

**WHITE PAPER**

Faster pathways to support increasing,  
critical data at the edge



**JUSTIN MYERS**

Principal Product Manager  
Mercury Systems

**VINCENT CHUNG**

Principal Product Manager  
Mercury Systems



**Today’s modern defense systems exploit massive volumes of data to provide warfighters with more accurate, complete intel.** Examples include advanced radar object detection and AI-driven threat identification. Solving bandwidth challenges at the system interconnect helps ensure mission success.

Commercial semiconductor innovators are continuously developing new capabilities in silicon that defense system designers leverage to build open architecture-based processing solutions. When intelligently packaged for reliable deployment at the edge, the latest data center technologies can advance mission capabilities for the defense and aerospace sector.

This paper examines the critically important interconnect technology and ecosystem built around Peripheral Component Interconnect Express (PCIe) standards and specifically, how the emerging PCIe 5.0 protocol benefits a new generation of rugged, deployable application platforms. Gen 5.0 doubles the bandwidth over the Gen 4.0 spec, which similarly doubled bandwidth over ubiquitous Gen 3.0 devices, and is an integral part of the latest devices from industry giants Intel® and NVIDIA.

**THE CURRENT NEED FOR MORE INTER-COMPONENT BANDWIDTH**


The nature of modern defense systems makes interconnect bandwidth increasingly critical to overall system performance. Field-deployed servers receive inputs from multiple types of sensor systems, including radar, lidar, visual imagery, EO/IR and EW. They also need to access stored data for comparison and analysis purposes; for example, identifying artifacts in a current image that were not present in yesterday’s image.

Each sensor interface is a data path that must be able to keep pace with the data source. Then, to process those data streams, the server will typically employ multiple components of various types – CPUs, GPUs, FPGAs and DSPs – with interactions and data movement between them. Moving data quickly from one component to the next is critical to overall server performance. At a higher level, deployed servers must communicate with each other and with displays that form the human interface.

**RES X08 RUGGED RACK SERVERS**

These latest-gen rugged rack servers optimize the latest PCIe 5.0 compute, memory, storage and networking to eliminate bottlenecks, support hybrid computing architectures and drastically speed compute-intensive applications like sensor processing and artificial intelligence (AI), all in a smaller or equivalent form factor.

RES X08 servers integrate the latest commercial technologies from Intel and NVIDIA including data center-class 4th gen Intel Xeon Scalable processors (formerly Sapphire Rapids), low-latency PCIe 5.0 fabrics, NVIDIA H100 GPUs, high-speed DDR5 memory and Compute Express Link (CXL) to infinitely scale system memory.



RESX08-2U18RA Rugged Rack Server



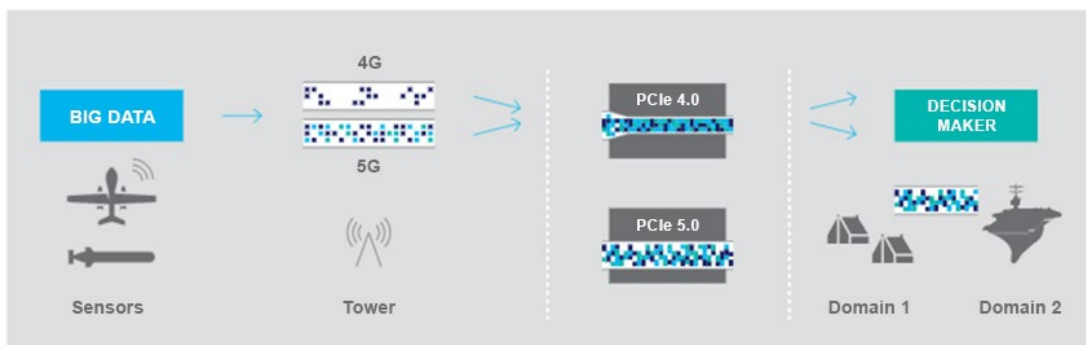


AI applications need low-latency access to a range of data types, coming from multiple sensors on multiple platforms.

### AI WILL DRIVE DEMAND FOR EVEN MORE BANDWIDTH

Looking beyond current requirements, deployed servers will soon need to host AI workloads, as defense systems use deep learning and inference engines to extract critical information from a torrent of incoming sensor data.

Applying AI effectively to real-time applications will demand high-bandwidth, flexible data movement without bottlenecks. This need is illustrated by an inference engine that is working backward from an assumption to see if there is any available data that supports the assumption. Low-latency access to a range of data types is essential for effective use of this technique in real-time systems.



PCIe 5.0 offers 2x the data throughput for application efficiency.

## PCIe'S ROLE IN DEFENSE SYSTEMS

PCIe is the computing industry's most common interface standard for connecting high-speed components. In the commercial world, virtually every PC motherboard provides PCIe slots. PCIe is also an essential communication interface in commercial servers that support data centers and cloud applications. In these servers, PCIe is the primary data path from CPUs to both storage devices and to auxiliary coprocessing components, such as GPUs.

The defense electronics industry adapts commercial technology to meet its needs, exploiting the continual performance increases driven by enormous market forces. As such, rugged server architectures are analogous to commercial server architectures, with PCIe occupying a central role in moving sensor data to archive storage, retrieving stored data for analysis and moving data to, and between, specialized computing components.

PCIe 3.0 was a mainstay in rugged servers over the past decade; the majority of modern servers support Gen 3.0. However, expanding volumes of sensor data have pressured users to consider more capable protocols, such as PCIe 4.0, now implemented in most new deployed systems.

While Gen 4.0 delivered twice the bandwidth of Gen 3.0, it is expected to be short-lived. Sensor data volumes continue to increase exponentially, while AI-based applications, with even greater bandwidth demands, are moving into active use.

Another factor is the increasing speed of network connections, the primary mechanism for intersystem data communication. Ethernet is the universal fabric connecting data centers and rugged deployed systems. As Ethernet evolves into a terabit/s capable protocol, architects can expand the data paths between individual systems that may have previously been bottlenecks.

Fortunately, the larger world of commercial electronics faces a similar set of challenges and is already moving forward with the greater bandwidth of PCIe 5.0.

## A STANDARDS-BASED APPROACH TO RAPID GEN 5.0 INTEGRATION

Defense system architects need a consistent pathway for technological upgrades, a way to rapidly integrate Gen 5.0 devices as they become available. To meet that need, Mercury is embracing the Open Compute Project (OCP) OpenEdge architecture.

We will soon be introducing a portfolio of rugged, deployable "plug-and-pull" blade products based on that architecture, as an evolution to the highly successful RES HD server product line. Highly configurable, these products will exploit the snowballing ecosystem of OCP-based hardware, an ecosystem also implementing PCIe Gen 5.0. Building on an open standards-based platform will allow our customers to deploy rugged, Gen 5-capable servers without vendor lock-in.

Mercury is committed to making our OCP-based platforms the fastest path to deploying PCIe Gen5 technologies on the battlefield.

## MOVING TO PCIe 5.0

PCIe 5.0 delivers twice the data throughput of Gen 4.0, reaching an effective bandwidth of 64 GB/s for a 16 lane (x16) link. While the Gen 5.0 specification was released in 2019, it has taken time for peripherals and semiconductors supporting this new spec to be developed. The first such devices were launched in early 2022, offering impressive capabilities.

Two industry giants were among the first to bring Gen 5.0 products to market. Intel supports Gen 5.0 with its 12th Gen Core™ platforms and will soon be shipping even higher-performance Gen 5.0 CPUs in the Xeon® Scalable Processor™ line aimed at server systems. In parallel, NVIDIA is introducing new generations of GPUs and DPUs that use Gen 5.0 for both internal and external data paths.

## THE EMERGING GEN 5.0 ECOSYSTEM

It is this Gen 5.0 ecosystem of devices that will enable the next generation of high-performance servers for both commercial data centers and deployed defense applications.

**intel.** **Intel** - Intel Xeon Scalable Processors, formerly codenamed Sapphire Rapids, support leading industry standard technologies, including PCIe 5.0 as well as Compute Express Link™ (CXL) 1.1 to expand memory using PCIe slots. At 64 GB/s, Xeon Scalable PCIe gen-5 bandwidth not only helps today's applications run faster, it also enables new applications in artificial intelligence (AI)/machine learning (ML) and cloud computing. With AI permeating more and more workloads, Intel has committed to deliver a 30x total AI performance gain through hardware and software optimizations for AI workloads.

**NVIDIA.** **NVIDIA** - NVIDIA is implementing PCIe 5.0 interconnects in its latest generations of both GPUs, DPUs and network interface modules.


- **GPUs:** The new NVIDIA H100 Tensor Core GPU has 32 lanes of PCIe 5.0 communications for 128 GB/s of data bandwidth. The H100 also delivers a huge leap in general purpose GPU (GPGPU) compute performance over the previous generation A100 GPU; fast Fourier transform (FFT) execution is six times faster for HPC applications, while AI apps are accelerated by factors ranging from nine to thirty.
- **DPUs:** In a complementary product line, the recently announced NVIDIA BlueField®-3 DPU has the same 32 lanes of PCIe 5.0 to drive 128 GB/s of bandwidth, so NVIDIA GPUs and DPUs can be configured together without bottlenecks. The BlueField-3 also has a PCIe switch supporting up to 16 downstream ports. The third-generation BlueField-3 delivers a five-times higher SPECINT score than the previous BlueField-2 device.
- **Networking:** Using technology acquired via its purchase of Mellanox, NVIDIA is also shipping the ConnectX®-7 SmartNIC, which delivers 100 GbE, 200 GbE, and 400 GbE Ethernet networking with a PCIe 5.0 x16/x32 interface.

**Other vendors** - Intel and NVIDIA are not the only vendors bringing Gen 5.0 products to market. Astera Labs, Broadcom and AMD are just a few of the other companies building Gen 5.0 support into their newest solutions.

## DEPLOYABLE GEN 5.0 SERVERS FROM MERCURY SYSTEMS

Mercury's lightweight and compact rugged rack servers employ cutting-edge commercial technology, including Intel processors, NVIDIA GPUs, PCIe 5.0 and 200-400 Gb Ethernet networking, to tackle the most challenging workloads in the most remote environments - where the mission is critical and the cost of failure is high.

Mercury has expanded its rugged server lines, with capabilities beyond the PCIe 4.0 architectures offered in RES X07 Intel-based rack mount servers and the RDP B02 NVIDIA GPU/DPU-based server line, adding a new generation with a new generation of deployable RES X08 servers that support double the interconnect bandwidth. These will be based on PCIe 5.0 and exploit the powerful capabilities of Intel's third generation of Xeon Scalable Processors, NVIDIA's H100 Tensor Core GPUs, NVIDIA's ConnectX-7 NICs, and NVIDIA BlueField-3 DPUs. The new servers will deliver advanced data center-caliber performance to the tactical edge, eliminating latency bottlenecks and enabling compute-intensive applications to operate with real-time response.



RES X08 servers incorporate the latest commercial technologies, ruggedized for the harshest environments.

**MORE DETAILS ON THE PCIE STANDARD**

PCIe is a standard defining communication between two devices over a point-to-point interconnect, or link, with each link having one or more lanes. Every lane has two signaling pairs, one pair for receiving and one for sending; that means every lane consists of four physical wires, or semiconductor traces.

The PCIe standard defines links with 1, 4, 8, 16 or 32 lanes, though 32-lane links are rare. The lane count is written with an ‘x’, so a 16-lane link is ‘x16.’ Lane count is often called ‘width,’ as in ‘16 lanes wide.’

There are two different ways of measuring the performance of a specific link. One way uses the raw-bit transfer rate, usually expressed in terms of ‘Gigatransfers per second’ (GT/s). The other, more useful measurement reflects the effective transfer rate of actual data, or bandwidth. The data bandwidth is less than the raw transfer rate because a portion of the bits transferred are used for ensuring data integrity (encoding) and for other communication overhead tasks. The chart below shows these two measurement values for PCIe versions 1 through 6.

Generation	Bandwidth (x16)	Gigatransfer per lane
PCIe 1.0	4 GB/s	2.5 GT/s
PCIe 2.0	8 GB/s	5 GT/s
PCIe 3.0	16 GB/s	8 GT/s
PCIe 4.0	32 GB/s	16 GT/s
PCIe 5.0	64 GB/s	32 GT/s
PCIe 6.0	128 GB/s	64 GT/s

As evidenced by this chart, the available data bandwidth doubles for each successive generation.

And while PCIe 5.0 backplanes are backwards compatible with PCIe 4.0 devices and earlier devices, PCIe 5.0 devices must be used at both ends of a communication path in order to achieve full bandwidth afforded by the technology.

**ENGAGE WITH MERCURY TO IMPLEMENT THIS NEXT-GEN TECHNOLOGY**

Start architecting your next generation of upgradable defense systems, based on the highest-bandwidth platforms available and using an ecosystem of the latest semiconductors.

Our close collaboration with commercial technology leaders in the semiconductor industry means latest-generation innovations, like PCIe 5.0, can be rapidly integrated to push server-class processing to the tactical edge.

To ensure rapid, cost-effective delivery of deployable solutions, Mercury has invested in scalable, U.S. manufacturing operations—a valuable trusted and secure section of the supply chain.

We also see close cooperation with defense electronics visionaries, program managers and engineering teams as key to achieving maximum value for your solutions. Our system experts are focused on understanding your applications to develop an optimized system solution.

Engage with the Mercury team to see demonstrations of our PCIe 5.0 technology and explore how we can best meet your deployed server requirements.

## About Mercury

Mercury Systems is a technology company that makes the world a safer, more secure place. We push processing power to the tactical edge, making the latest commercial technologies profoundly more accessible for today's most challenging aerospace and defense missions. From silicon to system scale, Mercury enables customers to accelerate innovation and turn data into decision superiority. Headquartered in Andover, Massachusetts, Mercury employs more than 2,300 people in 24 locations worldwide.



### Corporate Headquarters

50 Minuteman Road  
Andover, MA 01810 USA  
+1 978.967.1401 tel  
+1 866.627.6951 tel  
+1 978.256.3599 fax

### International Headquarters Mercury International

Avenue Eugène-Lance, 38  
PO Box 584  
CH-1212 Grand-Lancy 1  
Geneva, Switzerland  
+41 22 884 5100 tel

Learn more about RES X08 Rugged  
Rack Servers: [mrcy.com/rack-servers](https://mrcy.com/rack-servers)

Contact us to collaborate on your rugged  
server solution: [servers@mrcy.com](mailto:servers@mrcy.com)



The Mercury Systems logo and the following are trademarks or registered trademarks of Mercury Systems, Inc.: Mercury Systems, Innovation That Matters, and BuiltSECURE. Other marks used herein may be trademarks or registered trademarks of their respective holders. Mercury believes this information is accurate as of its publication date and is not responsible for any inadvertent errors. The information contained herein is subject to change without notice.