June 2026

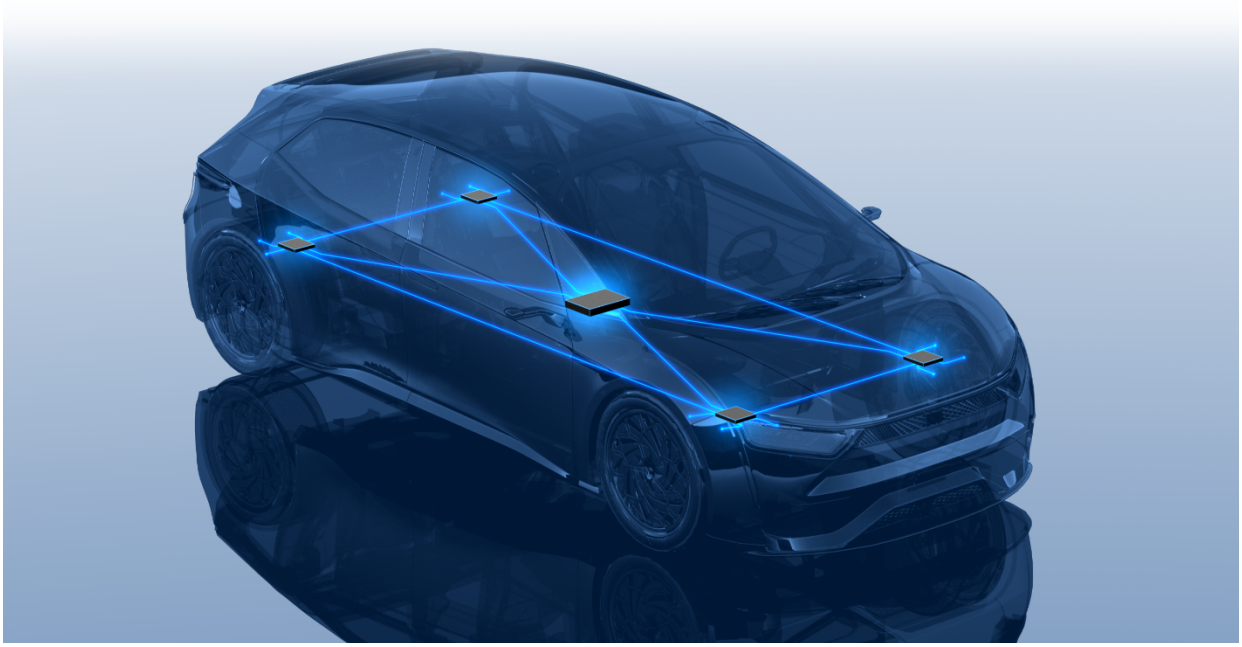

Architecting Centralized Automotive Compute Platforms

with PCIe® Switches and Centralized Storage

Russell Ruben, Director, Automotive, Industrial & IoT Marketing, Sandisk
Stephan Ahles, Marketing Manager, Socionext

Contents

1 Executive Summary

Automotive system architectures are rapidly evolving from distributed and domain-based designs toward centralized, high-performance compute platforms. This transition is driven by the increasing demand for advanced driver assistance systems (ADAS), immersive in-vehicle infotainment (IVI), connectivity, and data-intensive applications.

The combination of the Socionext SC1900 PCIe® switch and the SANDISK® iNAND® AT EN610 NVMe™ SSD forms a validated foundation for these next-generation platforms. Together, they deliver high bandwidth, low latency, and scalable resource sharing across multiple compute domains.

This paper examines the architectural principles behind centralized automotive compute, highlights the role of PCIe® as a high-speed interconnect, and outlines the storage performance requirements necessary to support modern vehicle workloads.

2 Introduction: Evolution Toward Centralized Automotive Compute

Automotive electronic/electrical (E/E) architectures have undergone significant transformation over the past decades. Early vehicle designs relied on distributed architectures, in which each function was implemented in a dedicated electronic control unit (ECU). While effective for isolated features, this approach tightly coupled hardware and software, making system updates or functional expansions after production extremely difficult. As feature content increased, so did system complexity, wiring effort, and overall cost.

To address these challenges, the industry transitioned to domain-based architectures, consolidating related functions - such as powertrain, ADAS, and infotainment - into dedicated domain controllers. This approach reduced ECU count and wiring complexity but still limited flexibility and cross-domain coordination.

The next step in this evolution is the emergence of zonal architectures, where vehicle electronics are organized by physical location rather than function. In this model, zonal controllers manage local sensors and actuators, while centralized high-performance compute (HPC) platforms execute compute-intensive workloads. This separation significantly reduces wiring complexity while enabling more flexible system design and facilitating over-the-air (OTA) updates.

3 The Role of PCIe® in Automotive HPC Platforms

As compute functions converge into centralized platforms, the underlying interconnect becomes critical to system performance and scalability. PCI Express® (PCIe®) has emerged as the preferred high-speed interface due to its bandwidth capabilities, low latency, and mature ecosystem.

3.1 PCIe® as the Backbone of Centralized Compute

In modern architectures, multiple workloads - including ADAS, IVI, connectivity, and body control - operate on shared compute resources. These systems require efficient, deterministic data exchange. PCIe addresses these requirements through:

- High bandwidth (up to 16 GT/s per lane in PCIe® Gen4[1]), enabling high-throughput data transport
- Low latency, essential for real-time and safety-critical applications
- Flexible lane configurations (x1 to x16) to support scalable system design
- A well-established ecosystem that reduces integration risk

These characteristics make PCIe® an ideal backbone for consolidating vehicle compute workloads.

3.2 PCIe® Switches in Multi-Processor Architectures

In complex automotive systems, a single root complex (bus master) is insufficient to support multiple processors and shared resources. PCIe® switches - such as the Socionext SC1900 - extend system capabilities by providing:

- Scalable connectivity: Enabling multiple SoCs (Systems-on-Chip) and accelerators to connect within a unified fabric
- Multi-host access: Allowing several processors to share resources such as GPUs or storage
- Peer-to-peer communication: Supporting direct data exchange between compute nodes without CPU intervention
- Flexible topologies: Supporting star, tree, or hybrid configurations
- Traffic management and isolation: Enabling quality of service (QoS) and functional safety mechanisms.

3.3 Deterministic Performance and Real-Time Behavior

Modern automotive systems demand predictable timing behavior, particularly for ADAS and autonomous driving applications. PCIe® switching fabrics support this through:

- Consistent and bounded latency
- Prioritization of safety-critical data traffic
- Isolation between compute domains

These features enable compliance with functional safety requirements and support deterministic system operation.

3.4 Enabling High-Performance Shared Storage

As automotive workloads become increasingly data-intensive, centralized storage architectures gain importance. PCIe® provides the bandwidth and scalability required for shared NVMe™ storage, offering:

- High-throughput access to large datasets such as maps and AI models

[1] See “PCI Express® Base Specification Revision 4.0 Version 1.0”, page 56

- Concurrent access from multiple compute nodes
- Reduced data duplication across the system.

3.5 PCIe® as a Long-Term Automotive Standard

PCIe® has been widely adopted across multiple industries - including data centers, mobile devices, and networking - for over two decades. Its forward-compatible roadmap helps ensure that it will continue to scale with future automotive performance requirements, making it a reliable long-term choice for vehicle architectures.

4 NVMe™ Storage in Modern Automotive Systems

Storage has always been a crucial part of automotive systems from digital audio and early navigation systems where application software and map data of a few gigabytes were stored to today's vehicles that have multiple storage devices with over 2TB of storage in high-end vehicles. The role of storage has changed from simply a read only use case to read/write intensive and to write intensive use cases. Not only have capacity points increased, and the use cases changed, but the speed required to move the data in and out of storage has also increased. SD, eMMC, UFS, and NVMe™ offer progressively high performance.

In addition, in order to optimize resources and reduce unnecessary duplication of data sets, the need for a shared storage solution has arisen. Shared storage where multiple systems or hosts can access the shared data while also maintaining their own storage space for data that should not be shared. This is similar to a drive on a computer where you partition it into A, B, C… drives and set permissions on each "drive" although it is still physically one SSD. One example of needing shared storage would be when the ADAS and the navigation systems both need to access the map data. Without a shared storage solution, there would need to be two separate storage devices each containing the same data sets.

5 Reference System Architecture

In this chapter, we introduce a central compute architecture based on PCIe® connectivity, using a switch and a central storage SSD.

In a basic high-performance computing (HPC) setup targeting Level 2 automation, two main architectural approaches are commonly used:

- Separated ADAS and IVI SoCs: Sensor data is processed by an ADAS SoC and often duplicated or forwarded to an IVI SoC for visualization purposes.
- Integrated ADAS/IVI SoC with external accelerators: A single SoC combines workloads, while discrete AI accelerators (e.g., NPUs) are attached via PCIe®.

While effective for initial implementations, these approaches present limitations, including limited scalability, inefficient data movement, and difficulty supporting redundancy.

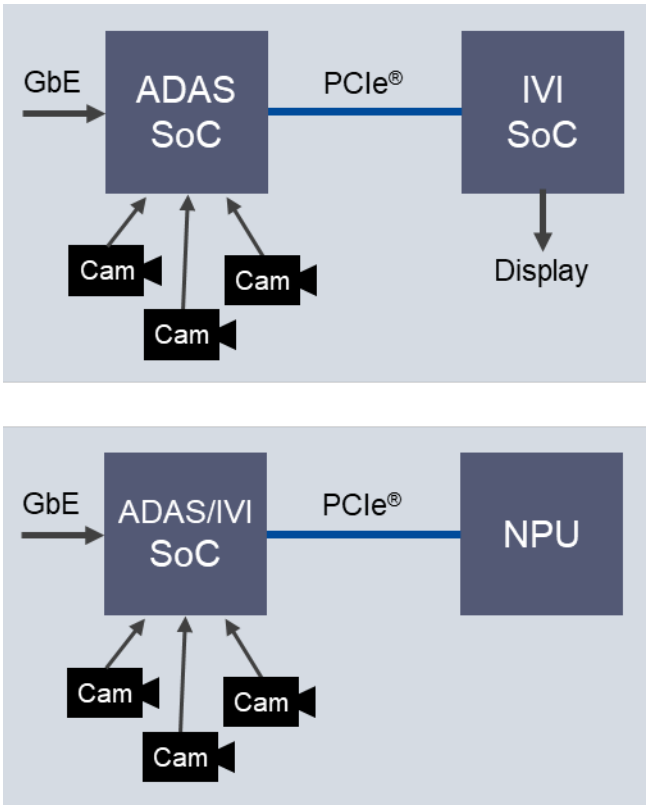

Figure 1: Simple HPC architecture

Level 3 architectures introduce redundancy and higher compute requirements, typically by adding a second ADAS SoC as a fallback system.

This creates several challenges:

- Sensor data distribution: Camera and sensor streams must be delivered simultaneously to multiple ADAS processors
- Indirect communication paths: Additional SoCs may not be able to communicate directly, requiring intermediate routing
- Increased workload on primary compute nodes: For example, one ADAS SoC may need to forward data to other processors
- Physical separation: Redundant compute units may be located in different parts of the vehicle, complicating connectivity

A PCIe® switch-centric architecture as shown in Figure 2, introduces a central high-speed switching fabric that interconnects all compute elements, sensors, and peripherals.

At the core of this approach is a multi-root capable PCIe® switch, which fundamentally changes how system components interact:

- Multiple SoCs can operate as Root Complexes (RCs) simultaneously
- Devices (endpoints) can be dynamically accessed by different compute domains
- Direct peer-to-peer communication between SoCs becomes possible

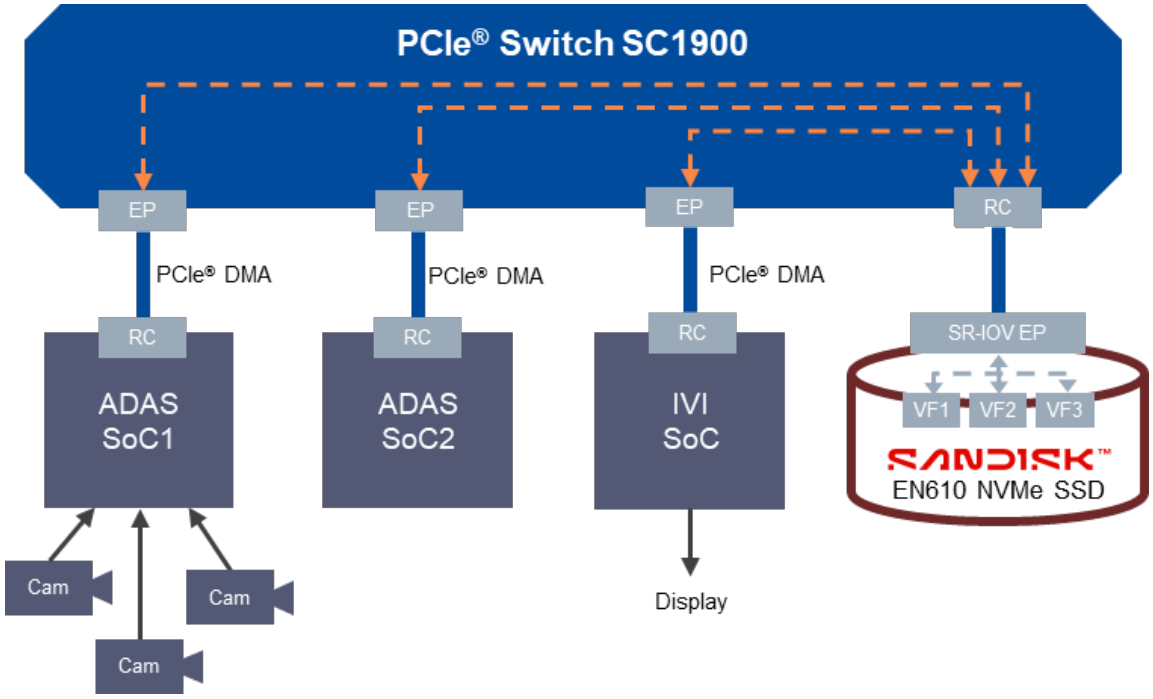

Figure 2: PCIe® -switch centric architecture

This architecture provides several key advantages:

- Full system scalability: Compute resources can be added or removed without redesigning the entire platform. This is very useful as the same architecture can be used for entry-level as well as premium vehicles. Also, software development is simplified as only one software version is required for all hardware configurations.
- The switch enables high-bandwidth, low-latency communication between all nodes.
- The distribution of data is partially done by the switch, reducing the load on the SoCs. The camera data arriving at SoC1 can be distributed to multiple other SoCs via the switch feature “multicast”. The SC1900 features multiple DMA controllers that are used for multicasting.
- The IVI and ADAS domains can be separated cleanly via the SC1900 switch. It is possible to configure communication from ADAS domain into IVI domain but not in the opposite direction. That would allow camera data to be distributed to IVI for display while helping ensure that the ADAS domain is sufficiently secured from unauthorized access via the IVI domain.
- Through a protocol called SR-IOV (Single Root IO virtualization), which is supported by the switch and SSD, a central storage element can be introduced to the system. The SSD can be virtualized by the hardware without using software solutions like a hypervisor.

Having a central storage has the following advantages:

- Cost savings through hardware consolidation
  - Using a single high-capacity NVMe™ device (instead of one dedicated per SoC) helps reduce BOM (bill of materials) costs, board space, power rails, cooling requirements, and overall system complexity. Fewer physical storage devices mean lower procurement, assembly, and maintenance expenses.
- Shared partitions
  - for read-mostly or collaboratively accessed data (e.g., maps, firmware updates, shared logs, machine learning models, or media assets across SoCs).
- Private/isolated partitions
  - per SoC - for sensitive or SoC-specific data (e.g., calibration data, security keys, personal user data, crash logs, or ECU-private databases). This provides logical isolation (via namespaces or access controls) while physically sharing expensive flash media, improving flash wear leveling and endurance across the entire capacity.
- Improved resource utilization and reduced stranded capacity
  - In multi-SoC designs, individual small drives per SoC often lead to under-utilized space or mismatched sizing. A single larger NVMe™ allows dynamic allocation and better overall flash usage.
- Simplified data sharing and reduced duplication
  - Common datasets no longer need replication across multiple local drives → saves storage space, bandwidth for syncing, and power.
- Lower power consumption and better thermal profile
  - Fewer storage controllers/chips → reduced overall power draw and heat generation.
- Easier over-the-air (OTA) updates and maintenance
  - Centralized storage makes firmware/OS image distribution and rollback simpler (one device to update vs. coordinating multiple).

6 Conclusion & Outlook

The transition to centralized and zonal compute architectures is redefining automotive system design. As workloads converge onto shared compute platforms, the underlying data infrastructure becomes a key enabler of performance, scalability, and safety.

The combination of the Socionext SC1900 PCIe® switch and the SANDISK® iNAND® AT EN610 NVMe™ SSD provides a robust solution for these emerging requirements. Together, they enable high-performance data exchange, efficient resource sharing, and scalable system architectures.

Looking ahead, increasing levels of automation, more advanced sensor systems, and growing AI workloads will further drive demand for bandwidth and deterministic latency. PCIe®, with its established ecosystem and clear roadmap, is well positioned to meet these future needs.

At the same time, centralized storage architectures will continue to gain traction as OEMs aim to reduce system complexity, improve efficiency, and support software-defined vehicle platforms.

Emerging technologies - including future PCIe® generations, CXL integration, and advanced storage virtualization - will further enhance these architectures. The solutions presented in this paper already align with these trends, offering a scalable and future-ready foundation for next-generation of automotive systems.

7 References

- https://cariad.technology/de/en/news/stories/enabling-centralized-architecture-with-pcie.html
- https://www.eletimes.ai/automotive-pcie-to-switch-or-not-to-switch
- https://www.cavliwireless.com/blog/not-mini/automotive-high-performance-computing-hpc-architecture