Special Report

Perspectives from HPE's Slingshot Innovation Forum at SC25

Mark Nossokoff and Earl Joseph
May 2026

HYPERION RESEARCH OPINION

With systems supporting thousands of scientists and researchers running diverse applications simultaneously, optimizing advanced technical computing infrastructure for traditional HPC scientific workloads has always been a challenge. Introducing modern AI training and inferencing workloads into the mix has dramatically exacerbated the situation. Architecting a balanced system performance between compute, storage, and networking to prevent any one resource from sitting idle and being underutilized, especially expensive GPUs, is a key priority for system architects, with interconnects factoring prominently.

With ISC26 fast approaching, it seems appropriate, if not exactly timely, to spotlight one last item from SC25. Highlighting its ethernet-based Slingshot interconnect technology, HPE sponsored a Slingshot Innovation Forum at SC25, inviting Hyperion Research to contribute an overview of the market to kick it off. This paper presents key topics from the forum, along with additional HPC and AI market perspectives. Providing an initial glimpse into the over HPC and AI interconnect adoption trends, Figure 1 summarizes interconnect technology adoption within the Top500 top 20 systems since June 2020.

FIGURE 1

Top500 System Performance by Interconnect Family

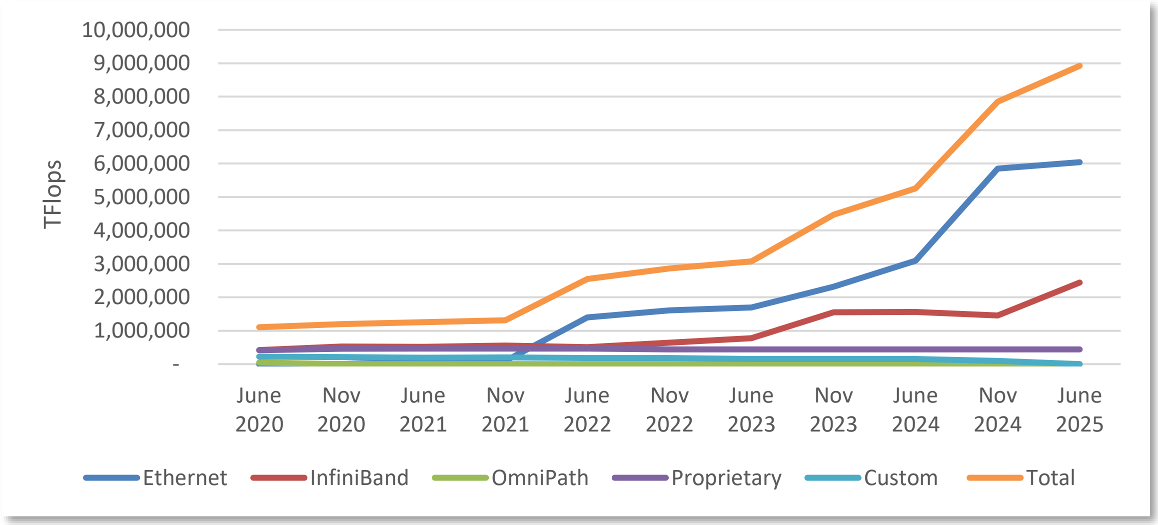

Source: Hyperion Research, 2026

OVERVIEW OF HPC MARKET AND TRENDS

On-Premises HPC Market

The on-premises HPC market continues to project healthy growth. The forecast $60B total HPC market in 2025 represents an annual growth of 23.4% in on-premises servers and 20.6% in on-premises storage. Global plans for more deployment of exascale-class systems and the buildout of AI factories are contributing to much of the growth projections over the forecast period. Interconnect contributions to market growth are included in both the Server and Storage broader market areas, with interconnect revenues historically accounting for 12%-15% of the on-premises HPC server revenues. Table 1 summarizes the forecast.

TABLE 1

On-premises Revenues for the Broader HPC Market Areas ($M)

	2024	2025	2026	2027	2028	2029	2030	CAGR '25-'30
Server	$25,557	$30,393	$34,741	$39,229	$43,900	$48,909	$53,318	11.9%
Storage	$11,015	$13,284	$15,282	$17,378	$19,470	$21,725	$23,693	12.3%
Middleware	$2,875	$3,360	$3,815	$4,306	$4,816	$5,363	$5,850	11.7%
Applications	$8,175	$9,564	$10,822	$12,151	$13,521	$14,980	$16,336	11.3%
Service	$3,215	$3,605	$3,860	$4,304	$4,756	$5,234	$5,709	9.6%
Total Revenue	$50,837	$60,205	$68,520	$77,367	$86,463	$96,210	$104,906	11.7%

Source: Hyperion Research, February 2026

Interconnect Adoption

There are a number of system interconnect technologies available for leadership-class advanced computing architectures, including ethernet, InfiniBand, OmniPath, and Bxi. There are also numerous segmentation methods to compare how each is adopted in the market. One method is to look at the aggregate performance supported by each technology within the Top500 list of global supercomputer deployments. Figure 2 summarizes interconnect technology adoption within the Top500 top 20 systems since June 2020, highlighting when three exascale systems based on HPE's Slingshot interconnect were fully accepted for production workloads.

FIGURE 2

Top500 Top 20 System Performance by Interconnect Family

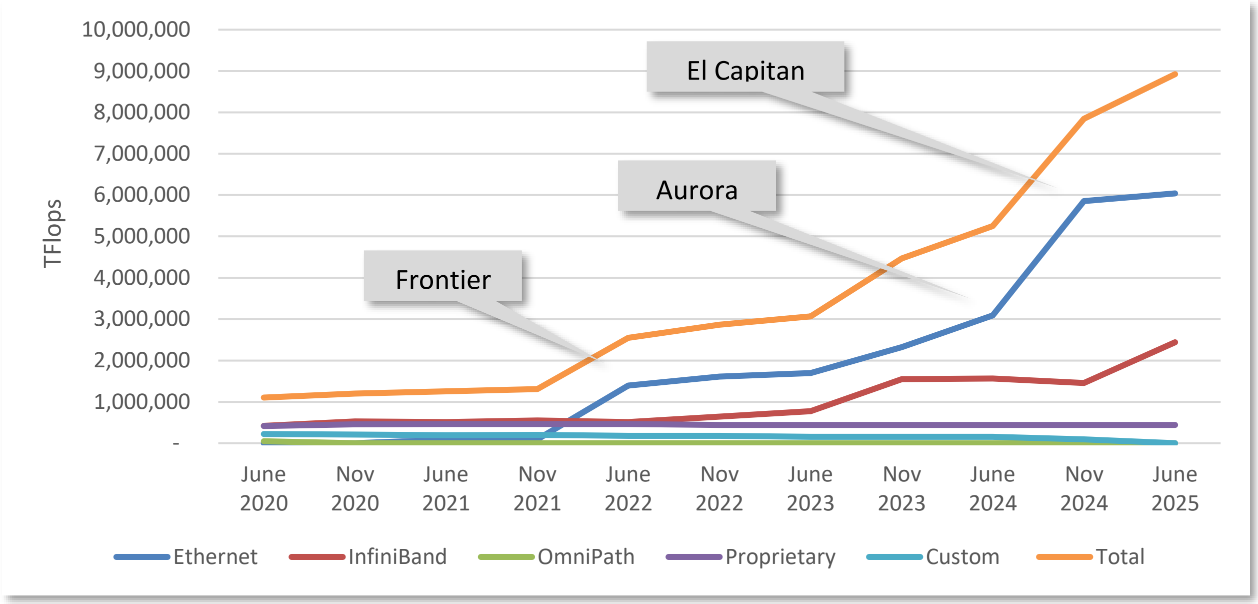

Source: Hyperion Research, 2026

INTERCONNECT ARCHITECTURES - ONE SIZE DOES NOT FIT ALL

I/O Profiles

On the surface, system interconnects all do the same thing – connect various elements within the overall system architecture such that data is provided to the requesting element or delivered to the target element as applications and workloads warrant. However, many application profile parameters impact the performance and feature requirements of the interconnect, particularly when looking at the AI data pipeline. Each step throughout the pipeline places a diverse range of requirements on the interconnect across a broad set of parameters. Table 2 shows the diversity of I/O profiles across the AI data pipeline.

TABLE 2

Diversity of HPC and AI I/O Profiles

Parameter	Ingest	Prep (ETL*)	Train	Checkpoint	Inference
Access Pattern	Sequential	Sequential or Random	Random	Sequential	Sequential
Access Type	Writes	Reads and Writes	Reads	Writes	Reads

TABLE 2

Diversity of HPC and AI I/O Profiles

Parameter	Ingest	Prep (ETL*)	Train	Checkpoint	Inference
Access Frequency	Idle → Intense	Moderate	Idle → Intense	Idle → Intense	Moderate to Intense
Data Size	Small to Large	Small to Large	Mostly Small	Small to Large	Small to Large
Locality	Edge	Edge, Cloud, On-premises	Cloud, On-premises	Cloud, On-premises	Edge, Cloud, On-premises

*ETL = Extract, Transform, Load

Source: Hyperion Research, February 2026

Other factors impacting system and interconnect feature and performance requirements:

- Training frequency (new foundation, RAG, pre-trained)
- Model type and size
- Data type (structure, unstructured; file, block, object)
- Data mode (text, image, video)
- Security (encryption, authentication)
- Compliance (what data to save and for how long)
- Sovereignty
- Physical transport (copper, optical)

Interconnect Hierarchy

Modern architectures employ a complex hierarchy of interconnect architectures throughout the system. Whether connecting chiplets within a semiconductor device to interfacing system elements within and between racks to bridging between datacenters, different interconnect technologies provide a diverse set of options for networking within each hierarchical level.

At a high level, the networks are loosely defined by three categories:

- Scale across – Connects independent datacenters, whether on the same campus or across geographically dispersed regions.
- Scale out – Connects elements within a datacenter. The elements can be rack to rack within a data center or between elements within a rack (e.g., compute nodes, storage nodes)
- Scale up – Connects elements within a node in a rack (e.g., CPUs, GPUs, memory)

Figure 3 defines the interconnect hierarchy levels and lists interconnect technologies employed at each level.

FIGURE 3

Hierarchy of Interconnects

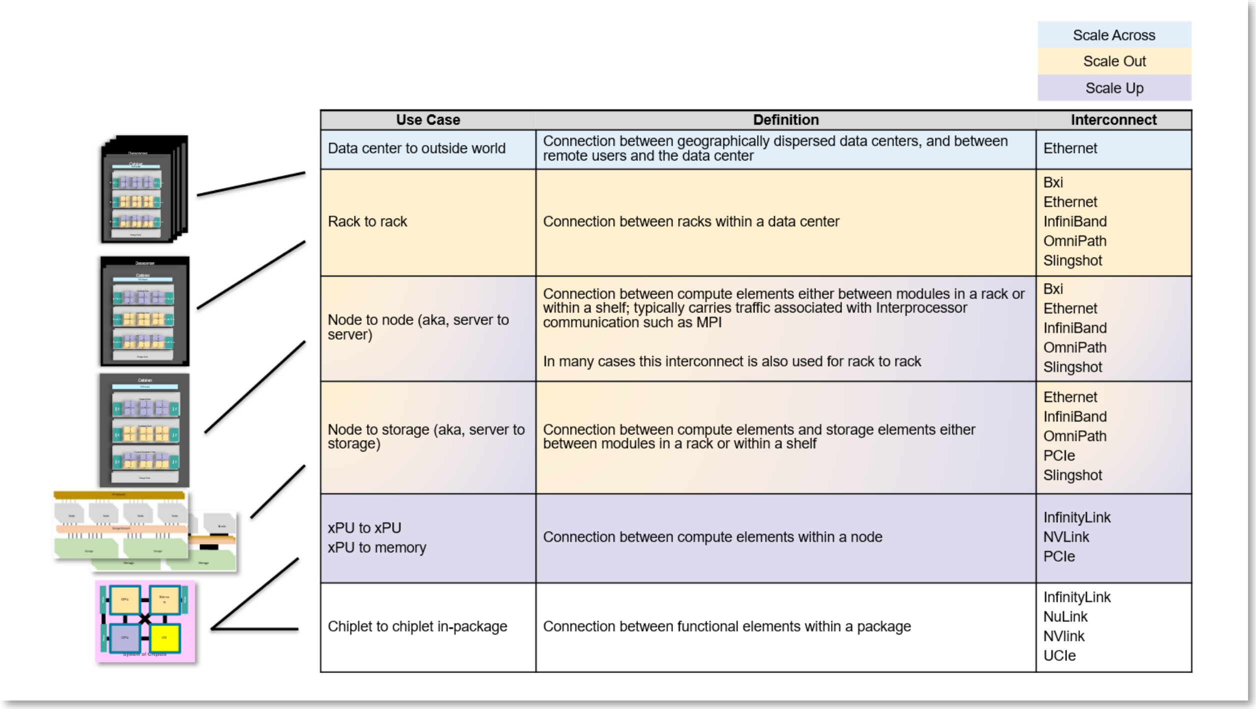

Use Case	Definition	Interconnect
Data center to outside world	Connection between geographically dispersed data centers, and between remote users and the data center	Ethernet
Rack to rack	Connection between racks within a data center	Bxi Ethernet InfiniBand OmniPath Slingshot
Node to node (aka, server to server)	Connection between compute elements either between modules in a rack or within a shelf; typically carries traffic associated with Interprocessor communication such as MPI In many cases this interconnect is also used for rack to rack	Bxi Ethernet InfiniBand OmniPath Slingshot
Node to storage (aka, server to storage)	Connection between compute elements and storage elements either between modules in a rack or within a shelf	Ethernet InfiniBand OmniPath PCIe Slingshot
xPU to xPU xPU to memory	Connection between compute elements within a node	InfinityLink NVLink PCIe
Chiplet to chiplet in-package	Connection between functional elements within a package	InfinityLink NuLink NVlink UCIe

Source: Hyperion Research, 2026

HPE SLINGSHOT

Slingshot has evolved from a proprietary Cray interconnect into a full featured, Ethernet based fabric engineered for converged HPC-AI workloads. HPE is positioning Slingshot not only as a high performance network, but also as a holistic interconnect platform combining hardware, software, congestion control, and RDMA innovation. Key elements of Slingshot highlighted by HPE at SC25 are further described here.

Hardware Architecture and Packaging

HPE has taken a vertically integrated approach to the hardware, including:

- Purpose built silicon: switch ASIC (Rosetta) and Network Interface Card (NIC) ASIC (Cassini)
- NICs: PCIe Gen5 and mezzanine cards
- Switch: integrated switch blade

While each of these elements are interoperable with other Ethernet vendor's components, the vertical integration into a complete system may give HPE more control over latency, congestion behavior, and RDMA semantics than Ethernet solutions built solely on other's silicon.

Software Ecosystem and Operational Tooling

HPE is pulling Slingshot under the umbrella of its broader system software portfolio:

- Slingshot AIOps for application-level performance insights
- Slingshot Monitoring Software for real time fabric health

HPE revealed several dashboards for fabric performance insights – by application or use," and real time fabric health & troubleshooting insights.

This aligns with a broader industry shift toward intent based, telemetry driven networking, critical for AI workloads where tail latency and congestion events can severely impact multi hour training runs.

Middleware Compatibility

HPE has invested and tested to ensure that Slingshot is compatible and interoperable with the leading HPC and AI middleware software stacks including:

- Parallel programming libraries (e.g., Cray MPI, SHMEM)
- Open MPI packages (e.g., Intel MPI, OpenMPI, MPICH)
- GPU communication libraries (e.g., NCCL, RCCL)
- Parallel file systems (e.g., Lustre, Storage Scale [formerly GPFS], DAOS)

Key Features

HPE highlighted several key technical features they identify as key differentiators to InfiniBand, OmniPath and standard ethernet.

Lossless Ethernet and reliability

Optional link-level retry and credit-based flow control, combined with low port switching latency and packet metering (credits) avoid pause-based collapse and expensive job stalls from lost packets — a critical capability for long, tightly synchronized HPC jobs and for checkpoint/IO heavy AI workflows.

In-network congestion management

Slingshot implements per-flow, in-network mechanisms that detect "offenders" vs "victims," apply stiff backpressure at sources, and dynamically re-routes around congested mid-fabric links. HPE claims this to be up to 10× faster control of tail latency compared with round-trip endpoint approaches. This is important for protecting GPU utilization and avoiding cascades (RAID rebuilds, GC, power throttling).

Connectionless RDMA transport

The Slingshot NIC transport avoids per-process qpair state explosion by using dynamically created, lightweight flow states. This reduces NIC memory requirements, avoids ordering constraints tied to qpairs, supports out-of-order delivery options, and exposes extensive NIC HW offloads for messaging primitives. The benefits include smaller NICs, simpler memory footprints, and better scalability in multi-tenant and large-process scenarios.

Operational ecosystem

Libfabric support, MPI/SHMEM stacks (Cray MPI, Intel/OpenMPI/MPICH), NCCL/RCCL for GPU collectives, monitoring (fabric health, AIOps), and storage ecosystem integrations (Lustre/DAOS, VAST, DDN). HPE couples Slingshot with its Juniper/HPE networking stack, Apstra automation, and Marvis AIOps to accelerate design, deployment, and troubleshooting.

Business and operational implications

HPE asserts several business and operational advantages enabled by Slingshot's technical capabilities.

Convergence to ethernet

HPE is leading industry efforts to position Ethernet as the unifying fabric across front-end, back-end, DCI and scale-up use cases (slides highlight Ultra-Ethernet initiatives). For centers seeking a single fabric to support checkpointing, training, inference and storage traffic, Slingshot can provide a strong balance between ethernet economics and ecosystem and HPC leadership-class performance and features.

Cost and efficiency

HPE's internal comparisons suggest Slingshot can reduce overall network cost (NICs + switches + cabling) vs "brute force" alternatives while avoiding expensive overprovisioning. Connectionless transport and reduced NIC state can reduce NIC BOM and memory requirements that otherwise poorly scale with queue pairs at large workload process counts.

Maturity and footprint

Slingshot is deployed in 5 of the top 10 systems and 171 of all 500 systems in the November 2025 Top500 list of highest performing supercomputers installed worldwide. The breadth and depth of major global HPC centers relying on Slingshot for their scientific, engineering, and research workloads is a testament to its stability and reliability.

Operational advantages

Integrated hardware packaging (NIC mated to switch blade, liquid cooling) plus management and observability tools (AIOps, fabric monitoring) lower integration burden for dense GPU deployments and speed time to value.

ANALYST COMMENTARY

HPE Slingshot is designed to be a pragmatic, engineering driven approach to bridging the divide: between Ethernet's ecosystem and economics and the deterministic behavior and scaling characteristics historically associated with specialized interconnects. For organizations building converged HPC + large scale AI infrastructure – especially those that need predictable runtimes across 100s-1000s of GPU nodes – Slingshot aims to address real pain points (lost packet recovery, congestion, queue-pair state explosion) with integrated hardware and software features.

Despite challenges expressed by some (e.g., ecosystem interoperability, verification of end to end application performance at very large scale), large HPC and AI datacenters may perceive that

Slingshot presents a credible path to unify networking under a high performance Ethernet umbrella while protecting application efficiency and lowering total system cost.

Slingshot aligns with the industry's shift toward Ethernet-based AI and HPC networking. Its exascale proven capabilities, custom silicon, and standards alignment (e.g., UEC) position HPE to influence the next generation of open, high performance Ethernet fabrics. As AI workloads continue to dominate datacenter growth, Slingshot's congestion control and RDMA innovations could give it a compelling advantage over both commodity Ethernet and proprietary alternatives.

HPE is not alone in advancing interconnect technologies to address heterogenous HPC and AI workload demands. NVIDIA (InfiniBand), Cornelis Networks (OmniPath), and Bull (Bxi) are each taking unique approaches to overcome congestion management and tail latency challenges and other requirements from converged HPC and AI workloads. Standard ethernet through the work of industry consortiums such as Ultra Ethernet Consortium (UEC) and Ethernet Scale-Up Network (ESUN) is being funded and supported by a broad spectrum industry leaders to bolster its technical competitiveness to augment its current leadership in interoperability and multi-vendor support. Founding members for each organization include:

- UEC: AMD, Arista, Broadcom, Cisco, Bull (Eviden/Atos), HPE, Intel, Meta, Microsoft
- ESUN: AMD, Arista, ARM, Broadcom, Cisco, HPE Networking, Marvell, Meta, Microsoft, NVIDIA, OpenAI,Oracle

FUTURE OUTLOOK

Scientists, engineers, and researchers are developing applications that are placing greater demands on their HPC infrastructure, stressing all elements of the machines. Workload completion times and expediting the time it takes for them to achieve results can lead to increased profits for the commercial sector and faster scientific discovery for the academic sector.

Much attention is paid to the compute elements of HPC architecture to meet these demands. However, if those elements are not appropriately connected in a balanced and performant manner to each other and to storage, then the optimizations provided by the server technologies will go unrealized. As a result, vendors are investing heavily in system interconnects and related technologies, including:

- High-speed interconnect links
- Acceleration and offload engines
- Techniques to minimize tail latencies
- Operational and related AIOps networking software
- Integration for compatibility with middleware software

HPE has developed and positioned Slingshot to be a fully integrated, end-to-end networking solution providing host and switch silicon, switches, interface cards, operational software, and complete high performance leadership-class systems to address the complex requirements of modern HPC and AI workloads. HPE has likely made further advancements in Slingshot functionality, performance, and operational software in the months since SC25.

ISC26 in Hamburg should provide the HPC community an opportunity to see Slingshot progress made by HPE, as well as provide visibility into competitive positioning and innovation around advancements in IB, OmniPath, Bxi, and standard ethernet.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). Hyperion Research provides thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com

Copyright Notice