NVMe Unlocks Data Access and Analysis at the Source

By Premio
www.Premioinc.com

Table of Contents

Executive Summary

Maximizing data value and usefulness requires a smart data strategy. As industries work to capitalize on the deluge of data generated by intelligent, connected devices, and intelligent data strategy has never been more critical. Enter edge computing, playing an essential role in addressing this challenge by enabling access and analytics close to the data source, a real transformation from the once pure cloud focus in Big Data processing. Yet the edge itself continues to evolve, with more demanding, real-time industrial computing redefining it as the **rugged edge** – environments where the combination of high- performance computing and real-time data is changing the way industrial businesses operate and compete. NVM Express™ (NVMe™) is a crucial enabling technology for the rugged edge, reducing latency and accelerating data processing to the point at which artificial intelligence, machine learning, and inference computing are realities anywhere a connected device can operate. This paper will explore the technical benefits of NVMe, elevated by advancements such as fanless design and hot-swappable capabilities, that add critical value to the rugged edge. Offering brief scenarios in settings such as industrial automation, telematics, transportation, surveillance, smart retail, and more, this paper will illustrate how real-time data access and analytics are turning rugged edge computing into a competitive necessity.

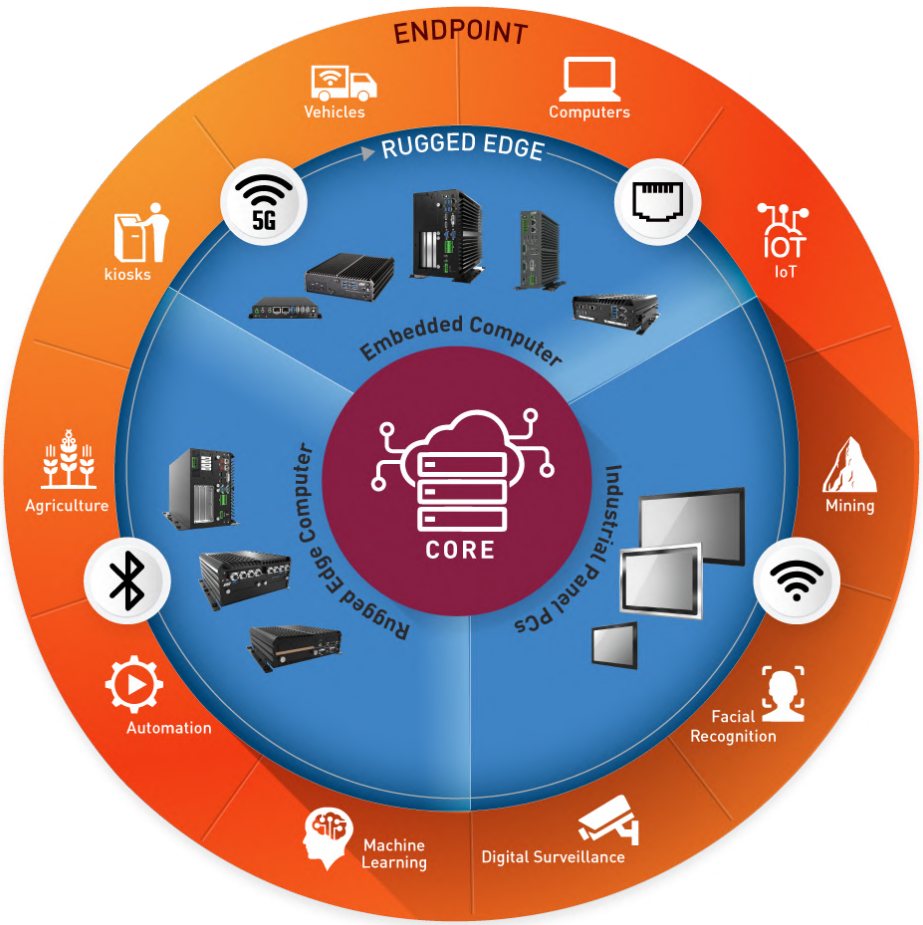

Moving to the Rugged Edge: *Data is the new gold, but only if you can mine it*

We are constantly generating data. People and machines generate information from all types of digital transactions and processes. The term "Big Data" accurately describes the result. Intelligent application of this data is what digital transformation is all about – uncovering it, accessing it, and using it to facilitate transactions, automate processes, increase safety and productivity, and improve services.

Industry analyst firm IDC recognizes that data is produced at three primary levels. These include the core (traditional and cloud-based data centers), the edge (enterprise-hardened infrastructure settings like cell towers and organizations' branch offices), and the endpoints (PCs, smartphones, and IoT connected devices). This "always-on" infrastructure increases the amount of available data to be captured – a global datasphere estimated to grow to 175ZB (Zettabytes) by 2025 – and increases reliance on data itself. The intersection of Big Data, connected devices, and edge computing has created a fast-evolving global business opportunity as faster and more intelligent analytics enable more significant insights at all levels of industry.

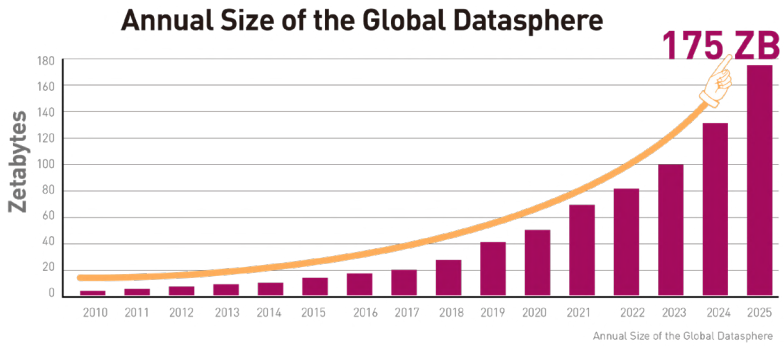

Yet as edge computing migrates technology and processing closer to where IoT sensors and devices generate data, these environments are inherently not as stable or regulated as data centers. Factories and warehouses are characterized by hot and cold extremes, processing plants where water and dirt are part of the job, or mobile environments that routinely experience shock and vibration. These and similar environments are the rugged edge. They demand specialized hardware to consolidate workloads that matter most, where operations benefit from smarter and more advanced automation. Rugged, high-reliability systems are breaking new ground in these rigorous settings by efficiently passing data back and forth to the core and capitalizing on reduced latency, accelerated processing, and increased data storage capacity right where it is needed. Where the harsh environment once prevented optimized, high-performance computing – or at least challenged it to the point of extreme limitation – today's rugged edge systems are accelerating a shift to machine learning and more intelligent decision making in real time in such environments.

Meeting the rugged edge challenge with connected, hardened systems

The rugged edge is a unique place, where heavy industry relies on computing systems that blend the latest high-performance technologies to accelerate data processing fed by a wide variety of sensor input data. At the same time, environments are often harsh and unstable, creating unique requirements for systems that are hardened at every level and developed to meet ever-growing needs for IoT connectivity and performance. Because the top priority of emerging IoT applications is to deliver continuous intelligence and insight, an effective supporting system must enable some level of machine learning and offer dedicated hardware to process and run sophisticated algorithms.

Critical access to wireless data communications via Bluetooth, Wi-Fi, as well as cellular 5G and 4G LTE technologies – along with the flexibility to embrace future connectivity technologies – enable rugged edge systems to meet this need. Wireless versatility makes a tangible difference, particularly as wireless sensors and devices become exponentially more diverse in their number and type. For industries in need of machine-to-machine communication, these options bring a greater level of real-world value to applications such as real-time data monitoring, GPS/GNSS location tracking, and all types of autonomous systems.

In these environments, mission-critical reliability matters – industrial-grade compute systems must be further enhanced to eliminate downtime and maintain connected performance no matter the ambient condition. Small form factor designs taking these into consideration simplify remote integration where applications often face space constraints and harsh operating conditions. Integrated watchdog system monitoring helps minimize downtime, reducing the need for on-site technical support, which can be difficult or impossible to access depending on the deployment.

A rugged edge computer is a fanless industrial grade computer that brings computation and data storage physically closer to where an application is operating in harsh usage conditions. Local IoT sensors and devices can process data through local edge computers for faster response times with lower latency instead of sending it to a distant cloud data center for off-site processing. Rugged edge computers are fortified and proven to operate reliably in extreme environments where typical computers are guaranteed to fail due to either severe hot or cold temperatures, dust or debris, humidity, widely varying voltage sources, and shock or vibration. With a higher MTBF (mean time between failure), rugged edge computers provide reliable mission critical operation for industrial IoT application systems.

Achieving data center performance at the rugged edge

Taken together, the factors defining rugged edge design are creating a new competitive challenge for system engineers. They often face the need to design compute solutions for unique and complex scenarios in which they lack real-world expertise. Premio plays a pioneering role in this landscape, demonstrating the blend of technologies optimized for rugged edge computing on a global stage. Company deployments are proven in a compelling slate of harsh computing environments, ranging from fish packing plants to underground mining to onboard fleet management.

By applying singular focus on key factors that benefit these exceptionally demanding settings, Premio has established that data center performance and architecture can be deployed at the rugged edge. Decisions centering on compute, storage, wireless connectivity, security, and system level ruggedization are crucial – together, these create solutions explicitly defined as rugged edge systems. This is not your average edge computing and even redefines 'rugged' to mean environments that have not been fully-served by smart automation due to the hurdles of achieving reliable machine learning performance.

Optimized rugged edge systems incorporate powerful multi-core processors with low-power options and advanced process nodes such as 14nm. Deploying Intel®'s 8th and 9th generation of Core™ processors balances performance and thermal regulation for maximum reliability. Ideally, CPU, GPU, and VPU support are balanced to accommodate machine learning and inference requirements at the edge, further enabled by fast 10Gbe wired connectivity for data uplink and downlink. Along with rugged high-speed storage, these performance accelerators allow real-time local processing without dependence on cloud technology, eliminating the bottlenecks from pass-through of unwieldy Big Data sets.

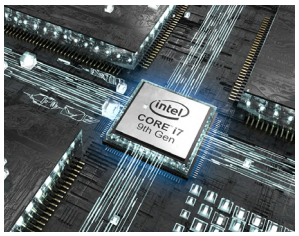

Consider the diversity of industrial controllers and deployed machinery, typically purpose built for non-stop heavy industrial use and often deployed for decades. Versatile I/O at the edge adds value and accommodates this range of legacy and modern devices. Designers can access COM, USB, GPIO/Digital I/O, LAN/PoE, and also remain prepared for new options with PCI/mini-PCIe/PCIe expansion ports.

Hardware-based security features must be embedded via TPM 2.0 hardware security, with cryptographic keys used to protect the digital data communications between devices from external attacks. These protocols, coupled with edge autonomy, safeguard sensitive data and signals, allowing processing without rerouting to less secure cloud resources. This is a critical advantage of rugged edge computing, as machines communicate bidirectionally, and data requires constant protection. TPM 2.0 further prevents authentication keys from being exported while recording and autonomous reporting software uploads during the boot process to thwart malware attacks.

Initiated by the Trusted Computing Group (TCG), a non-profit international standards organization, the concept of hardware security was promoted to remedy vulnerabilities that quickly became prevalent with the rise of the Internet and connected transactions. Premio's rugged edge systems include TPM 2.0, ensuring trusted status that combats increasingly sophisticated cyber-attacks on industrial systems.

Protecting the trusted status of rugged edge systems with hardware-based security

TPM 2.0 provides a secure foundation, functioning as an integrated module that enables on-board hardware security in edge computing systems. Initiated by the Trusted Computing Group (TCG), a non-profit international standards organization, these modules are physically fixed to the device motherboard by the manufacturer, establishing a chain of trust in securing firmware and software subsequently applied to the system. Expansive security features implemented at the OS level are then built on top of security-enabled hardware and firmware.

TCG's Opal 2.0 Storage Specification adds further value for rugged edge systems with high data performance enabled by SSDs. Opal 2.0 enhances security features of data storage devices, for example, defaulting to 'always-on' hardware-based encryption, essentially requiring that encryption is applied to stored data. Data at rest cannot be accessed or viewed without proper authorization and decryption, even if a drive was removed, lost, or stolen. This tamper-proof approach most commonly uses AES encryption, powering self-encrypting drives (SEDs) that require Logical Block Address (LBA)-specific permissions. Only password-authorized users are granted access, which is further restricted to permitted actions.

It is important to note that Opal-compliant drives deliver advantages over software-based encryption options, such as improved performance, security, and overall management. No host device or system is overloaded with encryption processes, which take place in the SSD itself. This is inherently valuable for rugged edge deployments – not only is high performance security present and protecting all data communications and signal transmissions but also these demanding processes are handled effectively and without performance bottlenecks.

Capitalizing on NVMe for fast, high-performance data transfer, access, and storage

NVM Express™, or Non-Volatile Memory Express (NVMe™), is a proven data center protocol that also benefits edge workloads. NVMe represents a breakthrough in reducing latency and accelerating data processing – critically important in enabling advanced operations such as artificial intelligence, machine learning, and inference computing. Used in hot-cache storage for quick data access in optimized data centers, the same protocol brings low-latency, high-speed storage access, and high capacity to the rugged edge. For instance, machine learning algorithms require fast non-volatile storage for AI training and inference. At the same time, valuable data must be filtered back to the data center for advanced IoT management. These requirements, coupled with the continued growth of IoT applications, drive the need for high-performance storage at the core as an imperative, not just a 'nice to have' feature.

NVMe was designed from the ground up to unlock the true performance of high-speed storage media such as solid state drives (SSDs) via the PCIe bus. This approach was developed in response to the limitations of traditional interfaces primarily design for interfacing with hard drive storage, challenged by their inability to take advantage of PCIe lanes and proximity to the CPU. Via the PCIe Gen 3 protocol that connects directly to the CPU architecture, NVMe executes data transfers 6 times as fast as SATA 6Gb/s bottleneck speeds. Data and command processing are executed in parallel, and it leverages low latency paths in an architecture similar to a high-performance processor. As a result, CPU cycles are streamlined – eliminating the I/O bottlenecks that persist in legacy storage protocols like SATA and enabling the scalable performance required for reflexive inference analysis at the edge.

With reduced latency and much higher transfer bandwidth that significantly increases read/write capabilities performance and IOPs throughput, not only are workloads optimized and accelerated in existing applications, but new real-time processing options are enabled. NVMe SSDs also increase ruggedness as well, storing data in flash memory with no moving parts, something that mechanical spinning disks lack. The architecture itself contains custom- developed features to regulate the power of SSDs, ensuring effective power management as an advantage that lowers the total cost of ownership (TCO) and extends battery life.

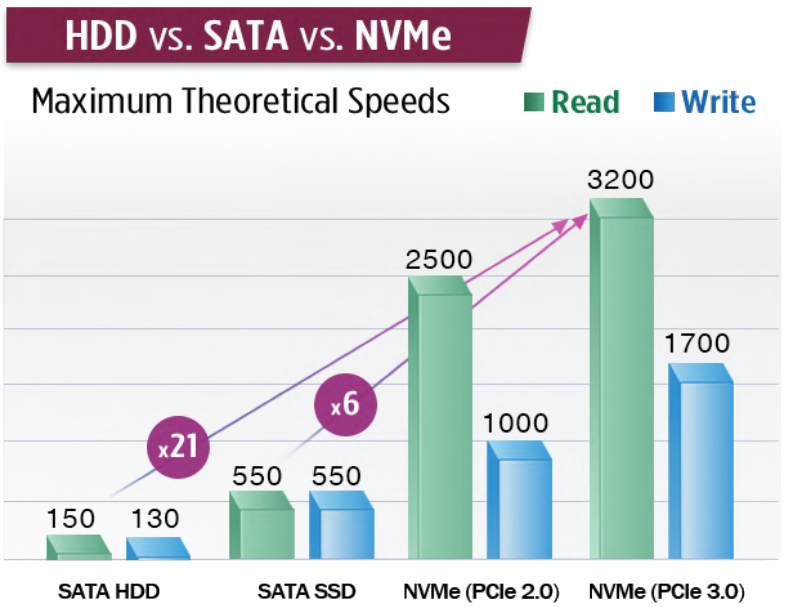

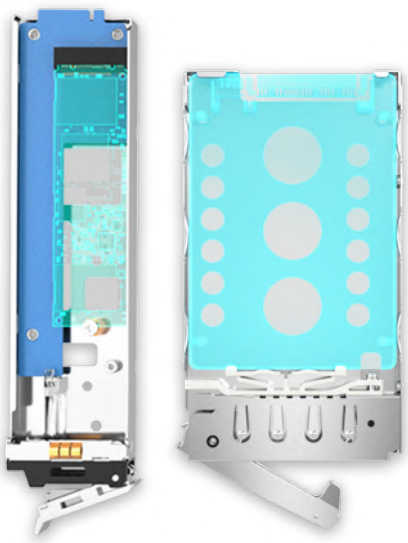

Accessible M.2 NVMe Drive Carriers

Accessible U.2 NVMe Drive Carriers

NVMe-enabled rugged edge advances AI-powered automation and services

Premio is first to market with an easy access, hot swappable NVMe storage (M.2) solution in a semi-ruggedized design. The company's patented carrier board design enables an NVMe M.2 storage tray as a host, increasing flexibility by allowing quick drive access and replacement. This approach eliminates read/write bottlenecks that can suppress real-time processing at the rugged edge, an advancement over competing solutions that offer only a single onboard NVMe drive.

Premio also supports the U.2 interface for connection with NVMe storage. Where small, flat M.2 drives are placed directly on the motherboard, the U.2 cable can be used to access separate SSDs in the 2.5" form factor common to traditional SATA SSDs. Depending on the deployment environment, this may provide a thermal advantage by separating the drive and its generated heat from the motherboard. Both options communicate to the NVMe storage system's motherboard directly via the PCIe bus, and both options use the same type of flash memory storage. The U.2 option provides this additional flexibility to end-users, allowing a somewhat higher storage capacity compared to M.2 (~8TB compared to M.2's maximum of 2TB). Overall, U.2 enables a faster data pipeline by using up to four PCIe lanes and two SATA lanes.

Product Brief
AI Edge Inference Computer

RCO-6000-CFL-4N-2060S, RCO-6000-CFL-4NH, RCO-6000-CFL-8NS

Compute, Storage, Connectivity, Ruggedness: Necessary Ingredients for Inference

Industry 4.0 is driving the need for improved intelligence at the rugged edge. Businesses and organizations want more compute power to perform inference analysis at the edge because of the benefits it provides to their operations. Deploying powerful inference computers at the edge to process data and make decisions in real-time enables organizations to reflexively respond to situational data, glean keener insights into their operations, achieve low latency data processing and decision-making, and reduce their network demands.

Premio has combined next-generation processing and high-speed storage technologies with the latest IoT connectivity features to create a solution designed from the ground up to deliver holistic inference analysis at the rugged edge.

The RCO-6000-CFL Series AI Edge Inference Computer incorporates advanced performance Intel 9th Generation processors, rich GPU support, and scalable, hot-swappable NVMe capacity into versatile hardware that's designed to withstand deployment in challenging industrial environments. Systems deployed in industrial environments are often exposed to dust, debris, shock, vibration, and extreme temperatures. Premio's AI Edge Inference Computers are hardened to withstand exposure to these harsh environmental factors, providing optimal and reliable performance regardless of the deployment environment.

Key Applications

- AI & Machine Learning
- Edge Computing & IoT
- Industrial Automation
- Telematics & Transportation
- Metrology & Defect Detection
- Surveillance
- Smart Retail
- ADAS and Autonomous Vehicle Data Capture & Storage

Inference Analysis at the Rugged Edge

The RCO-6000-CFL AI Edge Inference Computer supports GPUs (graphic processing units), allowing systems to process and analyze large volumes of complex data to drive AI functions at the edge. AI edge inference computers are purpose-built for deployment in either controlled or dynamic environments to capture data from sensors, high-resolution cameras, and other devices and peripherals that use real-time data for machine intelligence.

These powerful computers are capable of processing data in real-time from a variety of IoT sensor inputs at the edge so that such data can be stored for training machine learning and deep learning models at a later time. Training a neural network involves feeding the collected data to the neural network, allowing the algorithm to predict what the data represents. If the neural network makes an inaccurate prediction, the training process goes on until the neural network achieves the desired accuracy without human intervention, creating a level of machine intelligence.

After a neural network is trained, the refined algorithms are deployed at the rugged edge to conduct inference analysis. The speed and accuracy of a properly informed algorithm greatly exceed those of human counterparts. Industrial and critical infrastructure sectors are awakening to the fresh insights and optimization potential for safer, more resource, and cost-effective automation.

The RCO-6000-CFL AI Edge inference Computers features blazing-fast NVMe storage and high-speed I/Os that drive more efficient data collection and processing for training neural networks more quickly than ever before. Additionally, systems can transfer the data that's gathered in real-time to the cloud for advanced processing for more machine learning. Operators can use the system's GPU to apply AI algorithms to perform inference analysis at the rugged edge in real-time.

Multi-Core Computing For the Evolving Edge

Intel 9th Generation Processors & Q370 Chipset

The RCO-6000-CFL AI Edge Inference Computer leverages rich performance enhancements provided by 9th Generation Intel CFL-R S Processors and Q370 Chipset support. Intel technology enables 16-way multitasking through hyperthreading all eight cores. The processor supports DDR4 RAM for up to 64GB of memory and 2666 MT/s transfer speed, while UHD graphics offer rich visual output for many applications using optical data.

The LGA1151 socket design is combined with Intel's Q370 chipset to deliver augmented peripheral performance for low-latency edge responsiveness. Gigabit wireless speeds, PCIe 3.0 lanes, SATA ports, and high-speed USB 3.2 Gen 2 grant the RCO-6000-CFL AI Edge Inference Computers excellent I/O integration options for transmitting data to and from sensory devices sitting at the edge.

Performance Blocks: 2 Piece Modular Design

Industrial Fanless PC on Top

The RCO-6000-CFL Series of Edge AI Inference is extremely modular and comes in three distinct configurations that act as performance building blocks. The Base Model RCO-6000-CFL Series is a fanless, rugged computing solution capable of performing powerful computing at the edge while enduring harsh environmental conditions. The base system can be configured with 2x internal 2.5" SATA SSDs in 9mm height, 2x hot-swappable SATA SSDs in 7mm height, and a single internal M.2 NVMe SSD on the motherboard. Several optimized technologies have converged to promote real-time, in-depth responsiveness at the edge. Multi-core CPUs and advanced GPUs stand ready to perform numerous parallel processes, while 5G, 10GbE, and speedier I/O technologies wait to receive and offload volumes of rich data. SSD (solid-state drive) storage offers a vastly quicker and structurally more rugged data repository than its HDD (hard disk drive) cousin.

Top
Industrial
Fanless PC

Bottom
EDGEBoost
Performance
Nodes (GPU and
/or NVMe SSD)

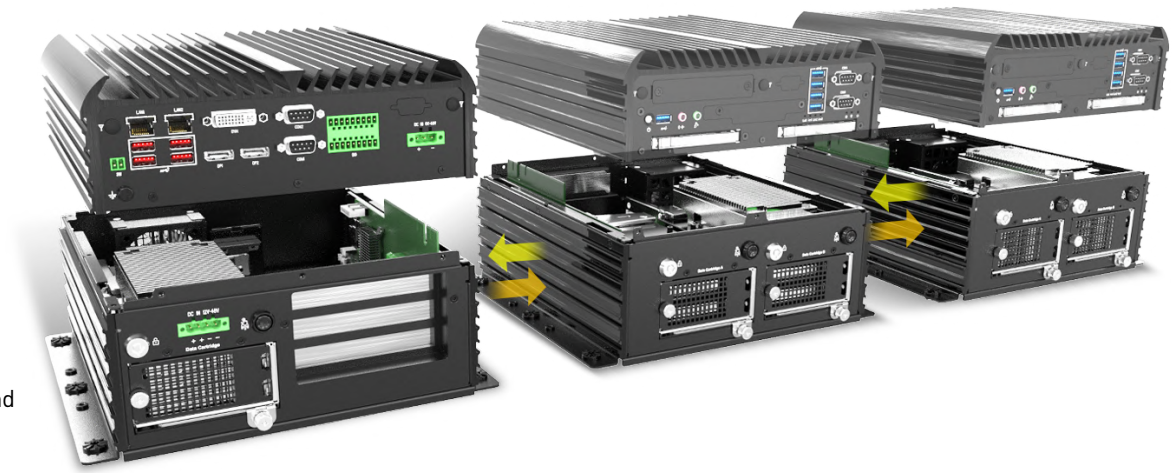

Flexible and Dedicated "**EDGEBoost Nodes**" on Bottom

A key differentiator of Premio's RCO-6000-CFL AI Edge Inference Computer is its modular add-on nodes for powerful inferencing capabilities and high-performance NVMe storage. Users can select specific nodes that provide high-speed, high-capacity NVMe storage and GPU performance acceleration directly for their edge application deployments. This innovative mechanical design provides the ability for a rugged and fanless industrial computer that pairs with performance acceleration nodes, or "EDGEBoost Nodes" dedicated for specific workloads at the edge.

Premio currently offers three types of EDGEBoost nodes that can expand the capabilities of the RCO-6000-CFL Series AI Edge Inference Computer:

EDGEBoost Node # 1

RCO-6000-CFL-2N-2060S

This EDGEBoost node attaches to a base RCO-6000-CFL Series and adds a hot-swappable NVMe SSD canister, capable of being populated with up to 2x lockable and hot-swappable 2.5" U.2 NVMe SSDs in 15mm height. This EDGEBoost node also adds PCIe expansion slots, enabling organizations and system integrators to add an Nvidia 2060 Super GPU for inference acceleration.

Exascend's PI3 series SSD is an industrial-grade (-40 to 85°C) PCIe Gen3x4 product with enterprise performance features (high sustained speed and low latency), available in standard M.2 2280 and U.2 form factors with up to 8TB capacity. The company's ultra-high capacity, high-speed PCIe Gen3 NVMe SSD products are based on Marvell's industry-leading PCIe NVMe SSD controller through a unique strategic partnership.

Performance - RCO-6000-CFL-2N-2060S (2x PI3 1.92TB)	
Maximum Sequential Read (MB/s)	6,680
Maximum Sequential Write (MB/s)	3,816
Maximum Random Read (IOPS)	600,227
Maximum Random Write (IOPS)	464,193

U.2 SSD Performance Benchmark

EDGEBoost Node # 2

RCO-6000-CFL-4NH

This EDGEBoost node focuses on ultra-high-speed NVMe storage but supports NVMe SSD storage media in 2.5" U.2 15mm form factors for higher capacities. This specific EDGEBoost node offers x4 lockable and hot-swappable NVMe SSDs configurable in RAID options in 0,1,5, and 10.

Exascend's PI3 series SSD is an industrial-grade (-40 to 85°C) PCIe Gen3x4 product with enterprise performance features (high sustained speed and low latency), available in standard M.2 2280 and U.2 form factors with up to 8TB capacity. The company's ultra-high capacity, high-speed PCIe Gen3 NVMe SSD products are based on Marvell's industry-leading PCIe NVMe SSD controller through a unique strategic partnership.

Performance - RCO-6000-CFL-4NH (4x PI3 1.92TB)	
Maximum Sequential Read (MB/s)	7,016
Maximum Sequential Write (MB/s)	4,044
Maximum Random Read (IOPS)	1,033,365
Maximum Random Write (IOPS)	760,615

U.2 SSD Performance Benchmark

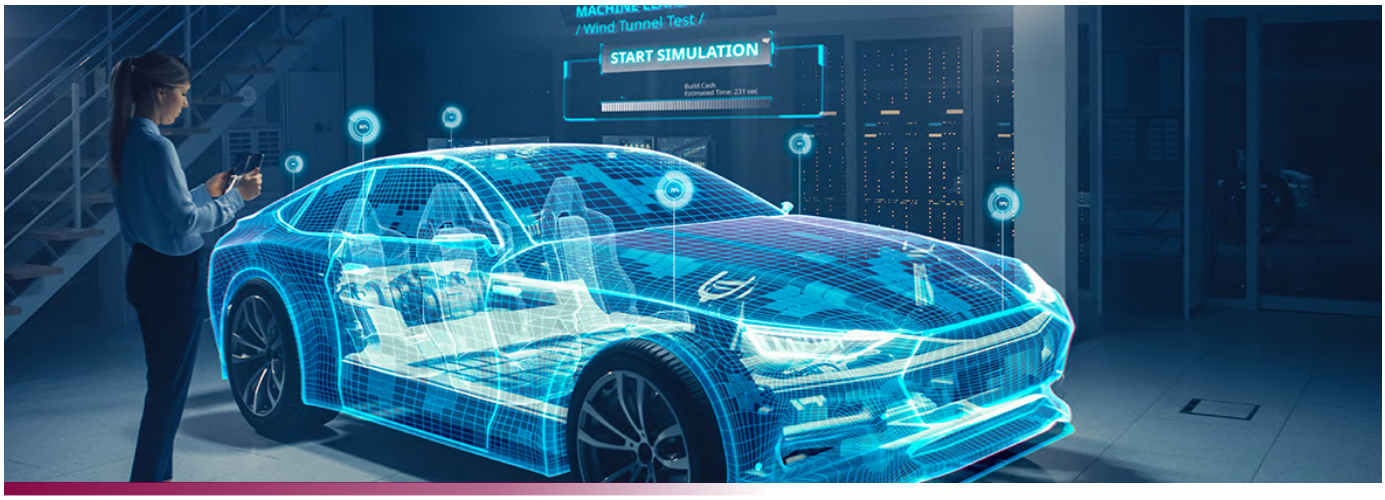

EDGEBoost Node # 3

RCO-6000-CFL-4N-2060S

This EDGEBoost node attaches to a base RCO-6000-CFL Series and adds a hot-swappable NVMe SSD canister, capable of being populated with up to 4x lockable and hot-swappable 2.5" U.2 NVMe SSDs in 7mm height. This EDGEBoost node also adds PCIe expansion slots, enabling organizations and system integrators to add an Nvidia 2060 Super GPU for inference acceleration.

Exascend's PI3 series SSD is an industrial-grade (-40 to 85°C) PCIe Gen3x4 product with enterprise performance features (high sustained speed and low latency), available in standard M.2 2280 and U.2 form factors with up to 8TB capacity. The company's ultra-high capacity, high-speed PCIe Gen3 NVMe SSD products are based on Marvell's industry-leading PCIe NVMe SSD controller through a unique strategic partnership.

Performance - RCO-6000-CFL-4N-2060S (4x PI3 1.92TB)	
Maximum Sequential Read (MB/s)	7,052
Maximum Sequential Write (MB/s)	4,012
Maximum Random Read (IOPS)	1,032,451
Maximum Random Write (IOPS)	784,324

U.2 SSD Performance Benchmark

EDGEBoost Node # 4

RCO-6000-CFL-8NS

This EDGEBoost node focuses on ultra-high-speed NVMe Storage and offers users the ability to add up to 8x lockable and hot-swappable 2.5" U.2 NVMe SSDs in 7mm height via two hot-swappable NVMe SSD canister bricks. The canister design allows organizations to quickly and easily remove all SSDs from the system to offload mission-critical data onto a central computer system. This allows for an easy and efficient way to transfer data at the edge and into a location with resources reserved for machine learning (ML) and deep learning (DL) training models.

Exascend's PI3 series SSD is an industrial-grade (-40 to 85°C) PCIe Gen3x4 product with enterprise performance features (high sustained speed and low latency), available in standard M.2 2280 and U.2 form factors with up to 8TB capacity. The company's ultra-high capacity, high-speed PCIe Gen3 NVMe SSD products are based on Marvell's industry-leading PCIe NVMe SSD controller through a unique strategic partnership.

Performance - RCO-6000-CFL-8NS (8x PI3 1.92TB)	
Maximum Sequential Read (MB/s)	14,300
Maximum Sequential Write (MB/s)	11,928
Maximum Random Read (IOPS)	2,440,000
Maximum Random Write (IOPS)	1,752,000

U.2 SSD Performance Benchmark

SSD Data Protection and Redundancy

Safety Storage Ejection Button

Adding to the ease of offloading data from the RCO-6000-CFL Series to the cloud is the availability of a physical button on the system that initiates the ejection of storage media for the safe removal of the SSD canister or individual SSDs. Pressing the button suspends all I/O operations and read/write operations to the storage devices to prevent the loss or corruption of data.

Having a robust and scalable data storage solution is a growing concern for today's hardware users. So, we've added the ability to hot-swap not only each individual NVMe drive, but multiple drives can be hot-swapped thanks to the availability of hot-swappable drive canisters. Drive canisters streamline the process of accessing and moving data from the Edge AI Inference computer to a central computer, especially for users that need to offload data from the PC frequently.

Configurable RAID

The Edge AI Inference computer comes with both software and hardware RAID, offering RAID 0, 1, 5, and 10. Configuring your system with a dedicated RAID controller can boost the performance of your system by offloading RAID functions from the host system's CPU to the RAID controller. Offloading RAID functions to a dedicated RAID controller allows the CPU to focus on running enterprise applications.

High-Speed Storage Primed for Edge Intelligence

NVME. Informing the Rugged Edge

Two removable canisters can be populated with up to eight hot-swappable U.2 NVMe SSDs, providing organizations with ultra-high-speed, solid-state storage at the edge and the ability to load and offload canisters for extremely quick data transfers. NVMe enables data-center equivalent data read/write speeds at the edge. High-speed data storage efficiently feeds integrated CPUs and GPUs with large volumes of data for complex applications. Additionally, rapid data speeds enable more reflexive inference analysis at the edge.

In the past, the main obstacle to immediate intelligence at the edge was the inability of storage technologies to efficiently read/write data beyond the inherent limitation in the connective architecture. This is where NVMe comes in and carries the day over more traditional data transfer protocols found in legacy products.

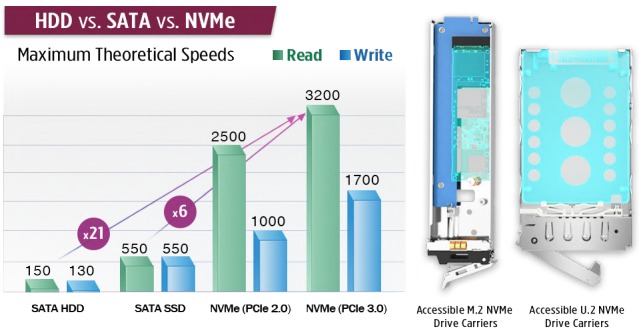

NVMe, or Non-Volatile Memory Express, is an SSD protocol that focuses on efficiency. As an SSD, it relies on semiconductor chips without moving parts to store and access data. NVMe's delivery system is via PCIe 3.0 Lanes. This streamlined interface alleviates data bottlenecking that can occur with other SSD technologies.

NVMe's theoretical speeds of 2,500MB/s are practically quintuple those of SATA. SATA, the other predominant SSD protocol, boasts write speeds of around 550MB/s, significantly slower than that of NVMe. AS far as HDD competition, NVMe delivers 16x the read speed of HDDs.

Deploying a rapid-fire storage media and setting a clear path for its transmission to PC components, NVMe provides the operative immediacy that the responsive edge intelligence demands. Applications leveraging inference analysis can access stored algorithms and mission-critical data at speeds exceeding human cognition.

Connecting From the Rugged Edge

Steady Wireless Connectivity

The RCO-6000-CFL Series enables seamless wireless connectivity for remote and mobile edge deployments. Systems feature both Wi-Fi 6 and Bluetooth 5 technologies to reliably connect to sensors and network systems through a wireless IoT enterprise. Additionally, systems can be configured with Dual External SIM sockets, providing 4G/LTE cellular connectivity for remote and mobile edge deployments. The RCO-6000-CFL Series are also 5G ready through a modular add-in card, providing edge deployments, vastly greater cellular speeds, and more granular network slicing options.

10GbE I/O Ready

The rugged edge inference computer supports two 10 Gigabit Ethernet Ports (Intel x710-AT2 Chipset) through its universal I/O bracket. The high-speed connections enable low-latency data transmission for advanced industrial inference analysis applications.

LAN/PoE Options

The RCO-6000-CFL Series comes configured with dual Gigabit Ethernet Ports. Four additional Gigabit Ethernet or PoE ports can be added to the system via expansion daughterboards. PoE supplies power and data to peripherals via a single ethernet cable. Optional locking M12 connectors ensure secure coupling in moving, volatile environments.

High-Speed USB Integration

Rugged Edge AI Inference PCs are equipped with several generations of USB connections to accommodate data traffic needs for various peripheral technologies. Systems come configured with four USB 3.2 Gen 1, offering 5 Gbps data transfer speed, and four USB 3.2 Gen 2, offering rapid 10Gbps data transfer speed.

CAN Bus For Vehicle Insights

The RCO-6000-CFL Series supports the CAN Bus Protocol to leverage vehicle telematics data for intelligent transportation systems, fleet management, process analytics, and system optimization.

4 LAN Module

4 LAN & PoE Module

4 Port M12 Module

4 Port M12 & PoE Module

Dual 10G Module

5G Module

Built Rugged. Built Ready.

Wide Operating Temperature Range

The RCO-6000-CFL Series has a wide operating temperature range, ranging from -25 C to 60 C, accommodating a wide range of challenging thermal conditions. Blistering steel foundries and ice-encrusted arctic mine sites can easily dispatch the AI edge inference computer for deep data inference analysis at the rugged edge.

Shock & Vibration Resistance

The RCO-6000-CFL Series is also hardened to endure exposure to impact and vibrations, common in rugged industrial environments. In fact, the system can withstand up to 50Gs of shock and 3GRMs of vibration, enabling deployment in environments where the system is exposed to frequent shock and vibration. AI Edge Inference computers can be deployed in a vehicle to collect and process vast amounts of sensor data that can be used later to train machine learning and deep learning models. Additionally, the data collected can be leveraged for telematics for intelligent fleet management or performing predictive analytics to alert of impending traffic hazards. Moreover, heavy industrial sectors can dispatch the RCO-6000-CFL Series to rough, volatile settings to conduct inference analysis on visual and situational input. Metrological application, quality inspection, and predictive maintenance algorithms make the AI edge inference computer an essential rugged utility for business optimization.

Wide Voltage Input

The RCO-6000-CFL Series is outfitted with a wide voltage input range, ranging from 9 to 48VDC, accepting a wide range of available power voltages. Additionally, the system supports AT/ATX power modes according to deployment requirements.

Power Ignition Management

Intelligent transportation deployment can harness the power of inference analytics safely through the RCO-6000-CFL's power ignition management. The power ignition management feature delays the system shutdown after engine shutoff for a pre-determined, programmable interval to ensure that no data corruption occurs as a result of a sudden and abrupt shutdown. The feature ensures that applications close properly, avoiding data loss or corruption.

Hot-swappable Blower

When the RCO-6000-CFL Series is equipped with an EDGEBoost node, the enclosure housing the NVMe SSD storage drives and/or GPU comes with a hot-swappable blower fan, delivering cooling where it counts. The integrated blower fan is necessary to remove the heat generated from the GPU and NVMe storage devices, neutralizing temperature spikes often experienced from high-performance NVMe and GPU technologies. The hot-swappable nature of the fans makes cleaning them and replacing them super easy and quick to eliminate unwanted downtime.

TPM 2.0 Security

An integrated trust platform module applies the TPM 2.0 standard to safeguard the RCO-6000-CFL Series. The microprocessor's root keys enable password protection, device authentication, and future-ready cybersecurity. The TPM defends the device, data, and transmission against malicious actors.

We Design, Manufacture, and Service customers around the world

NEW

AI Edge Inference Computer




	RCO-6000-CFL-2N-2060S	RCO-6000-CFL-4N-2060S	RCO-6000-CFL-4NH	RCO-6000-CFL-8NS
Processor	Support 8th / 9th Gen. Intel® CFL-R S Processor (LGA 1151, 65W/35W TDP)			
Memory	2x 260-pin DDR4-2400/2666MHz SO-DIMM, up to 64GB (Un-buffered and Non-ECC)			
GPU	RTX 2060S	RTX 2060S		
Display	1x DVI-I, 1x DVI-D, 1x HDMI, 3x DisplayPort	1x DVI-I, 1x DVI-D, 1x HDMI, 3x DisplayPort	1x DVI-I, 2x DisplayPort	
SATA Storage	4x 2.5" SATA HDD bay with RAID 0, 1, 5, 10 support (2x internal; 2x removable & hot-swappable), 1x mSATA (shared by 1x Mini PCIe)			
NVMe Storage	1x NVMe M.2 M Key, 1x Removable Module with 2.5" 2 Bay U.2 NVMe SSD (15mm) *Patented NVMe drives trays are tooless and hot-swappable	1x NVMe M.2 M Key, 1x Removable Module with 2.5" 4 Bay U.2 NVMe SSD (7mm) *Patented NVMe drives trays are tooless and hot-swappable	1x NVMe M.2 M Key, 2x Removable 2 Bay NVMe SSD Module with Hardware RAID 0, 1, 5, 6, 10 support (15mm) *Patented NVMe drives trays are tooless and hot-swappable	1x NVMe M.2 M Key, 2x Removable 4 Bay NVMe SSD Module with RAID 0, 1, 5, 10 support (7mm) *Patented NVMe drives trays are tooless and hot-swappable
Internal Expansion Slot	2x Full-size mini-PCIe (1 shared by 1x mSATA), 1x M.2 E Key			
I/O	4x USB 3.2 Gen 2 (10 Gbps), 5x USB 3.2 Gen 1 (5 Gbps), 2x USB2.0 header (internal), 6x RS-232/422/485 (2x internal), 16x isolated digital I/O, 2x GbE RJ45 (Support Wake-on-LAN and PXE)			
Power	9-48 VDC, AT/ATX Select, 3-pin Terminal Block and 4-pin Terminal Block			
Operating Temperature	-25°C to 60°C (35W /65W CPU)			
Certification	E-Mark, EN 50155, EN 50121-1, EN 50121-3-2			

Premio Inc. 918 Radecki Ct., City of Industry, CA 91748 Toll Free: 800.977.3646 Email: sales@premioinc.com **www.Premioinc.com**

Wide temperature, sustained performance and high capacity define industrial-grade M.2 and U.2 SSDs

Unpredictable operating environments, wide temperature fluctuations, as well as thermal and airflow restrictions pose significant design challenges for small-form factor, high-performance PCIe NVMe SSDs. Thermal dissipation in M.2 is far more challenging than in other form factors due to its compact size. Moreover, the higher the capacity, the more difficult thermal dissipation gets. Premio partners with global storage provider Exascend, integrating its 4TB M.2 that provides 1,450MB/s sustained write speed between -40 and 75°C, and 700MB/s sustained write speed between 75 and 85°C. This unique blend of industrial-grade thermal performance and enterprise-class storage capacity is a flash storage milestone for rugged edge system developers, combining two historically conflicting concepts into a uniquely capable high-capacity industrial flash storage product series.

Exascend storage devices also feature Adaptive Thermal Control™; Exascend's patented algorithm that continuously monitors drive thermals and regulates drive performance to improve SSD stability. Adaptive Thermal Control™ provides a ten-time enhancement of sustained performance and quality of service (QoS), and an up to three-time improvement of minimal performance. It brings an improved system thermal profile, minimizes the airflow and acoustic design complexity, and reduces the overall BOM cost – all while improving the long-term drive reliability and lifespan. Additional features include accelerated boot, which streamlines and layers the initialization by prioritizing system response and allowing background error correction and recovery.

About Exascend PI3 SSD

Exascend's PI3 series SSD is an industrial-grade (-40 to 85°C) PCIe Gen3x4 product with enterprise performance features (high sustained speed and low latency), available in standard M.2 2280 and U.2 form factors with up to 8TB capacity. The company's ultra-high capacity, high-speed PCIe Gen3 NVMe SSD products are based on Marvell's industry-leading PCIe NVMe SSD controller through a unique strategic partnership.

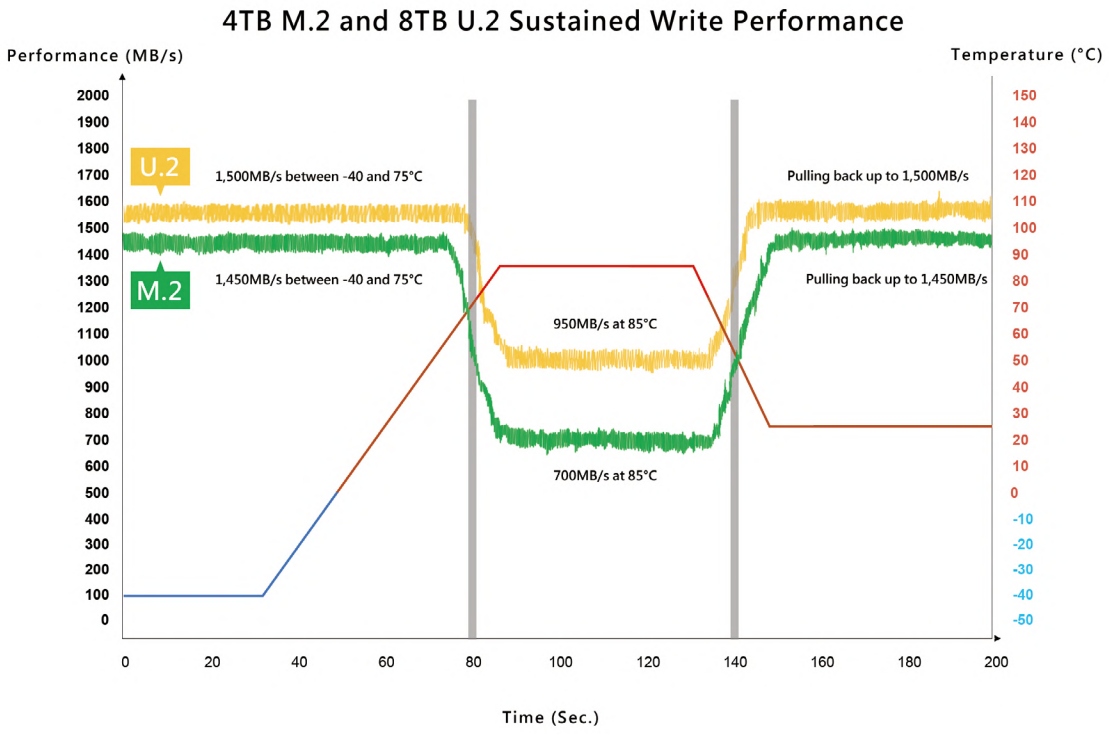

Part 3 Real-World Applications that Can Benefit

Real-World Application that Can Benefit

This design strategy defines a uniquely modular system, combining high-performance in compute, storage, and connectivity with rugged, fanless design. Real-world deployments best illustrate their competitive value in heavy industrial environments.

Inference analysis in real time, eliminating the challenges of bandwidth, latency, and cloud resources

Rugged edge systems enable smart 'object detection' software to process complex visual tasks more efficiently than the human eye. Based on deep learning, where neural networks are trained through repeated exposure to various data sources, this is a new level of advanced automation now available to the most challenging computing environments and applications. Historically the compute power necessary to aggregate and analyze this scope of data was relegated to cloud resources, and therefore unfeasible in many types of industrial deployments. Today, rugged edge systems are physically closer to their data sources, improving system responsiveness and protecting sensitive data with on-site processing. These high-performance systems tap into programmed sensors, software algorithms, and both CPU and GPU computing power capable of analyzing images at high speed, short intervals. Coupling these performance accelerators with high speed PCIe lanes for NVMe SSD data transfers provides results quicker and more accurately than plant personnel, improving productivity, efficiency, and market leadership.

Flexible data transfer options drive performance in telematics and transportation

On-vehicle sensors collect a monumental amount of data, poised to expand further with 5G capabilities. Backed by AI, this data transforms into learning that creates significant competitive impact, improving the accuracy of asset tracking, productivity, driving behaviors, and overall operations and costs. The industry also uses edge devices to manage its mobile workforce, with devices and data considered key to managing operations and assets. Flexibility in data handling is critical to intelligent fleet management, given the diversity of data, sensors, vehicles, and environments. For example, where vehicles are constantly in motion, robust connectivity is required for passing telemetry data back and forth through wireless links. Rugged edge systems with uniquely modular, hot swappable capabilities, such as those mounted inside trucks or other commercial vehicles, provide another alternative in how

data is moved from the vehicle and into storage or other master data management facilities. Once the vehicle returns to base, drives are easily unplugged and swapped, reducing downtime and keep vehicles on the road. Remote and mobile applications focused in vehicle telematics and transportation require high-capacity storage drives that can withstand environmental factors like wide-temperature, shock, and vibration. Therefore, solid-state-drives (SSDs) that do not use any spinning disks are great candidates in many transportation focused applications at the edge. Especially as the overall cost/GB on solid state NAND pricing is more affordable at scale, high-capacity enterprise SSDs help store mission critical data at the edge, without requiring an abundance of SSDs.

Intelligent vending taps workload consolidation to enable a full-featured customer experience

In smart retail, rugged edge technology can be used for anything from logistics and inventory management to forecasting and personalized shopping experiences. The advent of AI is transformative in this market, proving that 'building' intelligent stores – both online and brick and mortar – provides the foundation retailers need to compete and thrive in an industry where the customer is king. Self-service kiosks deployed in remote and rugged locations feature resistance to shock, vibration, and extreme temperatures in a balance of modular and semi-fanless designs. Workload consolidation blends the right combination of compute, storage, connectivity, and I/O variety to ensure efficient and reliable remote manageability at the edge.

Surveillance wins with fast data performance in a compact footprint

High-performance computing in public venues keeps infrastructure facilities and the people who visit them safe. Arenas, airports, and train stations – here, AI powered by rugged edge computing adds deep value to public safety, enabling systems to analyze live video without human intervention. Systems are both powerful and compact, validated to withstand damaging environmental elements like

wide temperature (-40°C to 70°C), wide voltage and power protection (9VDC – 50VDC), and shock and vibration (5GRMS vibrations and 50G Shock). These ruggedized features, along with support for a PCIe Gen 3 interface for the high speed and efficient NVMe protocol, allows embedded system developers to design applications for data logging, monitoring, and surveillance in environments mobile, harsh, and remote. These systems recognize a performance gap in SATA storage technologies and instead build on the best practices of high-performance data center design to achieve fast, enterprise-class low power edge performance.

Industrial automation is ripe for advancements in AI and machine learning

Today's manufacturers are embracing Industry 4.0 in a big way, deploying a vast range of electronic systems and devices featuring smart, connected operations. Machinery and equipment are now enabled with much more functionality, for example, advanced cameras with motion detection, depth perception, and signal transmission used to empower autonomous processes. Rugged edge computers harness and unlock the massive data streams from these devices and sensors – powering the AI and machine learning so critical to streamlining processes and advancing automation. Rugged edge nodes, tasked with analyzing data and acting in real-time, optimize workloads to improve quality control, accelerate production, maintain equipment, speed time to market, and meet, or even exceed customer expectations in the process. Workload consolidation also provides a strong foundation for mobile and remote deployments, benefiting from mission-critical reliability and its ability to reduce or eliminate the need for on-site technical support in remote applications and geographies.

Rugged performance, real-time benefits

Leveraging a slate of proven machine learning technologies, Premio's rugged edge systems enable AI inference computing. Ruggedized for industrial performance and deployed in more rigorous physical environments, these low latency systems avoid data bottlenecks and deliver real-time analytics with the power to change and improve how an organization operates.

Industrial leaders can now apply the proven NVMe protocol and glean its advantages via PCIe in edge level deployments. Premio furthers this value with multiple NVMe M.2 drives and its proprietary carrier board design, enabling easy data access, hot swappability, and non-stop performance.

Leveraging the deluge of data for the win

Deep, real-time data is foundational to AI; this means the rugged edge is everywhere data is. In this data-rich world, Gartner estimates that connected edge devices will grow to 25 billion units by 2021.2 Not surprising, given the framework's applicability to a wide variety of environments.
Hardware intervention – newly enabled by rugged edge systems – empowers all types of industrial organizations to harness, learn from, and use AI. Distributed systems, purpose-built for rigorous environments, can now effectively capitalize on rich insight and performance enhancements for greater quality control, safety, security, and efficiency.

Real-time decision making is empowered when these capabilities are localized – tapping into faster storage, faster network speeds, faster I/O peripherals in a proven, rugged system. NVMe adds critical value via real-time processing at the rugged edge. Always learning, data-driven environments can adapt quickly, reacting to new information and garnered intelligence. AI permeates the rugged edge's potential, offering new ways of using and benefiting from data in a spectrum of applications. And, as costs per NAND flash GB decrease, organizations have an increasingly viable path to leverage rugged edge computing to competitive advantage.

Smart system and mechanical design are bringing AI to life in new and ever more challenging industrial settings. To begin an engineer-to-engineer conversation with Premio and learn about the rugged edge as a new business imperative, connect via email **sales@premioinc.com** Consultation with our IoT deployment experts is free, and we're at the ready to help you maximize data storage and processing at the rugged edge.