Hyve Solutions and Astera Labs Collaborate to Bring AI Inference Capacity to Customers Faster

November 16, 2026

Hyve Solutions Apex™ platform with Astera Labs' Orion lineup designed to let customers deploy the rack once and choose accelerators as workloads evolve

News Highlights

- **Expanded Orion portfolio:** Hyve Solutions to integrate and deploy the Orion Compute Board (Orion-CB), Fabric Tray (Orion-FT), and Link Board (Orion-LB), validated building blocks for open, rack-scale AI inference.
- **More tokens from the same CPUs and GPUs:** Orion-CB's single-hop Scorpio fabric delivered up to 89% higher tokens per second in an 8-GPU configuration. [1]
- **Future-ready chassis:** Hyve Solutions Apex MGX 6U AI server chassis is built to support higher-power accelerators as they become available.

SAN JOSE, Calif. and FREMONT, Calif., Nov. 16, 2026 (GLOBE NEWSWIRE) -- Astera Labs, Inc. (Nasdaq: ALAB), a leader in semiconductor-based connectivity solutions for rack-scale AI infrastructure, and Hyve Solutions Corporation, a U.S.-based design and manufacturing partner for compute, storage, and networking systems in AI, cloud, and digital infrastructure, and a wholly owned subsidiary of TD SYNNEX Corporation (NYSE: SNX), today announced an expanded collaboration to help neoclouds and enterprises bring AI inference capacity online faster. Hyve will offer Astera Labs' Orion Compute Board, Fabric Tray, and Link Board as part of the Hyve Solutions Apex platform at rack scale. Because the components are pre-validated within Astera Labs' compute-agnostic reference architecture, customers can expect shorter integration and qualification cycles and the flexibility to choose the accelerators and CPUs that best fit their cost, lead time, and workload.

Hyve Solutions Apex platform offers Astera Labs' Orion line:

- **Orion Compute Board:** scales up 8 PCIe accelerators and peripherals under a single fabric switch
- **Orion Fabric Tray:** a common scale-up fabric for accelerators and memory pooling
- **Orion Link Board:** extends fabric connectivity from tray to rack scale

Hyve's Apex platform brings Orion's single-hop Scorpio fabric to customer deployments. In Astera Labs testing of an 8-GPU configuration, a single Scorpio X-Series 320-lane switch delivered up to 89% higher tokens per second than a baseline built on a conventional two-switch configuration, [1] helping operators lower their cost per token. Apex systems are air-cooled, so customers can deploy Orion in existing data centers without liquid-cooling infrastructure, and Hyve's future-ready chassis is designed to support higher-power accelerators as they arrive.

"Every GPU is waiting on something, and too often it is the fabric," said Pravin Patel, vice president and general manager of the AI Platform Solutions Group at Astera Labs. "Thanks to Astera Labs' advanced in-network compute and Hypercast™ features available in the Orion architecture, Orion is designed to let customers scale up to 160 accelerators in a pod with predictable, low-latency communication. Hyve turns that into a deployable chassis, and customers can engage with Hyve today to get better cost per token."

"The rack is now the unit of AI, and rack-level systems are what Hyve builds," said Dr. Dimitris Ziakas, chief technology officer and vice president of Architecture at Hyve Solutions. "We design, integrate, validate, and deploy at rack scale. Pairing Hyve Apex with Astera Labs' Orion gives customers scale-up AI that they can deploy once and evolve over time, swapping accelerators as their workloads change without rebuilding the system."

Hyve Apex systems featuring Orion rack-scale reference architecture are on track for availability in the first quarter of 2027. Contact Hyve Solutions at www.hyvesolutions.com or connect with Astera Labs for commercial inquiries.

Additional Resources

- **Six Five Media Interview:** 2026 OCP Global Summit interview on Orion with Sanjay Gajendra, President and COO of Astera Labs, and Dr. Dimitris Ziakas, CTO and VP, Architecture, Hyve Solutions
- **Blog:** Buy the Rack Once, Pick the Accelerator Later: How Orion and Hyve Solutions Aim to Speed Up AI Inference Capacity
- **Orion Product Page:** Orion Rack-Scale Reference Architecture

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.

About Hyve Solutions Corporation
Hyve Solutions transforms complex engineering challenges into production reality for technology innovators building AI, cloud, and connected infrastructure. As a U.S.-based design and manufacturing partner, the company rapidly delivers fast, agile execution through deep technical partnership and co-innovation, and supply chain clarity. Hyve's integrated ODM, CM, and SI capabilities eliminate vendor complexity while accelerating time-to-market with single partner accountability from design through scale. The company co-innovates with deep engineering expertise, treating customer success as its own while building tomorrow's digital infrastructure.

Hyve Solutions is a wholly owned subsidiary of TD SYNNEX Corporation (NYSE: SNX). For more information, visit hyvesolutions.com, email sales@hyvesolutions.com, or call (855) 869-6973.

About TD SYNNEX
TD SYNNEX (NYSE: SNX) is a leading global distributor, solutions aggregator, and original design and contract manufacturer that plays a central role in connecting the technology ecosystem. We support more than 150,000 customers across over 100 countries with a comprehensive edge-to-cloud portfolio spanning cybersecurity, analytics, artificial intelligence, mobility, and Everything-as-a-Service. We are a Fortune 100 company that helps partners maximize the value of technology investments and achieve measurable business outcomes through our global reach, expertise and enablement capabilities. Headquartered in Clearwater, Florida, and Fremont, California, the Company's distribution business brings together a broad portfolio of IT hardware, software and systems, providing access to products across the global IT ecosystem. The Company's Hyve Solutions business partners with technology companies to design, manufacture, and deliver traditional and accelerated compute, cloud, and connected infrastructure. For more information, visit www.TDSYNNEX.com, follow our newsroom or follow us on LinkedIn, Facebook and Instagram.

Safe Harbor Statement
This press release contains forward-looking statements within the meaning of Section 27A of the Securities Act of 1933 and Section 21E of the Securities Exchange Act of 1934. These statements include, among others, statements regarding the expected availability of Hyve Solutions Apex systems featuring Astera Labs' Orion rack-scale reference architecture; the planned integration, validation, qualification and deployment of Orion products in Hyve's Apex platform; the potential for the collaboration or products to accelerate deployment of AI inference capacity, shorten integration or qualification cycles, improve system performance or cost per token, or provide flexibility to select or change accelerators; and the scalability, interoperability, compatibility and ability of the products or platform to support future accelerators and customer workloads. Forward-looking statements may be identified by terms such as "built to," "designed to," "expect," "may," "on track," "plan," "potential," "to help," "will" and similar expressions. These statements are based on current expectations and assumptions and are not guarantees of future performance.

Actual results may differ materially from those expressed or implied by these statements as a result of risks and uncertainties, including delays or changes in product development, integration, testing, validation, qualification, manufacturing or commercial availability; the possibility that products or systems may not perform as expected or may encounter defects, limitations or other issues; differences among test configurations, customer workloads, system configurations, software, firmware, operating conditions and deployment environments; interoperability, compatibility or integration challenges involving third-party accelerators, CPUs, memory, networking, software or other components; changes in third-party product specifications, availability or release schedules; limitations in power, cooling, facility or other infrastructure that affect deployment or support for future accelerators; supply-chain constraints, component shortages or manufacturing delays; customers' or partners' decisions, requirements, adoption, purchasing or qualification timelines; and changes in industry standards, competition, economic conditions, trade restrictions, regulation or geopolitical conditions. The foregoing risks are not exhaustive. Additional information about factors that could affect Astera Labs' and TD SYNNEX Corporation's results is included in their respective filings with the Securities and Exchange Commission, including their most recent Annual Reports on Form 10-K, Quarterly Reports on Form 10-Q and subsequent filings.

Forward-looking statements speak only as of the date of this press release. Astera Labs and TD SYNNEX undertake no obligation to update or revise any forward-looking statement, whether as a result of new information, future events or otherwise, except as required by law.

[1] Testing compared two 8-GPU platforms identical except for the fabric: one Scorpio X-Series 320L switch versus a baseline of two 144-lane switches. Output tokens per second were measured across four large language models and three workload profiles at concurrency levels from 1 to 512 requests. Gains of up to 89% vary by model and workload. Results may vary depending on system configurations.

Copyright 2026 TD SYNNEX Corporation. All rights reserved. TD SYNNEX, the TD SYNNEX Logo, and all other TD SYNNEX company, product and services names and slogans are trademarks of TD SYNNEX Corporation. Other names and trademarks are the property of their respective owners.

Astera Labs Media Relations Contact

Peter Lo
Director, Communications
peter.lo@asteralabs.com

Hyve Solutions Corporation Media Relations Contact
Bill Fox
Hughes Agency
For Hyve Solutions
(864) 271-0718
billf@hughesagency.com

