



NEWS RELEASE

Cisco Advances Scalability and Openness at OCP 2026

2026-10-12

Delivers the resilient, open Ethernet foundations required to connect and scale massive, distributed AI factories

News Summary

- Cisco and NVIDIA plan to develop an open industry specification for a next-generation 400G-per-lane pluggable optics design, increasing bandwidth density while preserving serviceability and customer choice.
- Cisco advances open infrastructure across silicon, systems and optics to meet the demands of AI workloads, spanning high-density scale-out and distributed scale-across networks.
- The Cisco Silicon One portfolio continues to expand to serve all AI builders, with the first G300-powered systems shipping to customers, new 1.6T G302 silicon, and P200-powered systems now deployed in production hyperscale environments.

SAN JOSE, Calif., Oct. 12, 2026 /PRNewswire/ -- Today at the Open Compute Project (OCP) Global Summit 2026, Cisco (Nasdaq: CSCO) introduced new silicon, systems and optical innovations that advance open AI infrastructure from rack-scale compute to distributed AI capacity. Together, the announcements demonstrate how Cisco is engineering the network as a coherent, programmable architecture while preserving the openness and choice customers expect in the super intelligence era. The new capabilities deepen Cisco's continued commitment to deliver an open stack from silicon to Kubernetes, required to connect the AI foundations for hyperscalers, neoclouds and enterprises. Strong market adoption validates Cisco's leadership and underscores deep customer trust as demand for AI compute capacity rapidly accelerates.

"AI is redefining what it means to scale infrastructure. It's no longer just about individual data centers or clusters, but how we connect compute at massive scale without compromising performance or resilience. Cisco's strength across silicon, networking, optics, systems and security puts us in a unique position to deliver the open infrastructure customers need as AI continues to evolve," said Kevin Wollenweber, SVP and General Manager of Data Center and Internet Infrastructure, Cisco.

"Scale-Across" Connects the World's AI Capacity

Scale-across extends coordinated AI workloads across facilities, building on traditional data center interconnect (DCI) technologies while introducing significantly different workload requirements. When accelerators across sites participate in the same coordinated job, congestion, latency and recovery can directly affect compute utilization and job progress. Cisco draws on its silicon, routing and optical networking leadership to address those requirements as one packet-to-optical architecture.

Cisco's scale-across architecture brings together [Silicon One P200](#), deep-buffered systems and coherent optics to connect AI fabrics across locations to act as a unified cluster. It provides customers a single, unified compute fabric enabled by the performance, deep programmability and embedded security of the Silicon One architecture. Cisco is the first in the industry to announce its 51.2T systems, powered by the Silicon One P200, are deployed in a production environment by leading hyperscalers.

"Microsoft and Cisco share a longstanding partnership rooted in engineering collaboration and innovation. We are pleased to announce that today, Microsoft has moved into production on systems powered by the Cisco Silicon One P200. As our infrastructure demands grow, the P200's robust architecture provides the speed and enhanced buffering capabilities necessary to address the increasing scale of our cloud and high-performance services. This milestone reflects the next stage in our long-term collaboration, ensuring we continue to deliver the performance and reliability required for the modern digital era," said Anand Ramakrishna, Corporate Vice President, Engineering, Microsoft.

"Alibaba is proud to announce a major milestone for our next-generation backbone network, eCORE. We have deployed our in-house Quantum Cat router at massive scale, powered by the Cisco Silicon One P200. These routers bring 51.2 Tbps of routing capacity to our infrastructure. They feature a fully shared deep-buffer architecture, high programmability, and excellent power efficiency. This setup provides the exact performance and scale needed for today's heavy cloud and AI workloads. We paired Cisco's silicon with Alibaba's custom systems and the SONiC operating system. This combination proves the true value of open networking. It delivers the flexibility and operational power required for some of the world's most demanding networks. We look forward to deepening our work with Cisco, the Open Compute Project (OCP), and the SONiC community. Together, we will continue to advance open, high-performance networking for the AI era," said Dennis Cai, Vice President, Head of Network Infrastructure, Alibaba Cloud Intelligence Group

Scaling AI from Rack to Fabric for the 1.6T Era

Inside the data center, AI scale is driving the same need for greater bandwidth, density and power efficiency across silicon, systems and optics. To meet the radical bandwidth requirements of the next generation of AI, Cisco's Silicon One portfolio is delivering the next level of network performance designed for massive AI scale. As a member of OCP's ESUN, Cisco is committed to advancing open, interoperable, high-performance AI networking at scale and collaborating actively with ESUN on the 1.0 and future specifications for scale-up workloads.

[Announced earlier this year](#), the 102.4T Silicon One G300 is now available for customer sampling, and the first G300-powered systems are currently shipping to customers. Cisco is also expanding the Silicon One portfolio with the first 51.2T switching silicon with 200G SerDes designed for top-of-rack (ToR) deployments. The new Silicon One G302 delivers 51.2T capacity and features Cisco 1.6T optics and direct-to-chip liquid cooling integration to drive efficiency gains. The G302 is built to power ToR switch systems, which sit at the rack level and gather traffic from servers and connect it efficiently to the larger data center network. This optimizes how servers are linked and integrated, delivering unprecedented front-end scale, density, and throughput for hyperscalers. The G302 is planned to begin shipping to customers before the end of 2026.

Cisco and NVIDIA Introduce New Pluggable Optics

Together with NVIDIA, Cisco also [unveiled plans to develop](#) the Ultra-Dense Small Form Factor Pluggable (USFP), a new high density 400Gb/s-per-lane capable optical module. Compared to other approaches to increasing the bandwidth of optical modules, USFP relies on a proven operating model to scale to 400Gb/s-per-lane. USFP is expected to deliver 4x the bandwidth density and 1.3x the power efficiency of 200Gb/s-per-lane optics, improving the efficiency and scalability of AI factories without sacrificing serviceability, deployment flexibility or the operational simplicity of today's pluggable optics. To drive broad industry engagement and adoption, Cisco and NVIDIA plan to create an open industry specification for USFP through an upcoming community-led multi-source project.

Expanding Secure AI Networking

SONiC (Software for Open Networking in the Cloud) is the open-source operating system that forms the foundation of many hyperscaler and now neoclouds, service providers and enterprise deployments, across the AI builder ecosystem. Cisco customers can adopt Cisco-validated SONiC with [Nexus One](#), a next generation data center architecture, embracing an open approach for multi-site to rack-scale environments. This support will be extended with [Live Protect](#) for customers starting mid-year 2027. Utilizing Cisco runtime technology, Live Protect reduces the exposure customers would typically face between a vulnerability disclosure and a patch on supported Cisco products and releases. Live Protect will give Cisco SONiC customers time to test and deploy the permanent software fix without disrupting operations.

Control and Recovery at AI Scale

Out-of-band management systems form a network operators rely on to provide remote access even when the production network is down, allowing for fixes during a blackout or cyber incident. But at AI scale, out-of-band management systems can become a patchwork of legacy systems, tools and proprietary interfaces, making thousands of GPU nodes harder to manage and fix. The Cisco 8001, a new Silicon One A100-powered management system available as 10G or 1G, solves this challenge by upgrading performance and bringing SONiC to this critical network. Operators can use familiar APIs, automation and operating practices from the fabric to the rack, without building and maintaining a separate management environment. Customers can now deploy a simpler, programmable management network that scales with AI infrastructure, and remains available when it is needed most.

Cisco Silicon One: The Industry's Most Scalable and Programmable Unified Networking Architecture

[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. With today's innovations, Silicon One proves its ability to scale from 100G to 100T with a single, unified silicon architecture.

Supporting Resources:

- Executive Blog Post: [Cisco and NVIDIA Announce USFP: Expanding the Path to AI Connectivity](#)
- Executive Blog Post: [SONiC's Next Chapter with Cisco: Open Choice, Production Confidence](#)

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: Some products or features described may be in various stages of development and offered on a when-and-if available basis. Cisco reserves the right to change delivery timelines and will have no liability for any delays or failures to deliver.

View original content to download multimedia:<https://www.prnewswire.com/news-releases/cisco-advances-scalability-and-openness-at-ocp-2026-302905690.html>

SOURCE Cisco Systems, Inc.