2026 NDIA MICHIGAN CHAPTER
GROUND VEHICLE SYSTEMS ENGINEERING
AND TECHNOLOGY SYMPOSIUM
MODULAR OPEN SYSTEMS APPROACH (MOSA) TECHNICAL SESSION
AUGUST 11-13, 2026 - NOVI, MICHIGAN

SOSA Enabled Automotive Technologies Disrupt the Battlefield

William J. Pilaud[1]

[1]LCR Embedded Audubon, PA

ABSTRACT

Modern Command, Control, Computers, Communications, Cyber, Intelligence, Surveillance, and Reconnaissance (C5ISR) operations hinge on rapid, resilient decision-making at the tactical edge. This white paper presents a modular Artificial Intelligence (AI) compute solution: integrating next generation automotive technologies into rugged liquid cooled SOSA/VITA systems. The approach delivers low-latency AI inference, multi-sensor fusion, rapid field serviceability (level-2 maintenance), safety certifiable systems, and a clear path for rapid technology insertion with commercial off the shelf (COTS) solutions specifically designed for military platforms. The automotive marketplace is investing tremendous resources into Autonomous vehicle technology. Autonomous vehicles utilize open-standard hardware and software, provide containerized deployment for software agility, offer vast amounts of simulation tools for rapid training of AI inference systems, and support zero-trust mission assurance. By combining SOSA based systems with autonomous vehicle technologies, C5ISR platforms maximize lethality for ground, air, and maritime platforms that minimize the time to theater, while controlling lifecycle costs.

1. Operational context and mission drivers

1.1. Battlefield realities

Operational tempo and decision latency

Modern engagements unfold at machine timescales, compressing the window between detect and act. Human-in-the-loop processes remain essential, but they must be supported by automated triage, tracking, and recommendations that operate locally when links are contested. Systems that wait on remote compute create avoidable delays and risk; edge inference and sensor fusion integrated on platform reduce this latency, which preserves initiative and maximizes lethality [1].

Contested spectrum and bandwidth scarcity

The electromagnetic spectrum is congested and actively contested, with deliberate jamming, deception, and opportunistic interference. High-fidelity ISR feeds saturate links quickly, and backhaul capacity rarely

aligns with peak sensor demand. Prioritization, compression, and on-node product generation are required so only decision-grade artifacts traverse the network [2].

Positioning, Navigation, and Timing (PNT) degradation and timing resilience

GPS and other global navigation satellite systems face spoofing, jamming, and localized denial. Without resilient time and position, track fusion drifts, sensor alignment degrades, and weapon-target pairing suffers. PNT alternatives, disciplined time distribution, and algorithms tolerant to drift are necessary to maintain coherence under attack.

Distributed operations under degraded communications

Small, dispersed formations and unmanned systems must operate with intermittent connectivity and variable latency. Mission threads cannot assume persistent reach-back; nodes should collaborate opportunistically, synchronize when able, and function autonomously when isolated. This demands data models, protocols, and workflows that degrade gracefully rather than stall.

Sensor proliferation and data deluge

Platforms now carry heterogeneous sensors—EO/IR, RADAR, Electrical Warfare (EW), acoustic, LIDAR—each with unique sampling, latency, and fidelity. The challenge is less about collecting data and more about turning it into timely, trusted actions. Standardized ingest, deterministic pipelines, and scalable fusion frameworks are essential to convert raw feeds into actionable tracks at the point of need.

Adversarial pressure: cyber, EW, and deception

Adversaries target software supply chains, inference pipelines, and control links with malware, adversarial inputs, and protocol manipulation. Systems must assume compromise attempts and enforce least privilege boundaries, measured boot, and continuous attestation. Robustness testing against noise, spoofing, and edge-case inputs is a first-order requirement, not a nice-to-have.

Survivability, attrition, and graceful degradation

Nodes will fail from damage, environment, or wear. Platform Architectures must isolate faults, contain faults, and re-route services without cascading collapse. Mission value comes from the aggregate: losing a node (platform, system chassis or plug in card (PIC)) should reduce capacity, not disable the force, or compromise the mission.

Environmental extremes and size, weight, power, and capability (SWaP-c) constraints

Dust, salt fog, shock, vibration, and temperature extremes define the operating envelope. Sealed enclosures and conduction paths protect electronics but constrain airflow; thermal density rises as AI workloads grow. Reliable performance hinges on efficient heat extraction and materials that maintain conductivity and mechanical integrity under stress.

Power budgets and energy discipline

Vehicle buses and battery packs impose strict limits on peak and sustained draw. Bursty AI loads can exceed margins without power governance, causing brownouts, throttling, or noise on shared rails. Power-aware scheduling, configurable thermal design power (TDP) profiles, and telemetry-driven management keep compute within safe, predictable bounds.

Maintenance realities at the edge

Forward teams need clear fault isolation, fast swap, with minimal tools and the right to repair. Mean-time-to-repair (MTTR) should take minutes, not hours, and software reconstitution must be deterministic, in theater. Spares logistics, serialized configuration, and automated provisioning reduce downtime and cognitive load in high-pressure conditions. C5ISR systems nodes should be inter-changeable at the PIC level to

maximize system mission time and allow on the ground repair during mission time. Systems that need whole replacement at depot or maintenance facility create supply-chain expense and slow down deployments and mission readiness.

Interoperability and open systems

Mixed-vendor fleets and coalition operations require predictable electrical, mechanical, and software interfaces. Proprietary lock-in at system/platform level slows upgrades and increases lifecycle risk. Open standards for slots, fabrics, backplanes, and data models enable competitive sourcing, faster insertion, and reuse across platforms and missions.

Human-machine teaming and cognitive load

Operators face saturation from alerts, feeds, and chat channels. Automation should reduce—not add to—cognitive burden, surfacing prioritization. Explainable cues with confidence measures and provenance are essential to mission success. Trustworthy autonomy emerges from consistent behavior, transparent boundaries, and rapid feedback loops between users and systems. Different security and education levels in platform limit access to critical threat information, makes the war fighter blind to threats and are unable to act. Simplifying sensor information to easy-to-understand actionable cues sent to human machine interfaces vastly improves survivability and increase lethality.

Data stewardship and sovereignty

Classified, sensitive, and coalition data each carry unique handling constraints. Edge nodes must enforce policy at ingest, apply redaction and encryption automatically, and retain auditable trails. When links are scarce, nodes triage what leaves the platform, preserving sovereignty while still contributing meaningfully to the common operational picture.

Technology change and acceleration of new threats

Fundamentally, the technology change at battlespace is changing faster than the ability of current defense contractors to define, design, change, test and deploy. Design cycle speed is under extreme pressure.

2. Understanding HPEC Computing: The Ten Axioms

High Performance Embedded Computing (HPEC) is revolutionizing compute-intensive applications in defense and aerospace by merging advanced processing capabilities from High Performance Computing (HPC), embedded systems, and sensor integration. This transformation allows for superior signal, radar, and image processing that meets the rigorous standards of military environments. HPEC systems draw upon established technologies from commercial sectors, adapting them to maintain operational effectiveness in challenging contexts. Traditional HPEC systems leverage trends and technologies from several overlapping markets: Commercial HPC, embedded systems, cellphone technologies, and industrial sensors. By doing so, it capitalizes on the advancements from these sectors, which collectively get integrated into the design and functionality of HPEC systems. The applicability of technologies is guided by the **Ten Axioms of HPEC**, which delineate the critical characteristics that these systems must embody to effectively address the needs of demanding military applications.

2.1. The Ten Axioms of HPEC

1. **Maximized Floating Point Performance per Square Inch** (HPC): HPEC systems necessitate exceptional processing power, reflected in their ability to achieve high volumes of FLOPS (floating-point operations per second).
2. **Designed with Open Standards** (HPC): To ensure interoperability and lower costs, HPEC systems

leverage open standards in their architecture and design, drawing from the broader HPC landscape.

3. **Constrained by SWAP** (HPC): HPEC systems must balance Size, Weight, and Power (SWAP) to fit within military deployment constraints.
4. **Operated in Harsh Environments** (Embedded Systems): HPEC systems must function reliably outside controlled office or laboratory settings, enduring a range of extreme conditions.
5. **Secure** (Embedded Systems): Security is paramount in HPEC systems for protecting sensitive data from infiltration and espionage.
6. **Designed with Multiple Processors** (Embedded Systems): For enhanced computing capabilities, HPEC systems rely on a multi-processor architecture interconnected by optimized communication fabrics.
7. **High-speed Low-latency Fabric** (Embedded Systems): A requirement for high-speed, low-latency connections between processors and sensors to meet processing and responsiveness demands.
8. **Minimized Latency** (Sensors): Latency must be at a minimum to ensure real-time responsiveness which is necessary for military applications.
9. **Integrated with FPGAs** (Sensors): Sensors need Field Programmable Gate Arrays (FPGAs) for their adaptability and high bandwidth capabilities in processing sensor data.
10. **Co-located with Sensors** (Sensors): HPEC systems require processors to be situated close to the sensors they support to minimize data transfer delays.

2.2. Analyzing the HPEC Framework

The relationship between HPEC systems and their underlying technologies—HPC, Embedded Systems, and Sensors—illustrates a comprehensive approach to solving complex military challenges. The primary focus on maximizing floating-point performance and leveraging open standards allow HPEC to utilize commercial technology while enhancing its operational robustness through SWAP considerations.

HPEC systems selection depends on Floating Point Performance and adherence to Open Standards

The performance benchmarks in HPC heavily influence HPEC, epitomized by the usage of x86 architecture processors and advanced math accelerators such as GPGPUs. HPEC inherits these capabilities, addressing both SWAP constraints and a growing need for power optimization. Moreover, as the HPC landscape transitions from proprietary to open standards, HPEC systems benefit, gaining access to a broad array of tested software solutions and frameworks that streamline system development and integration from the commercial industry.

Operational Integrity in Contested Environments

Embedded systems tailored for rugged environments are inherently critical to HPEC's efficacy. The increasing threat of information-based warfare heightens the focus on security features, making it essential for HPEC systems to adopt robust protection measures across various operational layers—

from physical housing to individual components.

Processor Integration and Latency Management

Embedding a variety of processors interconnected via low-latency fabrics allows HPEC systems to achieve high levels of processing efficiency. Military response time depends on sensor processing time and the data transport time to the processing PICs; low latency fabrics performance is part of this do act loop. The integration of FPGAs plays a pivotal role in managing data throughput, acting as the interface for converting, decimating, and compressing data from sensors into actionable information.

Future Directions for HPEC

HPEC systems are positioned to evolve alongside advancements in sensor and network technologies, especially given the ongoing surge in data generation from high-bandwidth sensors. As network capabilities grow but often fail to keep pace with these data rates, HPEC's unique architecture may become indispensable for military applications that require near-instantaneous processing and response. In summary, understanding the Ten Axioms of HPEC not only highlights the critical elements of the driving system design but also equips developers and stakeholders with a framework for future innovation in high-performance, rugged computing applications. As the demand for processing power and security in deployed environments continues to escalate, the HPEC sector will become increasingly vital in shaping the next generation of defense and aerospace technologies.

2.3. Traditional technology driving the defense industry

Traditional HPEC technology leverages HPC, laptop, and cellphone industries. In the beginning, HPEC used specialized silicon chips like the TI DSP and customized ASICs. As FPGA technology improved these sub-systems slowly integrated into the fabric as I/O pre-processing and gradually improved to take more of the signal decimation and filtering. The disruption was the Altivec processor from the PowerPC AIM alliance – this processor capability embedded DSP capability in a general-purpose processor. This disruption allowed the HPEC industry to leverage off the shelf software and leverage the embedded SWAP optimized processor like PowerPC. The market leveraged this disruption for many years, until the Intel laptop processors gained processing capability at a better SWAP than PowerPC. Again, the Intel, ethernet switch, and FPGA era of HPEC ruled for many years. Shortly, after Mercury Systems developed "closer to the metal" software that leveraged GPGPUs from Nvidia and ATI (now AMD), these systems gradually, became an essential component of HPEC systems. The GPGPU opened AI enabled capabilities into C5ISR into a harsh environment.

3. HPEC examples

The following are real world examples of an HPEC sub-system for an airborne C5ISR application. This application is a typical example of a 10 slot PIC VITA 46 VPX backplane with a SOSA defined 100 gigabit per second Ethernet (GbE) central switch connected data plane. The first two VPX modules are power supplies, the second two are SOSA defined I/O single board computers (SBC) with solid state data storage (SSD) modules. Slots 5-9 are the traditional SOSA defined compute intensive profile slots that would host digital signal processing boards in some combination of Intel based SBC, FPGA or GPGPU resource (figure 1).

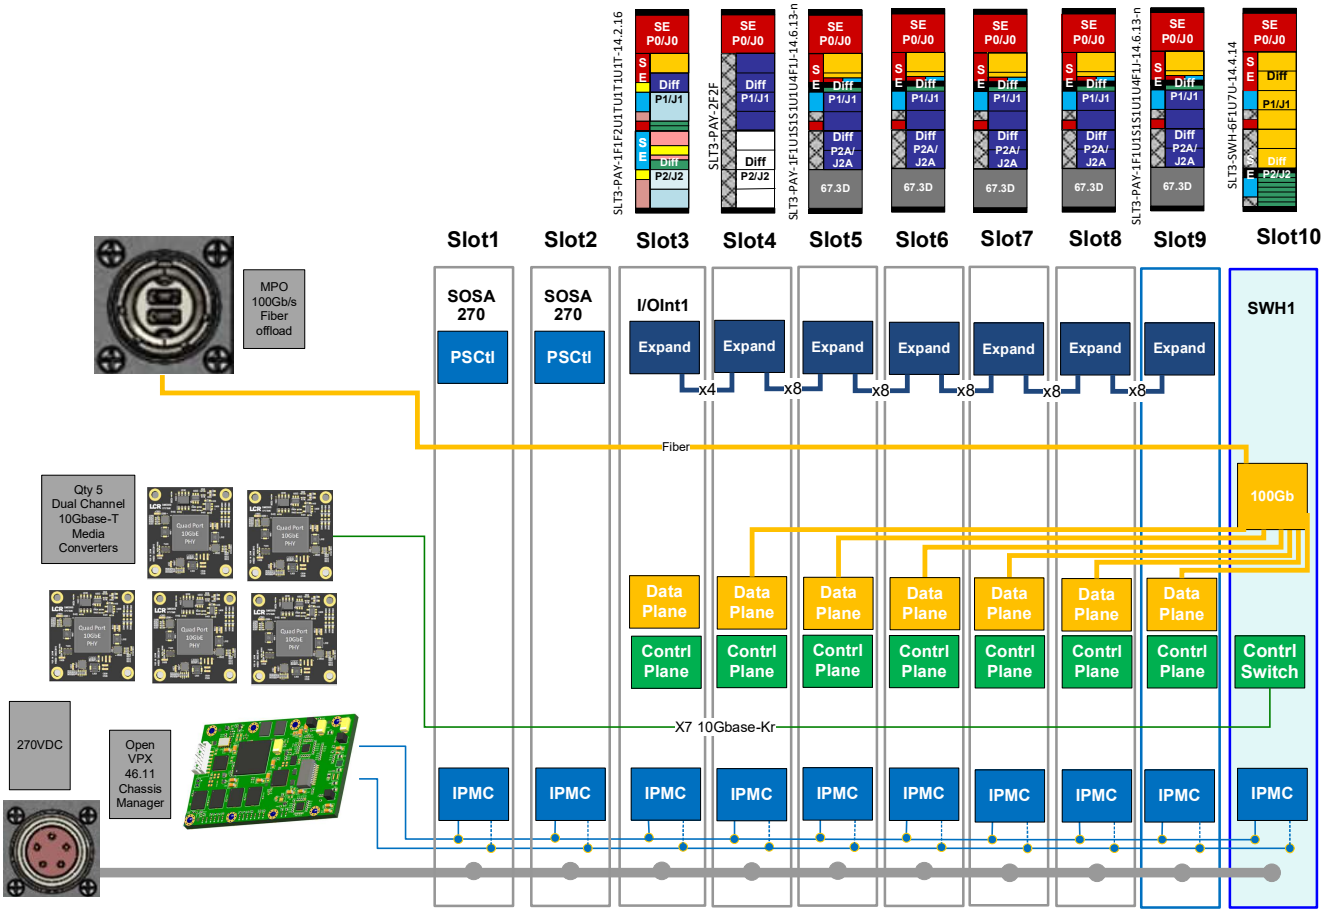

Figure 1: HPEC 10 VPX Slot SOSA aligned system architecture

Traditional HPEC systems would integrate an SBC with a GPGPU card combination connected via the Expansion plane with PCI-Express fabric. The SBC or GPGPU would connect the cluster of these cards together with a 100GbE NIC. The maximum performance of this system in terms of floating point per second (FLOPS) is 7.5 TeraFLOPS (TFLOPS) for NVIDIA Turing type GPGPUs on systems built last year and 41 TFLOPS for Ada Lovelace for designs built this year. In the future, when the Blackwell systems are available the cluster will improve to 160 TFLOPS. The Intel processor is not part of the performance equation because the GPGPU has much more floating-point processing power. This system architecture example demonstrates that system performance improves by replacing PICs without changing backplane, I/O, or enclosure. Improvement by sparing is possible by SOSA based PIC selection. The backplane shown above hosts three clusters of SBC/GPGPU sets which results in an impressive 480 TFLOPS.

SOSA based GPGPU PICs are not limited to a one-to-one SBC/GPGPU combination. Some C5ISR applications scale performance with the number of GPGPUS in a system. By segmenting the expansion bus or daisy chaining connections, a GPGPU cluster can remove two SBCs and improve the HPEC system to 900 TFLOPs of performance.

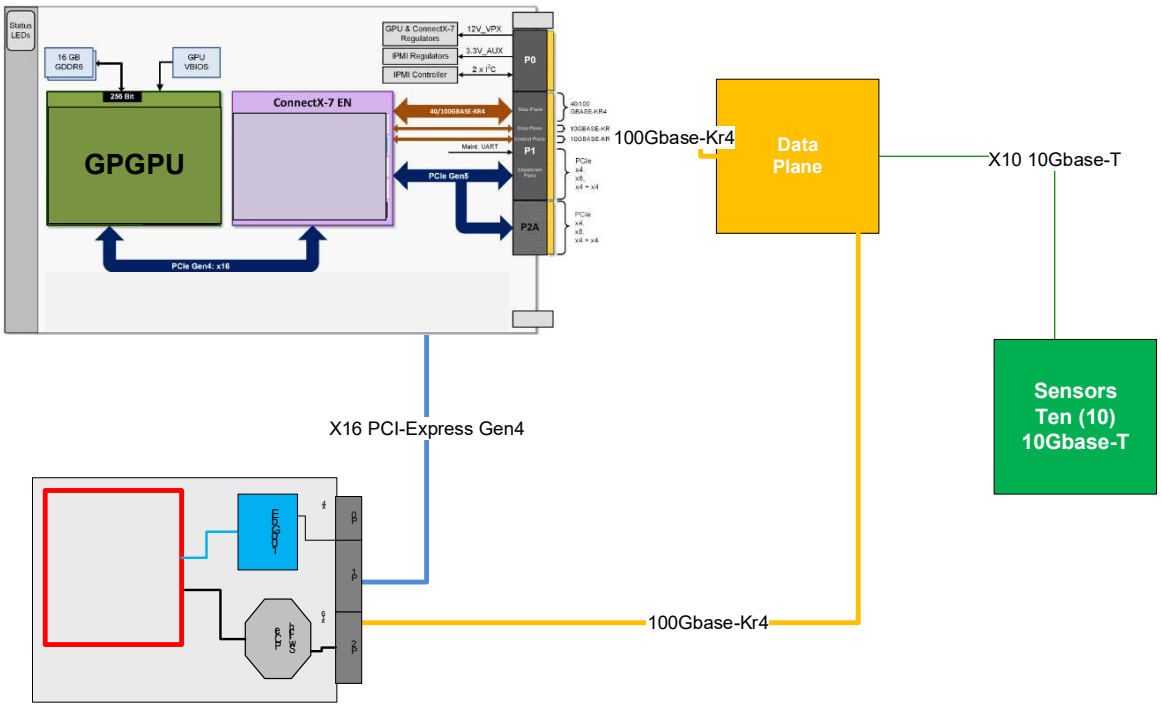

Figure 2: SBC/GPGPU VPX Cluster

Again, backplane, chassis and I/O are maintained, and this is simply a function of PIC selection. The power of HPEC systems that scale as technology improves is a powerful characteristic and has a side benefit of obsolescence prevention

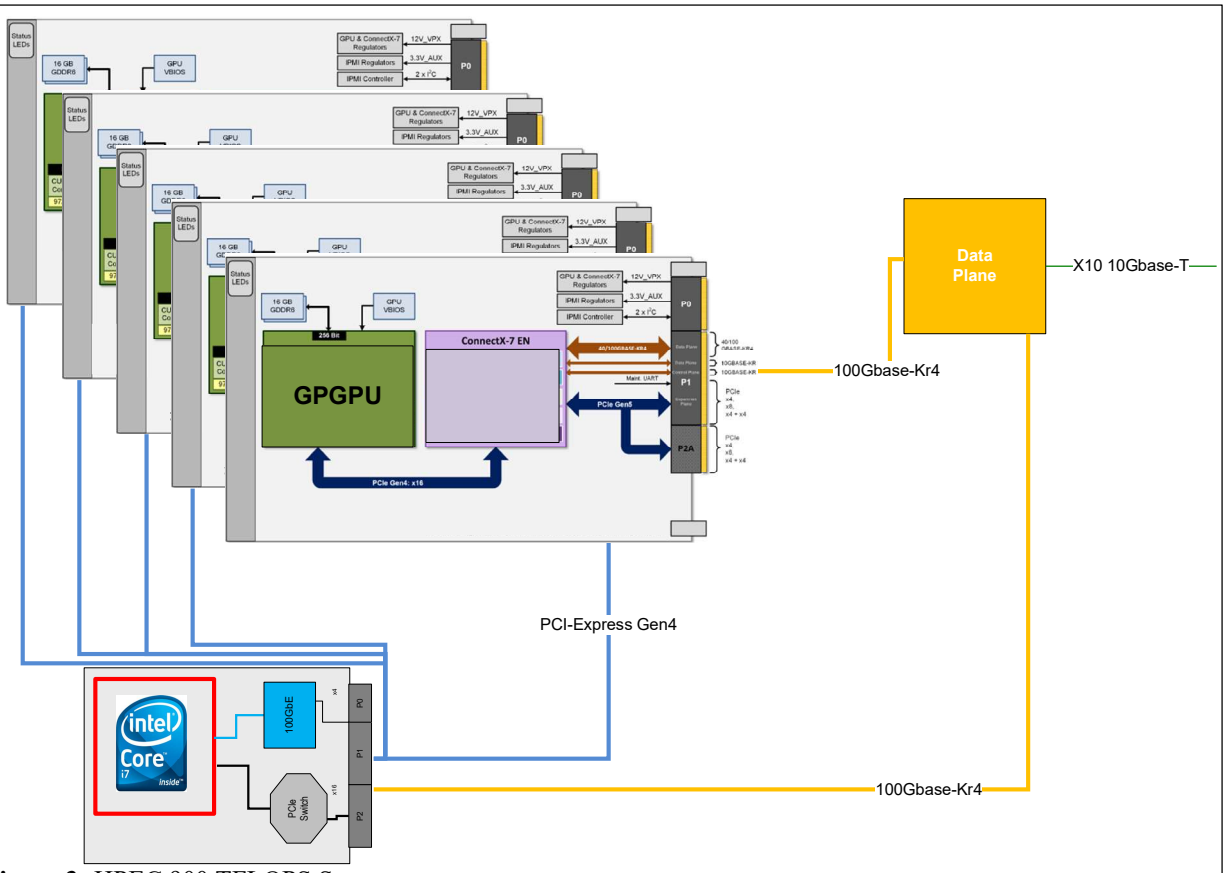

Figure 3: HPEC 900 TFLOPS System

.

4. Technological disruption from the automotive industry

The $4.3 trillion automotive market is in the middle of a once-in-a-century overhaul: electrification is scaling, autonomy keeps advancing despite setbacks, and vehicles are becoming software-defined, hyperconnected platforms. Executives now rank auto as the most disrupted major industry, driven by uneven electric vehicle (EV) adoption, reconfigured supply chains, and talent gaps at the intersection of software and hardware. These shifts are rewriting business models, compressing product cycles, and pushing intelligence to the edge, where rapid decisions and resilient systems matter most [3].

4.1. Vectors of change

- **Electrification at scale:** Consumers are pivoting toward EVs, with strong momentum across China and Europe; sustained battery supply, industrialization, and decarbonization are now strategic imperatives for incumbents and challengers alike [4].
- **Autonomy and advanced driver assistance:** While some players have retrenched, robo-taxis and robo-shuttles are expected to scale over the next one to two years, keeping autonomy squarely in the opportunity set even as liability and regulation evolve [3].
- **Software-defined vehicles and new revenue pools:** Connectivity, secure over-the-air updates, and feature unlocks expand revenue potential beyond unit sales; data and on-demand mobility services could lift total automotive revenue pools by roughly 30%, adding up to about $1.5 trillion [5][6].
- **Supply chains, tariffs, and localization:** Tariffs and geopolitical shifts are accelerating supply-chain rewiring and local investment, with executives citing materials costs and supplier visibility among top risks; firms are reshoring and pursuing transformational M&A to harness disruption [7].
- **AI-first engineering and cybersecurity:** Leaders see AI and software-defined functionality as revenue generators but also acknowledge compliance and overreliance risks; in parallel, cyber threats and supply-chain integrity demand continuous hardening of vehicle and enterprise software stacks [8].

4.2. Jetson Thor capabilities

NVIDIA Jetson Thor is purpose-built for this new edge reality, marrying datacenter-class AI throughput with sensor I/O and robotics software so decisions happen on-platform, not in the cloud. The module delivers up to 2070 FP4 TFLOPS with the Blackwell GPU, 128 GB LPDDR5X, and configurable power from 40 W to 130 W—over 7.5× the AI compute of AGX Orin with 3.5× better energy efficiency. It is designed for physical AI workloads—perception, planning, multimodal generative AI—where

tight control loops, multi-sensor ingest, and real-time streaming are nonnegotiable [8].

- **Compute and acceleration:** Blackwell-architecture GPU with fifth-gen Tensor Cores, Multi-Instance GPU (MIG) for partitioned workloads, and a 14-core Arm Neoverse V3AE CPU up to 2.6 GHz provide headroom for large transformers and classic robotics pipelines alike; Thor reaches up to 2070 TFLOPS (FP4 sparse) and about 1035 TFLOPS (FP8 dense) for edge-scale generative AI [9].
- **Memory and bandwidth:** 128 GB LPDDR5X on a 256-bit interface delivers 273 GB/s, supporting high-rate sensor fusion, map handling, and video analytics without starving accelerators.
- **Sensor ingest and networking:** Four 25 GbE links, a camera offload engine, and support for the Holoscan Sensor Bridge which enable high-bandwidth, low-latency capture and synchronization across many cameras and sensors, including raw frame transport over Ethernet for deterministic pipelines.
- **Vision and media engines:** Dual hardware encode/decode blocks support up to 4× 8Kp30 decode and 6× 4Kp60 encode across modern codecs, enabling on-node video summarization, compression, and dissemination without round-tripping to the cloud [9].
- **Power envelope and thermals:** Configurable 40–130 W operation balances SWaP constraints with sustained performance. Thor's efficiency improvements widen the viable envelope for dense on-edge AI under harsh conditions [9].
- **Software and security stack:** The platform integrates with NVIDIA Isaac for robotics, GR00T foundational models for humanoids, and provides end-to-end safety and security across models and edge-to-cloud pipelines to support trustworthy autonomy and managed updates.

Why does Jetson Thor matter for the defense industry?

Automotive disruption is normalizing an edge-first architecture: vehicles perceive, decide, and act locally, then share decision-grade products upstream. Jetson Thor aligns with that pattern. Its multi-25 GbE ingest, MIG partitioning, and media engines let you run perception, mapping, and VLM/Vision-LMM summarization concurrently, while preserving determinism for control loops. Its software stack shortens time-to-first autonomy and supports iterative feature delivery without forklift upgrades—exactly what software-defined fleets require in contested bandwidth and evolving regulatory environments.

For cross-domain adopters, Thor offers a practical path to automotive-grade patterns—high-throughput, power-governed AI at the edge; modular sensor pipelines with Ethernet-native capture; and continuous, secure updates. That combination turns

disruption into leverage: rapid insertion of new capabilities, resilience when links degrade, and lifecycle agility that keeps platforms relevant in fast-moving operational theaters [4].

4.3. Holoscan Sensor Bridge

The **NVIDIA Holoscan Sensor Bridge** is a hardware and software framework designed to connect heterogeneous, distributed sensors to high-performance AI computing platforms with ultra-low latency. Serving as a crucial data ingest layer within the broader NVIDIA Holoscan edge AI ecosystem, it enables a **"Bring Your Own Sensor" (BYOS)** architecture over Ethernet [10].

Key Technical Pillars

- **FPGA-Based Data Ingestion:** At the edge, peripheral sensor data (e.g., MIPI CSI-2 camera feeds, ultrasound, or microphones) is captured by a flexible, field-programmable gate array (FPGA) device interface. This FPGA packetizes raw data and streams it via UDP over standard Ethernet.
- **Direct GPU Memory Access:** For hosts equipped with accelerated RDMA-capable network interface cards (NICs), the bridge leverages technologies like GPUDirect RDMA and RoCE (RDMA over Converged Ethernet). This bypasses the host CPU entirely, allowing the Ethernet packets to be written directly into GPU memory, which slashes processing bottlenecks and latencies by up to 90%.
- **Host Agnostic Synchronization:** It provides hardware-based **Precision Time Protocol (PTP)** timestamping. This guarantees highly accurate synchronization between multi-sensor arrays and logs precise ingestion times to measure entire pipeline latency.
- **Modular Software Integration:** The host software component delivers custom operators (such as RoceReceiverOp) for Holoscan SDK. This enables seamless data injection directly into real-time C++ or Python AI pipelines for tasks like inference, medical imaging, and automated inspection.

The platform natively supports compute architectures like NVIDIA IGX, Jetson AGX Orin, Jetson Thor, and DGX Spark, and is widely deployed in latency-critical fields including robotic surgery, autonomous vehicles, and industrial automation. [10]

5. How this automotive technology disrupts C5ISR

Operational tempo and decision latency is increasing at machine speeds therefore, the battlespace strategy is to put as much act, decide do loop capability as close to conflict as possible. The autonomous automotive market is adding this capability into next generation silicon, and the military marketplace needs to leverage this capability to improve mission success. Fundamentally, these two markets align in technological need.

Neither the automotive marketplace nor the battlespace can depend on wireless connection to data centers to complete the mission or transport people because of contested spectrum or bandwidth availability or scarcity. This capability to connect and

adapt to new environments is essential to both systems, but the system must have some resilience to low or limited communication environments. This requirement also applies to the assured position navigation and timing (APNT) capabilities, both automotive and military applications need to have some ability to continue to operate without any external network connections for navigation.

Distributed operations under degraded communications with peer-to-peer coordination is not yet leveraged in automotive systems, however this capability will be is essential as systems in both marketplaces will become more autonomous. Even if information from central control is limited, environmental information can be shared from one system to another to better enhance operational efficiency for the short term. Both systems need to learn to change in real time and be able to alert local peer systems of these new challenges.

Sensor proliferation and data deluge is a huge problem with human operated platforms. Air force pilots sometimes complain about information overload and shut off non-essential sensor systems in battle. As sensor systems are added to automotive systems for safety, the newer automotive system architecture will add software tools and middleware to seamlessly adapt that data into world simulations and models. The automotive chips will add capability like holoscan and other sensor fabric connections to bring this new data into the inference engine in a timely manner. The military may have many more sensor systems than the average car and will probably need many more automotive chips than what is typically found in the automotive market, but SOSA and other standards will allow the Military to scale up to mission need for any platform.

Adversarial pressure: cyber, EW, and deception is unique to the military system while automotive systems will require some level of tamper resistant capability to prevent external mischief or crime, the EW and active deception of automotive systems probably won't have an analog to military platforms. However, the difference between sensor input and EW and cyber-attack systems would be like the simulated environment tools that the automotive market will develop. The military can tailor these tools to adapt deployed inference systems to use this capability in field.

Survivability, attrition, and graceful degradation is a capability that both automotive and the military systems need – as things fail a planned and coordinated method to continue to operate must be built-in, until the platform is at some safe state for repair. Both systems would benefit from being repaired at a high level so that fielded repair is possible without special depot or equipment to support repair. Military systems can benefit with SOSA standard formfactors to improve maintenance realities at the edge.

Environmental extremes and SWaP-T constraints. Military systems and automotive systems must be able to operate where humans need them to operate and must be able to survive these extremes. Military systems in airborne environments do differ than automotive applications, but the commercial airborne systems are looking at automotive technology the same as Military because it also wants to leverage that trillion-dollar electronics market. By leveraging the automotive market, the Military and airborne markets will benefit from access to the latest software and hardware that is rugged enough to deploy, but with the added benefit of having built-in safety and security to protect humans.

Power budgets and energy discipline. Military platforms mission time is dependent on available power and power consumption usage. C5ISR systems are competing with weapon system transportation, and

sometimes human life support. Therefore, the C5ISR will be limited by the electrical power generation the platform can provide. Because these power systems develop very slowly and lag other technologies in the platform, the power system will be the ultimate barrier for military systems to adopt the latest technologies developed for industries like HPC and datacenters. Whereas, the Automotive market has the same exact limitations as all C5ISR systems, and they align with AI need and power available to meet mission. Cerebus just recently announced a wafer scale inference engine needing 23KWs of power[11]. This technology, while possible in a data center, would exceed most power budgets for automobiles or C5ISR systems. Therefore, HPC and datacenter technologies will continue to diverge from automotive and C5ISR technical needs.

5.1. Token processing replaces HPEC flops/watt metric

Lastly, using the HPEC architecture above, the SBC is removed, and 6 Thor systems are easily integrated into the system resulting in a disruptive 12 Petaflops of FP4 performance. This architecture vastly improves processing capability by PIC replacement; the disruption is that the automotive processor combines the SBC scalar processor and the parallel processor into one component improving processing density and leverages automotive roadmaps that will improve this capability further. The key difference is this Thor system is optimized to process AI tokens as well as provide signal processing – it is built to run AI inference at the edge. General processing, even though it exceeds the latest laptop GPGPU system currently, is secondary to its ability to run the AI models trained by the large datacenters.

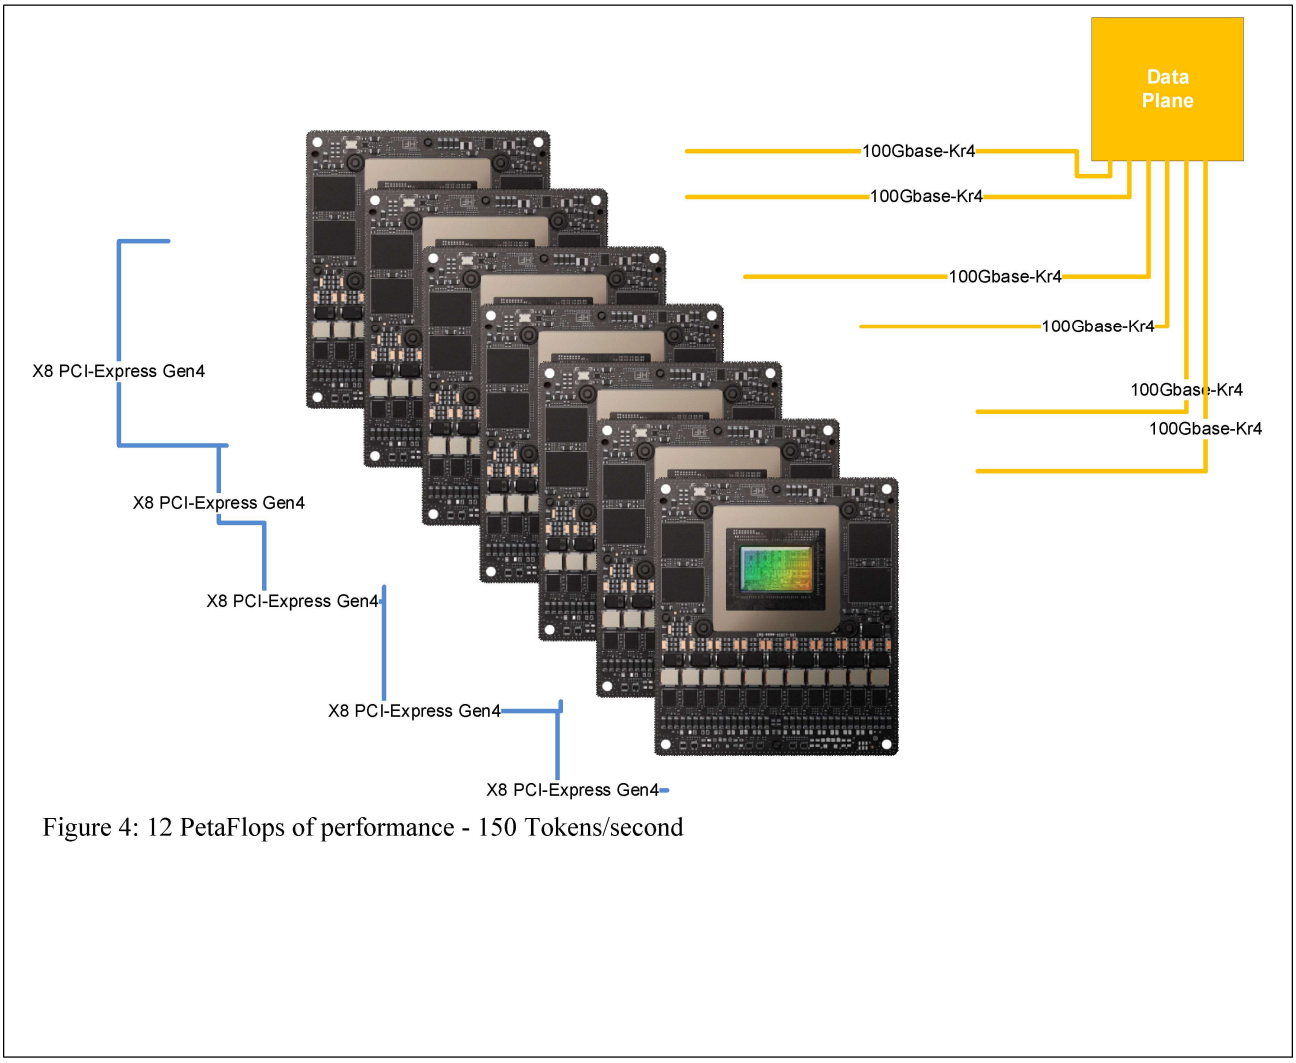

Figure 4: 12 PetaFlops of performance - 150 Tokens/second

6. CONCLUSION

Automotive shift to electrified, autonomous, and software-defined vehicles is cascading into every industry that depends on edge decisions. The winners won't just add compute; they'll adopt an architecture that treats bandwidth as scarce, latency as lethal, and updates as routine. NVIDIA Jetson Thor supplies the muscle and the plumbing for massive on-module AI, sensor native I/O, and a mature robotics stack to make that architecture real at the point of need.

To conclude this paper four more axioms are needed to define HPEC systems. Therefore, there are 14 Axioms to HPEC.

11) HPEC systems must be able to perform mission with reduced or designed network connections. HPEC systems require some level of autonomy in performing their main function. The example of a laptop applies here and is one of the simpler examples of HPEC. Network connections to a laptop are increasing in need so that the system can leverage all the cloud-based applications, but the laptop (unlike cell phones) can carry out its primary function without network connection. For example, entertainment, and some aspects of human work, is always possible with a laptop on top of a mountain. That is why laptop components are always going to be part of the HPEC sensor chain. Specifically for the human machine interface (HMI), the components of laptops provide a great basis for HPEC HMI-subsystems. Cellphones (currently) are enabled communication devices and are primarily focused on network connections, in time the cell phone and laptop technologies will merge such that cellphone, tablets and laptops will just be the same device and only the HMI will change based on need.

12) HPEC systems require some AI capability. The whole HPEC world requires AI. This world and the way HPEC systems are being deployed the man in the middle is becoming less part of the closed loop operation, therefore, see do and act may have no human in the loop for autonomous systems. The next war may evolve to be intelligence gathering and autonomous systems fighting other autonomous systems. HMI might be just a function to record the winners results and not be part of the sense act do loop. This activity is simply recognizing that humans cannot process at machine speed and when the war is won or lost in nanoseconds or time scales beyond human perception, AI has to be the responsible for the "do" in the sensor loop. Since the HPEC requires some ability of autonomy, the

13) HPEC sensors will require more intelligent fabrics to separate sensors and processing innovation. The level of innovation of digitizing the analog world to digital operates at a different pace than processing. Antenna, RADAR, LIDAR, and optics aka the digitization of light that make

up most of the C5ISR sensor cluster and the avionics or vehtronics for control of the platform are systems that change very slowly. For sensors, the technology to capture the analog world is a very slow development as the laws of physics (aka Heisenberg's uncertainty principle) are a very large barrier to enable new and disruptive ways of collecting information. Vehicles control is limited to the actuators and the low-level controls on those vehicles and these vehicles do not change every year either. Therefore, HPEC sensors need to be separated either physically or conceptually, so that if sensor innovation improves, the HPEC system can take advantage of that information without significant or very little change. Correspondingly, if signal processing systems vastly improve then these systems must be able to be changed quickly as well. Therefore, the only way to separate these two components is to connect with a low latency intelligent network that is sufficient in bandwidth to process and act on the sensor data so that the mission or application is useful. Without this separation innovation is slowed for adoption.

14) HPEC systems require liquid to be cooled. Perhaps the most controversial axion to add to the list. Not all electronic systems require liquid cooling, and many sensors and applications would not support liquid. However, all HPEC systems that have the entire sensor chain and have the ability for autonomy at this moment of time depend on liquid cooling because it is the best method to cool all the various components in the same manner. Liquid heat density (volume of heat transport) is 3000 times better than Air [12]. All the commercial autonomous AI systems use liquid as a method of cooling. Because HPEC is a cluster of resources to provide the sensor chain, liquid cooling can be standardized so that each component can cool the same way with the same mechanical interfaces and method which simplifies mechanical design. Fancy thermal simulations are back now to the envelope calculations on inlet and liquid flow rates and the transfer of this heat transients. For hotter systems or more thermally dense systems, will require more flow rate or cooling to a central SWAP optimized heat exchanger. Finally, components TDP and platform size will always limit the amount of capability on any HPEC system, but liquid is the method that best cools HPEC electronics.

Implications beyond automotive

What happens on the road never stays there. Automotive's tempo is setting expectations for other domains: modular platforms that can be refreshed quickly, pervasive sensor fusion at the edge, and continuous feature delivery without pulling systems from service. As vehicles evolve into mobile data centers, cross-market adopters from industrial to defense are inheriting proven patterns for real-time perception, resilient autonomy under degraded links, and rigorous telemetry for safety and maintenance.

These battlefield realities set the bar for any next-generation C5ISR node. They favor architectures that are modular, thermally disciplined, power-aware, maintainable, and able to move fast at the edge, fight through disruption, and evolve with minimal rework.

7. REFERENCES

[1] Stewart, Timothy *Engineering.com*. "How 'AI at the edge' enables decision superiority on today's battlefield." https://www.engineering.com/how-ai-at-the-edge-enables-decision-superiority-on-todays-battlefield. August 2025.

[1] Ferdinando, Lisa U.S., "Cloud Data to Give Warfighters Edge to 'Dominate' in Battlefield, DoD Official Says (John H. Gibs)." https://www.war.gov/Newsroom/News-Stories/Article/1502447/cloud-data-to-give-warfighters-edge-to-dominate-in-battlefield-dod-official-says/. April 2018.

[2] Blakeney, Ryan Over the Horizon Journal. "Edge ISR: Revolutionizing Intelligence in Contested Environments." https://othjournal.com/2025/04/25/edge-isr-revolutionizing-intelligence-in-contested-environments April 2025.

[3] Ayad, P., & Bultena, L. *A once in a century transition: How automotive's transformation is reshaping global strategy*. Hogan Lovells. Retrieved from Hogan Lovells https://www.hoganlovells.com/en/publications/a-once-in-a-century-transition. 2023

[4] McKinsey & Company. *New twists in the electric-vehicle transition: A consumer perspective*. https://www.mckinsey.com/features/mckinsey-center-for-future-mobility/our-insights/new-twists-in-the-electric-vehicle-transition-a-consumer-perspective. April 2025.

[5] Lapczynski, P. *Securing automotive over-the-air software updates*. Renesas. https://www.renesas.com/en/blogs/securing-automotive-over-air-software-updates. July 2022.

[6] McKinsey & Company. *Unlocking the full life-cycle value from connected-car data*. https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/unlocking-the-full-life-cycle-value-from-connected-car-data. February 2021.

[7] Foley & Lardner LLP. *Reshoring Accelerates Amid Trump's 2025 Tariff Surge*. National Law Review. https://natlawreview.com/article/what-every-multinational-should-know-about-use-reshoring-navigate-tariff August 2025.

[8] "National Highway Traffic Safety Administration (NHTSA)". *Vehicle Cybersecurity*. https://www.nhtsa.gov/research/vehicle-cybersecurity. 2025.

[9] "Introducing NVIDIA Jetson Thor, the Ultimate Platform for Physical AI" .https://developer.nvidia.com/blog/introducing-nvidia-jetson-thor-the-ultimate-platform-for-physical-ai/

[10] "What is a Holoscan Sensor Bridge – And How Does It Work?" https://forums.developer.nvidia.com/t/what-is-a-holoscan-sensor-bridge-and-how-does-it-work/345102
[11] "Breaking the Thermal Ceiling in Chip Design. "https://liquidstack.com/blog/breaking-the-thermal-ceiling-in-chip-design#:~:text=In%20April%20of%20this%20year%2C,it%2C%20Wafer%20Scale%20Engine%20(WSE).
[12] "Vertiv / Insights High Density cooling: A guide to advanced thermal solutions for AI and ML workloads in datacenters."https://www.vertiv.com/en-us/insights/articles/educational-articles/high--density-cooling-a-guide-to-advanced-thermal-solutions-for-ai-and-ml-workloads-in-data-centers/

8. CONTACT INFORMATION

William J. Pilaud
Chief Solution Architect
LCR Embedded Systems
2621 Van Buren Ave
Audubon, PA 19403
bpilaud@lcrembedded.com
978-231-3170
Linked In: linkedin.com/in/billpilaud

9. ACRONYMS

GVSC	Ground Vehicle System Center
PID	Proportional Integral Derivative
SVD	Singular Value Decomposition
UDP	User Datagram Protocol