



Product Brief AI Edge Inference Computer

- RCO-6000-CFL-2
- RCO-6000-CFL-4N-2060S
- RCO-6000-CFL-2N-2060S
- RCO-6000-CFL-2060S
- RCO-6000-CFL-8NS
- RCO-6000-CFL-4NH
- RCO-6000-CFL-4NS



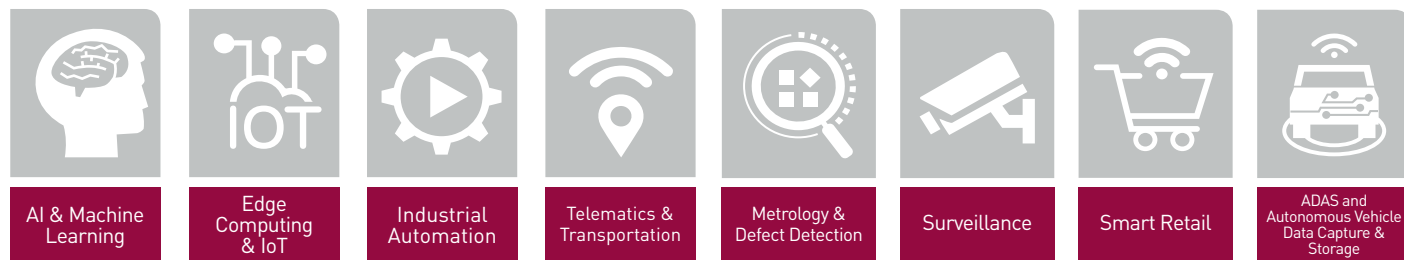
Compute, Storage, Connectivity, Ruggedness: Necessary Ingredients for Inference

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





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 five types of Edge Boost Nodes that can expand the capabilities of the RCO-6000-CFL Series AI Edge Inference Computer:



EdgeBoost Node # 1 RCO-6000-CFL-2N-2060S

The first edge boost 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. The edge boost node also adds PCIe expansion slots, enabling organizations and system integrators to add an Nvidia 2060 Super GPU for inference acceleration.



EdgeBoost Node # 2 RCO-6000-CFL-4NH

The second 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.



EdgeBoost Node # 3 RCO-6000-CFL-4NS

The third EDGEboost node focuses on ultra-high-speed NVMe Storage and offers users the ability to add up to 4x lockable and hot-swappable 2.5" U.2 NVMe SSDs in 15mm 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.



EdgeBoost Node # 4 RCO-6000-CFL-4N-2060S

The fourth edge boost 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. The edge boost node also adds PCIe expansion slots, enabling organizations and system integrators to add an Nvidia 2060 Super GPU for inference acceleration.



EdgeBoost Node # 5 RCO-6000-CFL-8NS

The fifth 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.



EdgeBoost Node # 6 RCO-6000-CFL-2060S

The sixth EDGEboost Node offers organizations with GPU acceleration via Nvidia's RTX2060 Super Graphics Processing Unit. The Nvidia 2060S sits comfortably in a PCIe Gen 3 x16 Slot, offering organizations with robust and powerful inferencing capabilities at the rugged edge.



EdgeBoost Node # 7 RCO-6000-CFL-2

The seventh EDGEboost Node provides organizations with expansion capabilities via the availability of two expansion PCIe/PCI expansion slots. Expansion slots allow users to populate the system with add-in cards for industrial computing workloads. The EDGEboost nodes provide three expansion capabilities:

- Option 1 : 1x PCIe x16 and x1 PCI Expansion Slot
- Option 2 : 2x PCIe x8 Expansion Slots
- Option 3 : 2x PCI Expansion Slots



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.

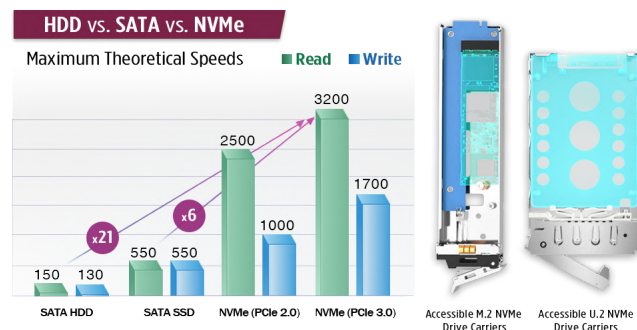
High-Speed Storage Primed for Edge Intelligence

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 edge boost 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



RICO-6000-CFL-2N-2060S



RICO-6000-CFL-4NH



RICO-6000-CFL-4NS

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		
Display	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 toolless 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 toolless and hot-swappable	1x NVMe M.2 M Key, 2x Removable 2 Bay NVMe SSD Module with RAID 0, 1, 5, 10 support (15mm) *Patented NVMe drives trays are toolless 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, EMC Conformity with EN 50155, EN 50121-3-2		

WE DESIGN,
MANUFACTURE, AND
SERVICE CUSTOMERS
AROUND THE WORLD



NEW

AI Edge Inference Computer



RCO-6000-CFL-4N-2060S



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	
Display	1x DVI-I, 1x DVI-D, 1x HDMI, 3x 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" 4 Bay U.2 NVMe SSD (7mm) *Patented NVMe drives trays are toolless 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 toolless 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, EMC Conformity with EN 50155, EN 50121-3-2	

WE DESIGN,
MANUFACTURE, AND
SERVICE CUSTOMERS
AROUND THE WORLD



NEW

AI Edge Inference Computer



RCO-6000-CFL-2060S



RCO-6000-CFL-2

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	
Display	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	
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), 1x USB 3.2 Gen 1 header (5 Gbps, 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)	-25°C to 70°C (35W CPU) ; -25°C to 60°C (65W CPU)
Certification	E-Mark, EMC Conformity with EN 50155, EN 50121-3-2	