



FOR IMMEDIATE RELEASE

Media Contacts:

media@chelsio.com

Chelsio Communications

1-408-962-3600

CHELSIO LAUNCHES AI INTERCONNECT PLATFORM WITH 400Gb 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 – July 16, 2026 – Chelsio Communications, Inc., a leading provider of high-performance (1/10/25/40/50/100/200/400Gb) Ethernet Unified Wire Adapters and ASICs, today announced its seventh-generation AI Interconnect Platform, comprising SmartNICs, Storage Controllers, and Data Processing Units (DPUs). Built on Chelsio seventh-generation Unified Wire architecture, the platform delivers native 400Gb Ethernet, Unified RDMA (iWARP and RoCEv2), storage acceleration, and AI infrastructure optimization for cloud, enterprise, storage, and hyperscale deployments. SmartNICs and Storage Controllers are available immediately, while DPUs are available today as an evaluation platform, with production availability planned for December 2026.

Designed for AI training and inference clusters, disaggregated storage, and high-performance networking, the platform combines SmartNICs for high-volume cloud and AI deployments, Storage Controllers for storage acceleration, and DPUs for advanced networking, storage, and security acceleration. Together, they share a common hardware and software architecture that preserves software continuity while enabling scalable, standards-based AI infrastructure.

Building on the proven success of Chelsio T5 and T6 families, the seventh-generation platform extends the company's leadership in Ethernet networking, delivering higher throughput, lower latency, improved efficiency, and software continuity across SmartNIC, DPU, and storage controller deployments.



With this launch, Chelsio extends its 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 seventh-generation AI Interconnect Platform 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.”

AI Interconnect Platform – Key Features & Highlights

- *Ethernet Performance:* Supports 1/10/25/40/50/100/200/400GbE with sustained line-rate throughput
- *AI & HPC Optimization:* Lossless Ethernet, advanced congestion management, and GPU-optimized data paths for large-scale AI training and real-time inference
- *400Gb RDMA Acceleration:* Native iWARP and RoCEv2 support enables 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 and enterprise storage
- *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/kTLS.



- *Virtualization & Cloud:* SR-IOV, vSwitch, and container offloads for multi-tenant performance isolation
- *Energy Efficiency:* Industry-leading performance-per-watt
- *Deployment Flexibility:* PCIe Gen5, OCP 3.0, and mezzanine form factors

Three Product Families

- SmartNICs – S7xx adapters based on S7x silicon, a memory-free single-chip, cost-optimized SmartNIC for cloud and AI deployments
- Storage Controllers – Dedicated storage acceleration controller adapters based on N7 Storage Controller silicon
- T7 DPUs – Fully programmable DPU adapters for AI networking, storage and security acceleration based on the T7 silicon.

Built on the seventh-generation Unified Wire architecture, the AI Interconnect platform introduces unified data paths that integrate networking, storage, and compute acceleration, creating purpose-built pipelines optimized for AI data movement, storage access, and distributed computing. With full software compatibility from T4 through T6, the platform protects existing software investments while enabling seamless migration to next-generation Ethernet and PCIe speeds.

Expanded AI Interconnect Product Family

The AI Interconnect Platform includes a broad range of SmartNICs, Storage Controllers, and DPU adapters built on Chelsio seventh-generation offload architecture:

Available Now

- *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 4×100 G and 2×200 G modes for ultra-scale fabrics.
- *T72200/T7450 Storage Controllers* — Dedicated storage acceleration controllers delivering hardware offloads for NVMe/TCP, NVMe-oF, iSCSI, RDMA, and enterprise storage workloads. Available in PCIe and OCP form factors.
- *T72200 / T7450 Storage Controller 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.

Early Access

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

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 400Gb RDMA, the Chelsio T7 DPU platform enables high-bandwidth, low-latency communication for distributed AI training and HPC workloads—reducing synchronization overhead, improving GPU utilization, and enabling efficient scaling of large clusters.

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

Industry Perspectives

“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

“As organizations scale AI and cloud deployments, DPUs provide a clear path to efficient, standards-based acceleration. Chelsio T7 DPU builds on proven architecture and open protocols to deliver performance and software continuity across generations.”

— **Brandon Hoff**, Executive Analyst at IDC.

“Chelsio T7 DPU is among the most advanced Ethernet-based acceleration platforms on the market. 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.”

— **Greg Schulz**, Server StorageIO™

Availability

The SmartNIC and storage controller adapter and silicon families are available through Chelsio global OEM, ODM, and distribution partners. The T7 DPU is available as an evaluation platform



for early customer engagements, with production availability planned for December 2026.

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)

[*DPU Enable Secure, Flexible, and Efficient Semiconductor Solutions in the Datacenter and at the Heavy Edge*](#) (IDC Analyst Brief)

[*T7 L2 NIC performance*](#)

[*T7 IPSec offload performance*](#)

[*T7 kTLS offload performance*](#)

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](#).

###