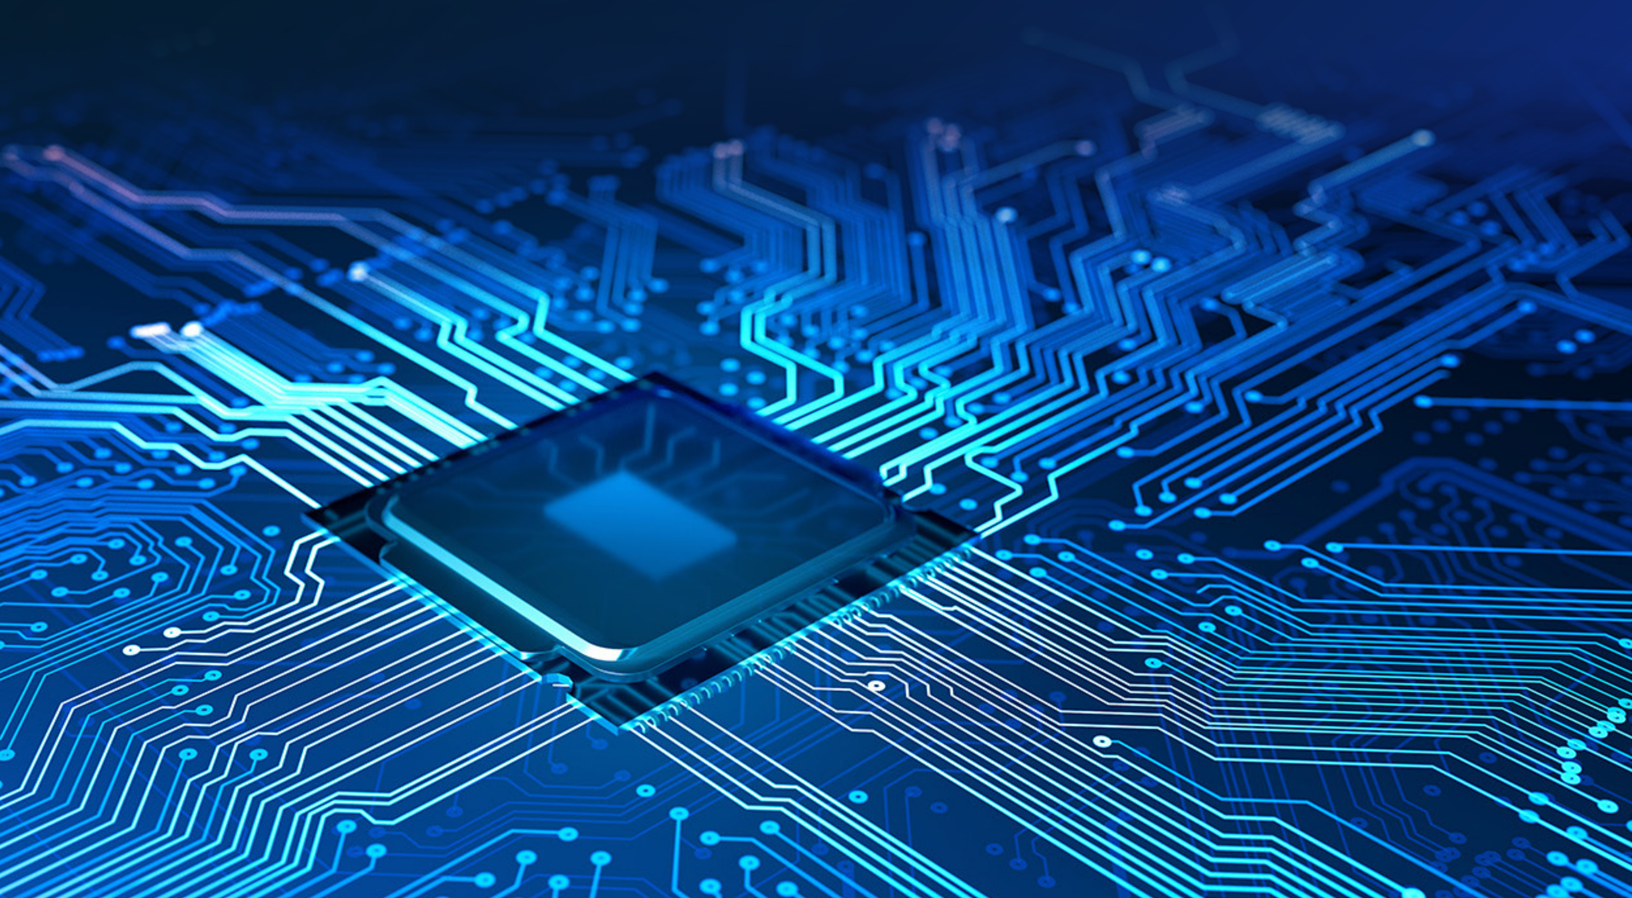

Differentiation in today's auto market is all about digital. Consumers value software-enabled features that reflect the wider "smartphonification" of passenger vehicles—including autonomous driving (AD), connectivity, electrification, and shared mobility (ACES). Meeting these needs means connected cars must continuously update software, engage with digital ecosystems, and provide instant access to onboard and offboard data. An enabler of these capabilities is electric/electronic (E/E) and software architecture, which is set to play a vital role in shaping the automotive and semiconductor value chain.

The stakes are high for both auto and chip makers. By 2030, the global automotive software and electronics market is expected to reach $460 billion, representing 6 percent growth (CAGR) from 2019 to 2030. Meanwhile, higher numbers of more complex features put heavy demands on systems, while software issues can cause launch delays or recalls. With 150 or more electronic control units (ECUs) in some premium vehicles, significant development, integration, and validation are required. In that context, E/E architecture is a critical variable that requires nuanced judgements over budgets and configuration.

Automotive E/E architectures have evolved from early generations based on decentralized models and single ECUs, to ever-increasing numbers of ECUs and closer integration. In the past few years, domain centralized architectures have emerged, through which a few domain computers replace several ECUs, enabling more complex cross-domain functions.

The most recent (fifth-generation) architecture introduces zonal controllers and central compute units (Exhibit 1). These have the potential to manage rising numbers of control units, as well as simplify software stacks and feature deployment. However, they also create challenges for car makers, including relating to latency and optimal controller topology, where there are significant design differences across the industry. In this article, we drill down into these competing dynamics and consider strategies for value chain participants as the market evolves.

Benefits of fifth-generation architectures

The single biggest design benefit of fifth-generation architectures is that they offer OEMs the ability to simplify the control unit landscape. Still, as OEMs transition, they need to gauge the potential benefits of different technical approaches, which can be broadly categorized under the following four banners:

1. **Facilitation of secure over-the-air (OTA) updates:** Large numbers of ECUs can cause update bottlenecks and complexity, leading to challenges around safety, reliability, regulatory compliance, and authenticity of individual images. The consolidation offered by zonal architecture is a significant step forward that will also facilitate rollbacks when updates fail.

Exhibit 1

Premium OEMs have adopted a variety of topologies for E/E architecture.

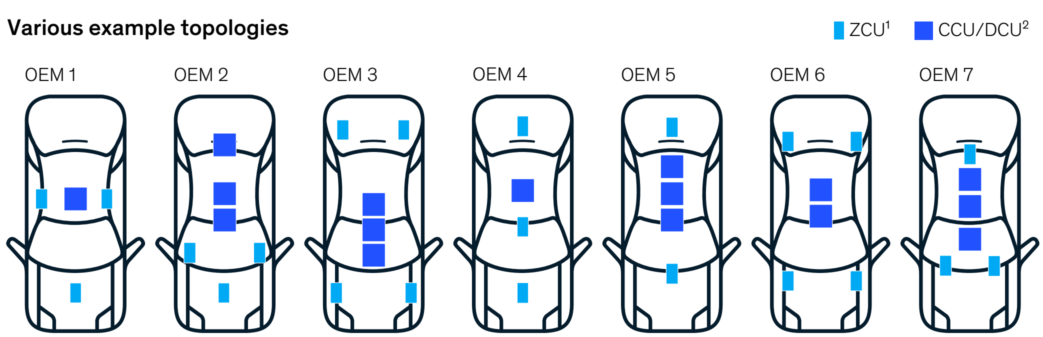

[1]Zonal compute units.
[2]Central compute units / domain compute units.

2. **Decoupling of hardware and software:** Decoupling offers the potential to speed up development times—enabling faster reengineering and adaption (reducing dependency on individual suppliers) and accelerating time to market, as well as boosting innovation flexibility. Decoupling can be achieved through abstraction layers in software stacks, so that one zonal controller implements the functionality of several, previously separate, ECUs. Abstraction layers can be constructed with less, and more commoditized, compute hardware.

3. **Increasing silicon consolidation and integration:** By adopting the functionality of several ECUs, zonal controllers achieve silicon consolidation and integration. Meanwhile, smaller node sizes boost power efficiency. Emerging system-on-chip designs bring together several central processing unit (CPU), memory, and dedicated hardware (HW) accelerator subsystems, the latter enabling deterministic routing latencies and freedom from inference. Unsurprisingly, this means that the latest system-on-a-chip (SoCs) for zonal controllers are based on node sizes of 16 nanometers (nm) and below.

4. **Reduction of wiring harness complexity:** Because zonal controllers serve as input/output (I/O) aggregators and are often placed at the mechanical rim of the car, they enable a less complex wiring harness, which in turn can foster standardization, support automation in the production process, and reduce costs due to lower employee skills needs. Given wiring harness costs in a modern vehicle often account for 20 percent of total E/E architecture budgets, this amounts to a significant benefit.

Zonal architecture to expand marketshare

Based on factors including announced E/E architecture launches, expert estimates, typical platform production times, advanced driver assistance systems (ADAS) adoption, and total cost of ownership benefits for different vehicles classes, we expect the global share of vehicles with zonal architecture to reach around 18 percent in 2030, while domain-based architectures (including light architectures with only one digital control unit (DCU) and full architectures, in which all domains compute have a DCU) will account for 48 percent and distributed architecture for 34 percent (Exhibit 2). The latter is due to the fact that we will still see vehicles produced on older architectures as well.

Exhibit 2

E/E architecture design distributions will vary over time.

Design distributions based on projected vehicle production volumes, %

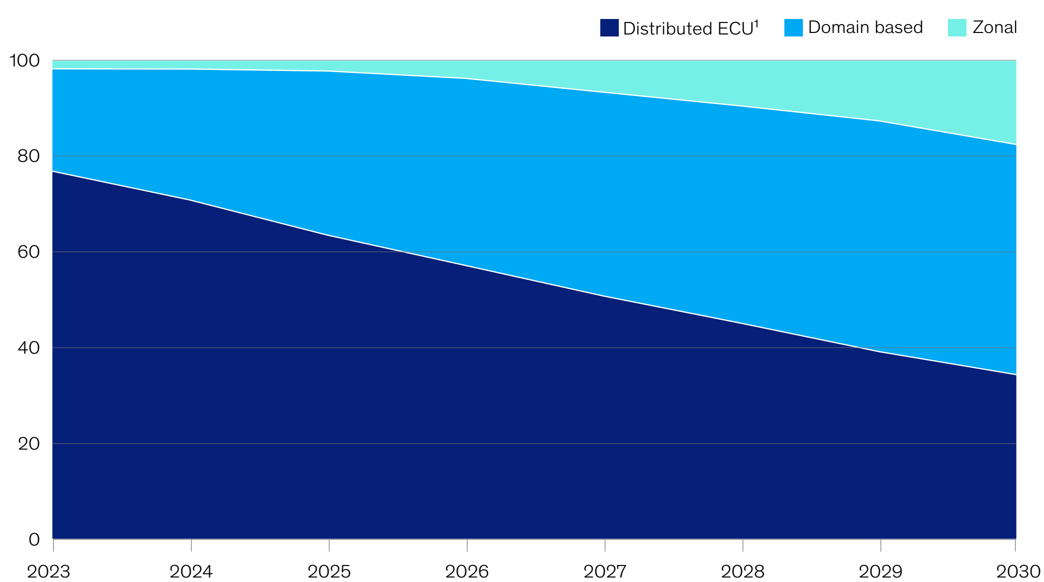

[1]Electronic compute unit.
Source: McKinsey Center for Future Mobility

The analysis carried out for this paper suggests the automotive compute unit market (distributed ECUs, DCUs, zonal compute units [ZCUs], and central compute units [CCUs]) will grow at a rate of around 6 percent annually between 2023 and 2030 (Exhibit 3). Because of consolidation in functionality, the ECU market will contract slightly (–1 percent annually) while the market for domain compute units (DCU), ZCU, and CCU will grow at around 30 to 40 percent per year. The significant market value for CCUs is the result of relatively high prices ($1,000 to $4,000 depending on the Society of Automotive Engineers' [SAE] level for the ADAS/AD part). By contrast, average selling prices for zone controllers are expected to price in the $50 to $70 range.

Controller design archetypes: What works best

When it comes to zonal controller design, value chain participants must consider a range of use case and application scenarios, identifying the most appropriate to match their propositions. There are four emerging archetypes, each of which offers a distinct set of benefits at different price points:

— **I/O aggregators:** In their simplest version, zone controllers act as I/O aggregators and gateways, passing on information from connected sensors and actuators to CPUs, and vice versa. They enable simplification of the wiring harness at the mechanical rim of the car. However, given the potential challenge of latencies in safety-critical applications, it is common to see networking functionality offloaded from the main CPU, and it becomes more important to implement routing and switching in dedicated hardware.

— **I/O aggregators with power control:** An enhancement on simple zone controllers is to combine I/O aggregator functionality with power control and distribution. Through smart fuses, this enables fine-grained power distribution to different parts of a vehicle—for example, enabling a tradeoff between safety and comfort requirements. Further, it permits predictive load balancing and comprehensive

Exhibit 3

The market for automotive compute units will see varying levels of value growth.

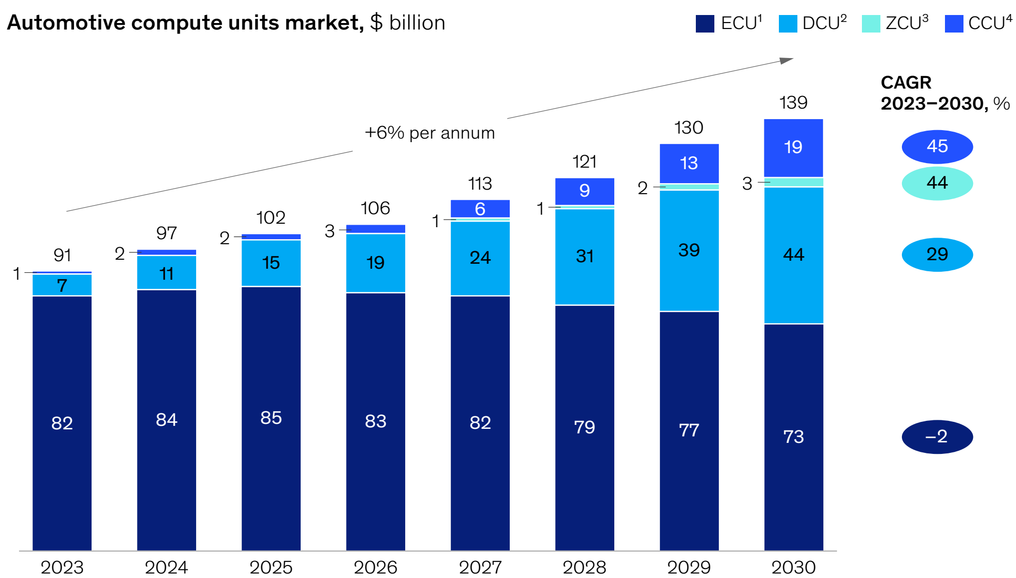

[1]Electronic compute units.
[2]Domain compute units.
[3]Zonal compute units.
[4]Central compute units.
Source: McKinsey Center for Future Mobility

circuit parameter monitoring. The approach provides a new way to control and measure power flow, and facilitates the powering on and off of specific features, rather than the whole vehicle. As a result, customers can expect longer ranges and improved data insights.

— **I/O aggregators with computing capabilities:** The zonal controller may provide computing capabilities, which are particularly relevant in the context of emerging satellite sensors—in which the (simple) physical sensor is separated from the compute unit that processes its raw signals. The separation means simpler sensors can be replaced relatively inexpensively if they fail or are damaged by impacts, for example, to bumpers. Meanwhile, the expensive processing unit is out of harm's way in the body interior.

— **I/O aggregators with computing capabilities and power control:** An advanced zonal controller would combine all the above-mentioned features and variations. This is the highest-cost option. Still, 53 percent of experts in a recent McKinsey survey said it is likely to emerge as the dominant model over time.[1]

Competitive dynamics along the value chain

As they consider how best to configure E/E architecture, value chain participants, including chip makers, tier-1 players, and OEMs need to make judgements across several dimensions. First, they must understand what domains will be most impacted, and what will change and not change in the move to zonal or central controllers. From a strategic perspective, decision makers must be clear on where they are best equipped to play. This will lay the groundwork for a plan on how to most productively engage.

The various vehicle domains have different requirements. The body, the powertrain, and the chassis are most suitable for zonal controllers, which consolidate functionality in areas that typically rely on a larger set of distributed ECUs, sensors, and actuators at the mechanical rim of a vehicle (Exhibit 4).

ADAS/AD and infotainment domains, by contrast, require significantly more compute power than zonal controllers can provide. Thus, these are likely to be implemented through a central or dedicated domain

Exhibit 4

Each automotive domain has distinct requirements.

Requirements by domain

	Zonal compute units			Domain/centralized compute units	
Domains	**Body**	**Powertrain**	**Chassis**	**ADAS/AD[1]**	**Infotainment**
Typical functionality	• Climate control • Interior/exterior lighting • Seat functions	• Motor control • Thermo management	• Steering • Suspension • Braking systems	• L2+: ACC,[2] AEB[3] • L3: ALKS[4]/traffic jam chauffeur	• Display control • Audio control • User applications
Compute requirements	Low-medium	Low-medium	Low-medium	High	High

High → Low
Integration of functionality into zonal controller

[1]Advanced driver assistance systems / autonomous driving.
[2]Adaptive cruise control.
[3]Automated emergency braking.
[4]Automated lane keeping system.

[1] McKinsey survey conducted as part of webinar, "Next generation of automotive E/E architecture: zonal computing," (November 2, 2022); n = 38.

compute unit. Connectivity also demands a separate unit, due mainly to cybersecurity requirements and its role in processing OTA updates.

As a result of the different requirements of each domain, zonal and central compute E/E architectures will need to be built on a range of software stacks (for example, AUTOSAR Classic, AUTOSAR Adaptive, or a tailored RTOS). Service-oriented architecture will play a key role, especially for zonal, domain, and central compute units, and will, for instance, support flexibility in software function deployment. As a result, the importance of automotive middleware solutions and players will rise, despite the fact that they are non-differentiating for customers.

As zonal and central compute units consolidate functionality across several domains, the transition to next-generation E/E architectures will likely accelerate the dissolution of functional siloes and boost research into end-to-end customer functionality. Tier-1 players and OEMs will need to collaborate more deeply, starting with common software architecture and middleware solutions, but also through joint development teams in agile settings to ensure frequent releases. This will be enabled by more stringent hardware and software separation and leveraging best practice in systems engineering. With cross-domain applications running on zonal controllers, OEMs will need to invest even more in software integration and validation.

Across the industry, we expect to see the formation of new ecosystems, alliances, and standards, especially for non-differentiating elements of the automotive hardware and software stack.

Meanwhile, semiconductor players will have their own priorities. Intellectual property (IP)/electronic design automation (EDA) and fabless operators may increasingly focus on scalable compute platforms that support different performance requirements. They may also think about new collaboration models to ensure easier integration of functional blocks—especially when chiplets become more common in the automotive segment toward the end of the decade. Chiplets separate the SoC into composite parts, creating the potential to build more powerful chips from smaller, standardized building blocks. This ensures higher yields and combat issues such as the reticle limit in extreme ultraviolet (EUV) lithography.

For IDMs or fabless players and wafer foundries, capacity reservations will become more common. These will increase transparency on planning horizons and strengthen supply chain resilience. We expect foundries to incrementally become more active in automotive, reflecting the sector's predicted 11 percent annual demand growth (CAGR) up to 2030. Meanwhile, OEMs will engage more actively with electronics manufacturing services (EMS) players and original design manufacturers (ODMs) as they diversify their supplier networks.

As OEMs turn increasingly to in-sourcing for E/E architecture, tier-1s will need to compete hard in architecture design and implementation. They can differentiate by building reference architectures and showcasing their capabilities. OEMs could also focus on customization of specific chips and accumulating know-how to enable tailoring, for example in the ADAS/AD domain. For IDMs this may create an opportunity for co-innovation and marketing of cutting-edge technologies. One impact will likely be more directed-buy arrangements, disrupting traditional value chain dynamics.

Across the board, software skill sets will be critical enablers. Amid intense competition, companies that take a strategic approach to hiring and retention will likely get ahead of their peers.

To meet rising customer demand for more digital automotive experiences, the automotive value chain must embrace vital enablers of transition. Zonal and centralized E/E architectures will play a central role, allowing for the simplification of control units, continuous updates, and through-the-lifecycle monetization. As part of this process, OEMs in particular need to balance business benefits against potential costs. For example, they must decide how much compute power is required now and in future, and strike a balance on trade-offs such as wake-up time and power management. As they ponder these variables, the most successful will prioritize a strategic approach, creating an optimal model for transition over time.

This work is independent, reflects the views of the authors (McKinsey & Company and GSA), and has not been commissioned by any business, government, or other institution.

Ondrej Burkacky is a senior partner in McKinsey's Munich office, where **Fabian Steiner** is a consultant and Martin Kellner is an associate partner; **Johannes Deichmann** is a partner in the Stuttgart office; and **Julia Werra** is an associate partner in the Berlin office.

The authors wish to thank members of the Global Semiconductor Alliance for their contributions to this article.