FOR IMMEDIATE RELEASE

Media Contacts:
media@chelsio.com
Chelsio Communications
1-408-962-3600

CHELSIO LAUNCHES T7 AI INTERCONNECT PLATFORM WITH 400GbE RDMA FOR NEXT-GENERATION DATA CENTER, STORAGE, AND AI INFRASTRUCTURE

Seventh-generation Ethernet-based AI interconnect platform accelerates AI training, inference, storage, and data center networking workloads with high efficiency and software continuity

Sunnyvale, CA – April 30, 2026 – Chelsio Communications, Inc., a leading provider of high-performance (1/10/25/40/50/100/200/400Gb) Ethernet Unified Wire Adapters and ASICs for storage networking, virtualized enterprise datacenters, cloud service installations, and embedded and cluster computing environments, today announced the launch of its seventh-generation T7 SmartNIC platform (based on Chelsio's DPU architecture), designed to support next-generation AI infrastructure, including training and inference clusters, disaggregated storage, and enterprise networking systems. The platform introduces native 400GbE with high-performance RDMA (iWARP and RoCEv2), enabling ultra-low-latency, high-throughput communication for AI training clusters and HPC environments. Building on the proven success of the widely deployed T5 and T6 families, the T7 delivers higher throughput, lower latency, and improved efficiency for storage arrays, JBOFs, and disaggregated storage architectures, as well as high-performance data center networking, while preserving software continuity.

With customers either in qualification or production, Chelsio T7 extends the company's leadership in Ethernet offload for AI, storage, and high-performance networking, enabling CPU-efficient, high-throughput connectivity across GPU clusters, storage, and network fabrics. By combining advanced congestion management, hardware offloads, and a programmable data path, T7 helps eliminate I/O bottlenecks across modern data center infrastructures. This positions Ethernet as a scalable, standards-based alternative to proprietary AI interconnects at 400Gb speeds.

"Chelsio T7 represents a milestone in Ethernet evolution, combining programmable flexibility with the proven performance our customers depend on," said Kianoosh Naghshineh, CEO of Chelsio Communications**.** "With T7, we are positioning Ethernet as a scalable, standards-based AI interconnect—delivering the performance required for large-scale AI training while preserving the flexibility and cost advantages of open networking."

T7 AI Interconnect Platform – Key Features & Highlights

- *Ethernet Performance:* Supports 1/10/25/40/50/100/200/400 GbE with sustained line-rate throughput
- *AI & HPC Optimization:* Advanced congestion avoidance, lossless Ethernet, and GPU-optimized data paths for large-scale AI training and real-time inference workloads
- *400Gb RDMA Acceleration:* Native iWARP and RoCEv2 support up to 400GbE, enabling ultra-low latency, high-bandwidth communication for AI training, HPC, and disaggregated infrastructure
- *Storage Acceleration:* Hardware offloads for NVMe/TCP, iSCSI, RDMA, and TOE, delivering peak efficiency for JBOFs, disaggregated storage, and enterprise arrays
- *Programmable Data Path*: Workload-specific acceleration, including AI data pipeline optimization and inline processing
- *Security:* On-chip crypto engines for QUIC, kTLS and IPsec/TLS, ensuring secure, high-performance connections without CPU overhead
- *Virtualization & Cloud:* SR-IOV, vSwitch, and container offloads for hyperscale cloud efficiency and multi-tenant performance isolation
- *Energy Efficiency:* Industry-leading performance-per-watt across AI, storage, and networking workloads

- *Deployment Flexibility:* Available in PCIe Gen5, OCP 3.0, and mezzanine card form factors to fit a wide range of server and appliance platforms

Platform Flexibility

The T7/S7 platform offers three ASIC variants for diverse deployment needs:

- *T7 DPU*: Full-featured silicon for AI networking, storage, and security acceleration
- *N7 Storage Controller*: Full-featured storage offload controller
- *S7 SmartNIC*: DRAM-less, cost-optimized for high-volume cloud and AI interconnect

The T7 architecture introduces unified data paths that integrate switch, FPGA, and compute functions to create purpose-built acceleration pipelines optimized for AI data movement, storage access, and distributed compute efficiency. With full software compatibility from T4 through T6, it ensures investment protection and seamless migration to next-generation Ethernet and PCIe speeds.

Expanded T7/S7 Product Family

The T7 platform introduces a broad range of T7 DPU and S7 SmartNIC adapters, each leveraging the seventh-generation offload architecture:

- *S7250 / S7450 / S7450-OCP SmartNICs* — Dual- and quad-port 1/10/25/50 GbE SmartNICs based on the S7 architecture, optimized for virtualization, embedded systems, and edge deployments.
- *S72200 / S72200-OCP SmartNICs* — Dual-port 40/50/100/200 GbE SmartNICs using the S7 silicon, ideal for high-volume cloud, HPC, and AI infrastructure.
- *S71400 SmartNIC* — Single-port 400 GbE SmartNIC supporting flexible 1x400, 4×100, and 2×200 G modes for ultra-scale fabrics.
- *T72200 / T7450 DPU Adapters* — Dual- and quad-port T7 DPU adapters featuring hardware NVMe/TCP and RDMA offloads, zero-copy, and kernel-bypass acceleration for AI and storage workloads.

- *T72200-DPU / T7450-DPU / T71200-iNIC-L DPUs* — Fully programmable T7 DPUs providing per-connection offload, direct data placement, and mixed stateful/stateless acceleration.

Configurations include L2 Ethernet NIC, storage-optimized offload adapters (NVMe/TCP, iSER/iWARP), RoCEv2 acceleration, and provisioning support such as UEFI PXE—across PCIe and OCP form factors.

Each adapter supports TCP/IP, UDP/IP, Unified RDMA (RoCEv2 & iWARP), iSCSI, NVMe-oF, NVMe/TCP, NVGRE, VXLAN, and TLS/IPsec/kTLS/QUIC/RSA offloads, providing unmatched flexibility across traditional and AI-driven data-center workloads.

AI Interconnect for Training and Inference

With support for 400Gb RDMA, the Chelsio T7 platform enables high-bandwidth, low-latency communication essential for distributed AI training and HPC workloads, reducing synchronization overhead, improving GPU utilization, and enabling more efficient scaling of large AI clusters. It enables a standards-based Ethernet AI interconnect, addressing the growing need for scalable, cost-efficient alternatives to proprietary AI fabrics. By combining lossless Ethernet, advanced congestion management, and full protocol offload, T7 delivers:

- 400Gb RDMA communication for distributed AI training and HPC scaling
- High-throughput, GPU-to-GPU and GPU-to-storage data movement for large-scale training clusters
- Ultra-low latency communication optimized for real-time inference and microservices
- Improved GPU and storage utilization through CPU bypass and full protocol offload
- Seamless scaling using standards-based Ethernet as an alternative to proprietary AI interconnects

Chelsio Communications
Accelerate

Customer Testimonials

IBM

"IBM customers increasingly want to consolidate high-performance networking and storage over Ethernet without sacrificing latency or efficiency. With Chelsio's T7, we can accelerate NVMe/TCP, iSER, NVMe/RDMA while maintaining a robust, standards-based Ethernet approach. The result is stronger performance headroom and simpler scaling for enterprise infrastructure."

— **Alistair Symon, Vice President**, IBM Storage

Seagate

"As disaggregated storage scales, NVMe/TCP efficiency becomes critical to sustain throughput and predictable performance. Chelsio N7 NVMe/TCP offload gives us the acceleration needed to reduce CPU overhead and sustain high-density storage performance under demanding data workloads. It's a strong fit for next-generation, software-defined storage architectures, while supporting all the current and legacy storage wire protocols."

— **Mohamad El-Batal, Chief Technologist**, Seagate

X-ES

"Our customers need compact, ruggedized servers that deliver deterministic performance, fast provisioning, and secure high-speed connectivity. Chelsio S7 provides the performance headroom and platform flexibility to support demanding workloads in constrained edge and embedded environments. It's a strong foundation for next-generation systems that must scale capability without adding complexity."

— **Rob Scidmore, CEO**, X-ES

Industry Recognition

"As organizations scale AI and cloud deployments, DPUs provide a clear path to efficient, standards-based acceleration," said Brandon Hoff, Executive Analyst at IDC. "Chelsio T7 DPU builds on proven architecture and open protocols to deliver performance and software continuity across generations."

"Chelsio T7 DPU is among the most advanced Ethernet-based acceleration platforms on the market," said Greg Schulz of Server StorageIO™. "By combining protocol offloads, programmable flexibility, and high-speed Ethernet in a single architecture, Chelsio is addressing critical infrastructure challenges for storage, cloud, and AI workloads. We expect strong adoption across hyperscalers, OEMs, and service providers."

Availability

The Chelsio T7 DPU and S7 SmartNIC platforms, available in a range of adapter form factors, are offered through Chelsio global OEM, ODM, and distribution partners. Shipments of T7 and S7-based adapters begin in Q2 2026 with all SKUs ramping in Q3.

Additional Resources

T7 and S7 Adapter Selector Guide

T7 DPU Product Brief

S7 SmartNIC Product Brief

AI Networking Solution: Chelsio T7 DPU and S7/T6 SmartNICs (Chelsio Solution Brief)

DPUs Enable Secure, Flexible, and Efficient Semiconductor Solutions in the Datacenter and at the Heavy Edge (IDC Analyst Brief)

About Chelsio Communications

Chelsio is a recognized leader in high-performance (1/10/25/40/50/100/200/400Gb) Ethernet adapters for networking, storage and AI infrastructure acceleration within virtualized enterprise data centers, public and private hyperscale clouds, and embedded and cluster computing environments. With a clear emphasis on performance and delivering the only robust offload solution, as opposed to simple speeds and feeds, Chelsio has set itself apart from the competition. The Chelsio Unified Wire and DPU solutions fully offload all protocol traffic, providing no-compromise performance with high packet processing capacity, sub-microsecond hardware latency, and high bandwidth. Visit the company at www.chelsio.com and follow the company on X and Facebook.

###