NEWS RELEASE

Cisco Announces New Silicon One G300, Advanced Systems and Optics to Power and Scale AI Data Centers for the Agentic Era

2026-02-10

New G300-powered Cisco N9000 and 8000 systems, advanced optics and management upgrades deliver hyperscale-level performance, reliability and efficiency for all AI network builders.

News Summary:

- As the AI ecosystem buildout expands beyond hyperscalers, Cisco AI Networking innovations address the critical considerations of the next phase of AI buildouts - increasing energy efficiency, lowering operating costs, and simplifying operations.
- The Cisco Silicon One G300 102.4 Tbps switch silicon can power gigawatt-scale AI clusters for training, inference, and real-time agentic workloads, while maximizing GPU utilization with a 28% improvement in job completion time.
- New Cisco N9000 and Cisco 8000, powered by the G300, offer 102.4 Tbps switching speeds in systems designed for hyperscalers, neoclouds, sovereign clouds, service providers, and enterprises.
- New systems available as a 100% liquid cooled design that, along with new optics, enables a customer to improve energy efficiency by nearly 70%.
- Nexus One now delivers a unified management plane allowing customers to stand up fabrics faster, scale predictably, and operate securely and efficiently.

AMSTERDAM, Feb. 10, 2026 /PRNewswire/ -- **CISCO LIVE EMEA --** Cisco (NASDAQ: CSCO) continues to transform the network into an AI innovation platform**,** today unveiling the Silicon One G300, a 102.4 Tbps switching silicon designed for massive AI cluster buildouts. The Cisco Silicon One G300 will power new Cisco N9000 and Cisco 8000 systems that push the frontier of AI networking in the data center. The systems feature innovative liquid cooling and support high-density optics to achieve new efficiency benchmarks and ensure customers get the most out of their GPU investments. In addition, the company enhanced Nexus One to make it easier for enterprises to operate their AI networks — on-premises or in the cloud — removing the complexity that can hold organizations back from scaling AI

data centers.

"We are spearheading performance, manageability, and security in AI networking by innovating across the full stack - from silicon to systems and software," said **Jeetu Patel, President and Chief Product Officer, Cisco**. "We're building the foundation for the future of infrastructure, supporting every type of customer—from hyperscalers to enterprises—as they shift to AI-powered workloads."

"As AI training and inference continues to scale, data movement is the key to efficient AI compute; the network becomes part of the compute itself. It's not just about faster GPUs – the network must deliver scalable bandwidth and reliable, congestion-free data movement," said **Martin Lund, Executive Vice President of Cisco's Common Hardware Group**. "Cisco Silicon One G300, powering our new Cisco N9000 and Cisco 8000 systems, delivers high-performance, programmable, and deterministic networking – enabling every customer to fully utilize their compute and scale AI securely and reliably in production."

Silicon One G300: The Networking Foundation for the Agentic Era
The new Silicon One G300 is a 102.4 Tbps switching silicon that exemplifies Cisco's rapid innovation and sets a new standard for AI backend networking. It is designed to power massive, distributed AI clusters with high performance, security, and reliability.

The G300 uniquely offers **Intelligent Collective Networking**, which combines an industry-leading fully shared packet buffer, path-based load balancing, and proactive network telemetry to offer better performance and profitability for large-scale data centers. It efficiently absorbs bursty AI traffic, responds faster to link failures, and prevents packet drops that can stall jobs, ensuring reliable data delivery even over long distances. With Intelligent Collective Networking, Cisco can deliver 33% increased network utilization, and a 28% reduction in job completion time versus simulated non-optimized path selection, making AI data centers more profitable with more tokens generated per GPU-hour.

Cisco Silicon One G300 is highly programmable, enabling equipment to be upgraded for new network functionality even after it has been deployed. This enables Silicon One-based products to support emerging use cases and play multiple network roles, protecting long-term infrastructure investments. And with security fused into the hardware, customers can embrace holistic, at-speed security to keep clusters up and running.

Cisco Silicon One is the industry's most scalable and programmable unified networking architecture, offering a complete portfolio of networking devices across AI, hyperscaler, data center, enterprise, and service provider use cases. Introduced in 2019, Cisco Silicon One is playing critical roles in major networks around the world.

New Systems, Optics: High-Density, Scalable Design for Power-Efficiency and Performance
To enable AI network builders of all sizes – hyperscale to enterprise – Cisco is introducing the next generation of Cisco N9000 and Cisco 8000 fixed and modular Ethernet systems, powered by Silicon One, and designed for the extreme power and thermal demands of AI workloads. Cisco is also introducing innovative optics that unlock even higher efficiency and greater reliability. These advancements deliver a major leap forward in power-efficient, high-performance AI infrastructure.

- **Cisco N9000 and Cisco 8000 102.4T Systems:** powered by Silicon One G300, deliver a new standard in data center performance and efficiency with liquid-cooled and air-cooled designs. The 100% liquid cooled systems enable significantly higher bandwidth density and achieve a nearly 70% energy efficiency improvement, offering the same bandwidth in a single system that would previously have required 6 prior generation systems.

- **1.6T OSFP (Octal Small Form-factor Pluggable) Optics**: deliver ultra-high bandwidth connectivity targeting AI scale out solutions for 1.6T switch to NIC links and 1.6T, 800G, 400G, or 200G switch to server links, offering customers high performance and reliability.
- **800G Linear Pluggable Optics (LPO):** driving greater efficiency for AI scale out networks, LPO reduces optical module power consumption by 50% compared to retimed optical modules. With new N9000 and 8000 systems supporting LPO, customers can reduce overall switch power by 30%, leading to more reliable and sustainable operations.
- **Expanded Portfolio of Silicon One P200-powered Systems:** building on the introduction of 51.2T systems for hyperscale deployments, new P200-powered N9000 systems and expanded OS support on 8223 systems deliver scale across, data center interconnect, universal spine, and core and peer routing capabilities to neoclouds, enterprises, and service providers. Cisco is also introducing new 28.8T modular line cards. This expansion of P200-powered offerings, combined with Cisco 800G ZR/ZR+ coherent pluggable optics, enables a wide range of customers to deploy a common architecture across multiple roles in their network.

Cisco Nexus One: Intelligent AI networking to drive AI infrastructure forward
Organizations need greater flexibility in where and how they run AI workloads. To address the diverse requirements of these environments, Cisco is advancing Nexus One with a unified management plane that brings together silicon, systems, optics, software, and programmable intelligence as a single integrated solution. We're also introducing AgenticOps for data center networking through AI Canvas — making it easier to troubleshoot through guided, human-in-the-loop conversations that turn complex issues into actionable resolutions. Key capabilities include:

- **Unified Fabric:** Nexus One allows customers to deploy fast and adapt their networks as demands shift, even across multiple sites. The Cisco N9000 systems serve as the common hardware for a diverse set of fabrics, including Nexus Hyperfabric, with a unified management plane to centralize operations. And API-driven automation and customization are built-in.
- **AI Job Observability and Native Splunk Platform Integration:** Nexus One delivers job-aware, network-to-GPU visibility that correlates network telemetry with AI workload behavior. With native Splunk platform integration coming in March, customers will be able to analyze network telemetry directly where data resides—without having to move it to external platforms. This is an essential capability for sovereign cloud deployments and compliance-sensitive environments where data locality is paramount.

Cisco's flexible and integrated approach enables more choice, stronger security, and deeper observability—making upgrades and innovation easier, regardless of where customers begin their AI journey.

Availability
The Silicon One G300, G300-powered systems and optics will ship this year.

Ecosystem Support
Cisco is proud to work with its strategic technology partners, including AMD, DDN, Intel, NetApp, NVIDIA, and VAST, to combine cutting-edge networking, compute, and storage to deliver optimized infrastructure. Cisco's partnerships across the AI ecosystem give customers confidence and choice in their investments.

Supporting Quotes
"AI at scale demands open, standards-based networking that customers can deploy with confidence across diverse environments. Our longstanding collaboration with Cisco helps advance high-performance, standards-based Ethernet fabrics while reinforcing end-to-end interoperability, from GPU and CPU platforms to AI NICs, DPUs, and the software stack. We're excited to continue to partner with Cisco across our enterprise and AI product stack, as we focus on giving customers the flexibility

and choice to build resilient, scalable AI infrastructure." ***– Yousuf Khan, corporate vice president, Networking Technology and Solutions Group, AMD***

"AI at scale demands both compute efficiency and high-performance AI networking fabric. Intel® Gaudi® 3 AI accelerator combined with Cisco Nexus 9000 switching delivers an optimized, open solution that lets customers build at scale LLM inference clusters with uncompromising cost-efficient performance." – ***Anil Nanduri, VP, AI Get-to-Market & Product Management, Intel***

"Our partnership with Cisco is built on a long history of innovation, starting with the world's first ACI deployment. As du Tech continues to expand, we rely on Cisco's partnership to meet our infrastructure needs across multiple generations. The N9000 series, with its Silicon One G300 NPU, gives us the 100Tbps capacity and efficiency we need as we move into the 1.6T Ethernet era to support our AI and cloud growth. By using Nexus One and AgenticOps, we can maintain high-performance, energy-efficient data centers. This collaboration ensures our infrastructure is ready to handle the UAE's future digital demands." ***– Jasim Al Awadi, Chief ICT Officer, du***

"Sharon AI has selected Cisco Nexus Hyperfabric to power our AI solutions because it delivers a comprehensive, end-to-end infrastructure that seamlessly integrates networking, GPU, and storage. This robust platform enables us to accelerate enterprise deployment with a fully Cisco-validated solution that aligns with NVIDIA NCP design and performance standards. By leveraging a unified, cloud-managed control point, we simplify operations, reduce management overhead, and ensure consistent, high-performance experiences—ultimately allowing us to innovate faster and deliver greater value to our customers." ***– James Manning, CEO, Sharon AI***

"Cisco's bold innovation is fueling a new era of intelligent, sustainable operations in our data centers. The Cisco N9000 series, now powered by G300 Silicon One for breakthrough 1.6T scale-out performance, along with the deep-buffer 800G P200-based switches for seamless scale across our network, dramatically expand what's possible for AI infrastructure. Combined with the intelligence of Nexus One with cutting-edge AgenticOps and 102.4T liquid cooling, Cisco is positioned to deliver high-performing, energy-efficient environments that keep us ahead of tomorrow's demands." ***– David Driggers, CEO and CTO, Cirrascale Cloud Services***

"As AI adoption moves beyond hyperscalers and scales across enterprises, neoclouds, and sovereign environments, network architecture is becoming a defining constraint on performance, cost, and sustainability. Cisco's approach—combining high-performance silicon, liquid-cooled systems, advanced optics, and integrated operations—speaks directly to the next phase of AI infrastructure, where maximizing GPU utilization, improving energy efficiency, and simplifying operations are critical to realizing real economic value from AI at scale." ***– Matt Eastwood, SVP, Enterprise Infrastructure and Datacenter Group, IDC***

"Silicon One G300 represents a transformational leap in networking silicon, finely tuned for the demands of large-scale AI clusters. By significantly boosting network utilization, more GPU tokens are generated per hour. AI networking is critical in reshaping the economics and performance expectations for AI data centers. Back-end scale-out networks will rapidly move to 1.6T and be a key driver to push the Ethernet data center switch market above $100B a year." ***– Alan Weckel, Founder and Technology Analyst, 650 Group***

"Networking has been the fundamental constraint to scaling AI. Cisco has emerged as an important player in advancing high-speed, efficient Ethernet solutions, and the 102.4T Silicon One G300 is a clear proof point of that progress. At this scale, networking directly determines how much AI compute can actually be utilized." ***– Dylan Patel, Founder, CEO, and Chief Analyst, SemiAnalysis***

"Cisco's Silicon One G300 sets a new bar for AI networking, and VAST ensures that bandwidth translates into real application throughput. The VAST AI Operating System unifies ingest, retrieval, vector search, and real-time analytics, removing the data-path bottlenecks that starve GPUs and stall jobs. Together with Cisco's new systems and operations model, we're delivering a validated, turnkey foundation that helps enterprises scale agentic AI from pilot to production." ***– John Mao, VP, Global Technology Alliances, VAST***

"At AI-factory scale, performance is no longer determined by the network or the data layer alone—it's defined by how tightly they work together. Cisco's Silicon One G300–powered Nexus platforms provide the deterministic, high-bandwidth fabric required for agentic and GPU-dense environments, while DDN's AI-native data intelligence ensures data is always in the right place, at the right time, at full speed. Together, we remove the hidden bottlenecks that starve GPUs, extend job completion times, and stall production AI. This validation underscores a shared commitment to delivering AI infrastructure that is not just powerful on paper, but proven in real-world, large-scale deployments." **– *Sven Oehme, CTO, DDN***

"NetApp and Cisco have collaborated closely to help enterprises store, access, and optimize their data to accelerate AI innovation. With NetApp AFX disaggregated storage and Cisco's G300 based N9000 systems with 102.4T switching, data can reach GPUs at the speed organizations need to power innovation. Combining world class networking from Cisco and NetApp's enterprise-grade data platform enables customers to scale AI without compromise." ***– Syam Nair, CPO, NetApp***

"WWT clients know and trust Cisco networking in the AI data center. With the G300-powered N9000 and Nexus One, we're extending that trust to AI workloads—102.4 terabits of capacity with the industry's largest on-chip buffer, managed through Nexus Dashboard. This is the fastest we've seen Cisco move, and it's exactly what our clients need to accelerate their AI journeys." ***– Neil Anderson, VP and CTO, Cloud, Infrastructure, and AI Solutions, WWT***

"Cisco's Silicon One–powered Nexus switches N9364E-SP2R-X (P200) and N9364F-SG3 (G300) switches redefine what's possible in the modern AI data center, delivering an unprecedented scale with 64 ports of 800Gb/1.6Tbs and 51.2T/102.4Tb total throughput. With operational flexibility across NX-OS and Cisco ACI with next-generation silicon, these switches position enterprises and neoclouds to embrace high performance platform-driven innovation, scale and digital transformation with confidence." **– *Brian Campbell, VP, Hybrid Infrastructure & Digital Experience, CDW***

"Our hyperscale and neocloud customers need networking that matches GPU density. Cisco's N9000 with NX-OS delivers programmability and telemetry to optimize every flow. The G300 silicon enhances this with industry-leading buffers, power efficiency, and 1.6T port density. Through our strategic partnership with Cisco, we deliver lossless, high-performance networking for AI training and inference. The Nexus One Platform ensures predictable performance—deep buffers manage bursty traffic, and Intelligent Packet Flow maximizes GPU utilization." ***– Thomas Berger, Director, Data Center Networking, Computacenter***

Additional Resources:

- Executive Blog Post: One platform for the Agentic AI era by Jeetu Patel, President and Chief Product Officer, Cisco
- Blog Post:Networking for the Agentic Era: Cisco Unveils New Innovations in Scale and Simplicity
- Blog Post:Cisco Silicon One G300: The Next Wave of AI Innovation
- Blog Post: More Scale, More Intelligence, and More Control: New Cisco Solutions for Accelerating AI Networking

- For more information about announcements from Cisco Live Amsterdam, please visit the Cisco Newsroom.

About Cisco
Cisco (NASDAQ: CSCO) is the worldwide technology leader that is revolutionizing the way organizations connect and protect in the AI era. For more than 40 years, Cisco has securely connected the world. With its industry leading AI-powered solutions and services, Cisco enables its customers, partners and communities to unlock innovation, enhance productivity and strengthen digital resilience. With purpose at its core, Cisco remains committed to creating a more connected and inclusive future for all. Discover more on The Newsroom and follow us on X at @Cisco.

Disclaimer: Many of the products and features mentioned are still in development and will be made available as they are finalized, subject to ongoing evolution in development and innovation. The timeline for their release is subject to change.

View original content to download multimedia:https://www.prnewswire.com/news-releases/cisco-announces-new-silicon-one-g300-advanced-systems-and-optics-to-power-and-scale-ai-data-centers-for-the-agentic-era-302683306.html

SOURCE Cisco Systems, Inc.