



## Astera Labs Extends Leadership in Open, AI Scale-Up Networking with New 320 Lane Scorpio X-Series Smart Fabric Switch

May 5, 2026

**Now Shipping to Leading Hyperscalers, Scorpio Smart Fabric Switch Family Delivers Breakthrough Accelerator Utilization Through Memory-Semantic Based Open and Platform-Specific Protocols**

### News Highlights:

- **Largest open, memory-semantic fabric switch:** The Scorpio™ X-Series 320 Lane AI fabric switch, shipping today, supports increased scale-up cluster sizes with low latency
- **Intelligent AI fabric with In-Network Compute to maximize token economics:** Hardware-accelerated Hypercast™ and In-Network Compute engines boost collective operations by up to 2x [1], to improve time to first token and tokens-per-watt performance
- **Broadest family of PCIe 6 fabric switches:** From 32 to 320 lanes, the newly-expanded Scorpio P-Series enables diverse accelerator optionality and system topologies

SAN JOSE, Calif., May 05, 2026 (GLOBE NEWSWIRE) -- Astera Labs, Inc. (Nasdaq: ALAB), a leader in semiconductor-based connectivity solutions for rack-scale AI infrastructure, today announced the Scorpio™ X-Series 320 Lane Smart Fabric Switch, the industry's largest open, memory-semantic fabric switch engineered to improve token economics and support large scale-up clusters with minimal latency. Astera Labs also announced an expanded Scorpio P-Series PCIe fabric switch family—now spanning 32 to 320 lane configurations—designed to give data center architects at AI labs, hyperscalers, and neo-clouds the flexibility to rapidly scale compute capacity across diverse accelerators for training and serving frontier AI models.

As AI deployments move into large-scale production and achieve operational efficiency, infrastructure teams face a new set of constraints: multi-trillion-parameter models, agentic workflows, and multi-step reasoning distributed across heterogenous compute infrastructure. The industry needs connectivity solutions purpose-built for these workloads—higher radix to simplify topologies, intelligent fabric capabilities to reduce communication overhead, open and platform-specific optimizations, and datacenter-grade diagnostics to maintain uptime when a single fault can cost millions of dollars in idle compute. Scorpio Smart Fabric Switches are purpose-built to address each of these constraints and meet the demands of production AI at scale.

Scorpio's software-defined architecture is designed to integrate seamlessly with leading merchant and custom silicon, enabling AI labs and hyperscalers to integrate and deploy new accelerator platforms for both training and inference. Its memory-semantic connectivity enables accelerators to access fabric resources through native load/store operations, eliminating software overhead and improving fabric efficiency at scale. Scorpio X-Series delivers simplified high-radix scale-up topologies, cutting hops and reducing end-to-end latency across the cluster. Newly introduced Hypercast™ and In-Network Compute engines accelerate collective operations by up to 2x to maximize GPU utilization and tokens-per-watt performance. Rounding out the family, the Scorpio P-Series complements the X-Series in front-end network and AI compute system deployments, to deliver resiliency, dynamic configurability, and broad interoperability.

COSMOS™ software unifies the platform with purpose-built resiliency and serviceability, delivering non-disruptive firmware updates, OpenBMC management, and real-time telemetry. Hypercast and In-Network Compute are configured through COSMOS, extending its capabilities to now include platform performance optimizations. These capabilities are designed to maximize uptime, accelerator utilization, and ensure operational reliability for continuous production workloads while accelerating rack validation and compressing time-to-deployment. COSMOS extends across Astera Labs' complete rack-scale portfolio of fabric switches, copper connectivity, and optical solutions, enabling composable AI infrastructure from a single unified management software stack.

"The frontier models driving today's most demanding AI applications require connectivity infrastructure that keeps pace with the accelerators powering them," said Jitendra Mohan, CEO of Astera Labs. "The Scorpio X-Series 320 Lane high-radix AI fabric switch replaces multiple legacy switches to enable larger scale-up cluster sizes in a single hop and reduce overall latency. Hardware-accelerated Hypercast and In-Network Compute engines further boost collective operations by up to 2x to improve tokens-per-watt performance. Together, with Astera Labs' complete rack-scale connectivity portfolio, including the expanded 32 to 320 lane Scorpio P-Series PCIe fabric switches, give infrastructure teams a platform to build and scale AI entirely on their terms."

The Scorpio X-Series 320 Lane Smart Fabric Switch is shipping into [a merchant scale-up switch silicon market projected to reach \\$20 billion by 2030](#), with production ramp in 2H 2026.

### Patrick Moorhead, Founder, CEO and Chief Analyst, Moor Insights & Strategy, said:

"As AI moves into production, the bottleneck is shifting off the GPU and into the fabric that connects these systems. The pace of accelerator refresh is exposing how rigid many of these interconnects are, especially the closed ones, where every upgrade turns into a forklift decision. With the broadened Scorpio family, Astera Labs is pushing toward a more open, high-radix approach that

reflects how AI clusters are being built and run."

**Dylan Patel, Founder, CEO, SemiAnalysis, said:**

"At the scale AI labs are operating, interconnect is where GPU utilization goes to die. Scaling effective interconnect bandwidth is built upon two key pillars: Implementing efficient in-network compute of collective operations and reducing latency by collapsing multi-hop topologies into single-hop topologies. AI labs that control their fabric architecture, closely co-design networking solutions with compute ASICs, and thus enable step-change improvements in scale-up networking will have a structural cost and performance advantage. The Scorpio Smart Fabric Switch Family is the first open-standard platform that credibly addresses both of these pillars."

**See Scorpio at Computex 2026**

Astera Labs will showcase the Scorpio X-Series 320 Lane and its complete rack-scale connectivity portfolio at Computex 2026 (Taipei, June 2-5), including industry-first PCIe 6 scale-up optics demonstrations leveraging COSMOS for end-to-end link management.

**Additional Resources:**

- Product page: [www.asteralabs.com/scorpio](http://www.asteralabs.com/scorpio)
- Hypercast™ Blog: [www.asteralabs.com/why-your-mixture-of-experts-model-is-only-as-good-as-your-switch/](http://www.asteralabs.com/why-your-mixture-of-experts-model-is-only-as-good-as-your-switch/)

**About Astera Labs**

Astera Labs (Nasdaq: ALAB) provides rack-scale AI infrastructure through purpose-built connectivity solutions. By collaborating with hyperscalers and ecosystem partners, Astera Labs enables organizations to unlock the full potential of modern AI. Astera Labs' Intelligent Connectivity Platform integrates CXL®, Ethernet, NVLink Fusion, PCIe®, and UALink™ semiconductor-based technologies with the company's COSMOS software suite to unify diverse components into cohesive, flexible systems that deliver end-to-end scale-up and scale-out connectivity. The company's custom connectivity solutions business complements its standards-based portfolio, enabling customers to deploy tailored architectures to meet their unique infrastructure requirements. Discover more at [www.asteralabs.com](http://www.asteralabs.com).

[1] Based on Astera Labs internal analysis. "Up to 2x" reflects at least 50% reduction in AllReduce collective operation latency compared to traditional Ring AllReduce, achieved by offloading ReduceScatter and AllGather operations to Scorpio's In-Network Compute and Hypercast™ engines, reducing per-GPU transmit and receive operations from N-1 to 1 per phase. Actual results may vary based on workload, system configuration, and cluster size.

**Forward-Looking Statements**

This communication contains certain forward-looking statements regarding Astera Lab's expectations with respect to its Scorpio Smart Fabric Switch, including product features, capabilities, benefits, production ramp, effects and relevant market opportunity, and the advantages of labs controlling their fabric architecture. Such forward-looking statements are introduced using words such as "designed to," "engineered to," "expected," "projected," "targeting," "will" and variations of such words and similar expressions. Such statements involve risks and uncertainties, many of which are beyond the control of Astera Labs, that could cause actual results to differ materially from those expressed or implied in the forward-looking statements, including, among others, the risk that the expected market opportunities may not materialize; we may be unable to achieve or incorporate the expected product features; delays, disruptions, challenges or increased costs in the ability to incorporate product features or achieve the expected product roadmap within the expected timelines; the complexities and uncertainties in developing and implementing solutions based on new features and technologies; litigation or disputes related to our products; macroeconomic conditions, including general semiconductor industry economic conditions; regulatory restrictions; international conflict and other risks and uncertainties described in Astera Lab's Form 10-K, Form 10-Q and other filings with the SEC.

Forward-looking statements speak only as of the date they are made. Readers are cautioned not to put undue reliance on forward-looking statements, and no person assumes any obligation to update or revise any such forward-looking statements, whether as a result of new information, future events or otherwise, except to the extent that disclosure may be required by law.

**Media Contact:**

Peter Lo  
[Peter.lo@asteralabs.com](mailto:Peter.lo@asteralabs.com)

