Cisco N9000 Series Switches with Nexus One for AI Networking

© 2026 Cisco and/or its affiliates. All rights reserved.

Benefits

- **Unified end-to-end AI networking:** Cisco Nexus® One, an open, unified networking platform that brings together silicon, systems, optics, softw are, and a single operating model into one integrated solution and industry-first AgenticOps streamlining operations for AI Infrastructure
- **High-speed and low latency switches:** Cisco N9000 Series Switches powered by Cisco® Silicon One®, cloud-scale ASICs, and NVIDIA Spectrum-X Ethernet switch silicon
- **Massive scale meets intelligence:** 102.4T Cisco Silicon One G300 based Cisco N9300 Series Switches powering 1M+ GPU clusters with intelligent load balancing, microsecond telemetry, and the industry's largest packet buffer for predictable AI performance
- **Power-efficient innovation:** introducing direct-to-chip liquid cooling and linear-drive optics
- **NCP RA compliance:** for customers who require NVIDIA Cloud Partner (NCP) reference architecture—compliant designs, Cisco N9100 Series Switches are powered by NVIDIA Spectrum-Ethernet switch silicon compliant designs, Cisco N9100 series switches are powered by NVIDIA Spectrum-X Ethernet switch silicon.
- **Choice of Operating Systems (OS):** Cisco NX-OS and SONiC operating systems are offered based on customers' requirements.
- **Flexible connectivity options:** With support for 400G, 800G to 1.6T in both QSFP-DD and OSFP form factors, Cisco Optics ensure broad compatibility.
- **Maximized performance for AI networks:** Cisco Intelligent Packet Flow incorporates live telemetry, congestion-aware load balancing, and fault detection to ensure efficient and reliable traffic steering across the fabric, thereby helping customers unlock consistent performance and reduce Job Completion Times (JCTs).
- **Unified management for AI fabrics:**
  - **Cisco Nexus Dashboard**, an on-premises solution, streamlines AI fabric provisioning with built-in templates delivering job-level insights, topology-aware visualization, and GPI/NIC observability. It ensures end-to-end AI fabric visibility and congestion analytics, driving operational intelligence and reliability.
- **Cisco Nexus Hyperfabric™**, a cloud-managed solution, offers streamlined AI fabric management and has a purpose-built, full-stack AI infrastructure available, enabling faster deployment, improved scalability, and accelerated AI outcomes.
- **Ultra Ethernet ready:** Cisco is one of the key steering members for Ultra Ethernet Consortium (UEC), and Cisco N9000 Series Switches are Ultra Ethernet ready to support mandatory network-side requirements.
- **Vendor-agnostic:** Cisco offers vendor-agnostic support with validated designs across leading AI ecosystem partners, including NVIDIA, AMD, Intel®, VAST, WEKA, and others.
- **Sustainable scaling:** high radix and smaller footprint fabrics enabled by Cisco AI networking for data centers support massive scale, minimizing energy consumption and optimizing resource utilization.

© 2026 Cisco and/or its affiliates. All rights reserved.

“In an era where AI workloads and real-time data processing define the competitive edge, high-performance data center switching is no longer a luxury—it’s a necessity.

Powered by Cisco N9000 Series switching, we are able to move massive volumes of data with unparalleled speed and low latency that is foundational to unlocking the full potential of Groq’s innovative fast inferencing solutions, ensuring organizations stay ahead in a data-driven world.”

Cameron Ferdinands
Head of Network Operations, Groq

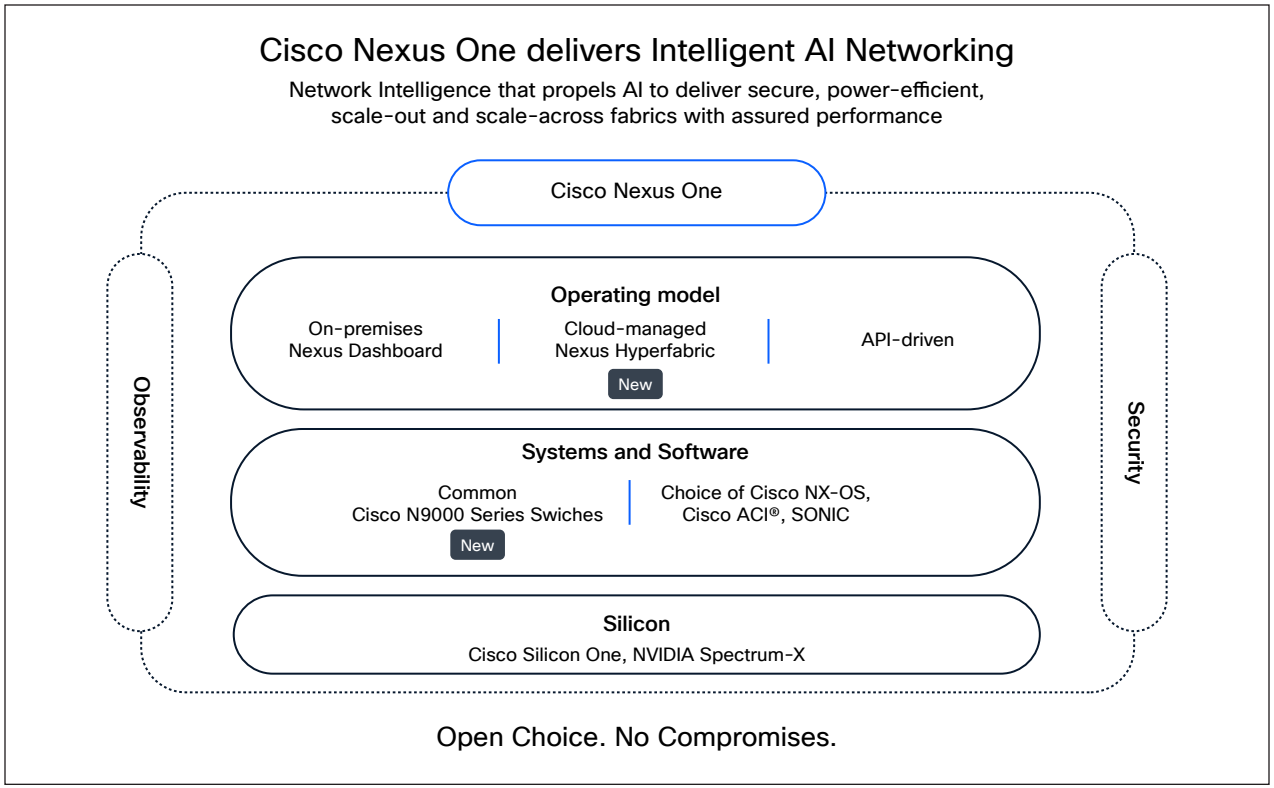

Figure 1. Cisco Nexus One: Intelligent AI Networking

Introducing Cisco Nexus One, network Intelligence that propels AI to deliver secure, power-efficient, scale-out and scale-across fabrics with assured performance.

© 2026 Cisco and/or its affiliates. All rights reserved.

“Cisco’s N9100 series powered by NVIDIA Spectrum-X Ethernet switch silicon, provides a solution for high-performance, open infrastructure to meet our AI cloud demands. The capability to run NX-OS or SONiC under a unified operating model on Nexus Dashboard delivers more flexibility to our customers with operational simplicity. Its enterprise-grade networking with the scale and agility of the cloud — exactly what the next generation of AI workloads requires.”

Leong Sun
Head of Infra, GMI Cloud

Product overview

AI has evolved from natural language-focused Large Language Models (LLMs) to agentic AI, and now to physical AI that operates in the real world. As AI becomes essential to enterprises, neoclouds, and hyperscalers, organizations need high-performance networking to support its explosive traffic demands.

Cisco is fueling the AI revolution with a unified data center network powered by Cisco N9000 Series switches and Nexus Dashboard.

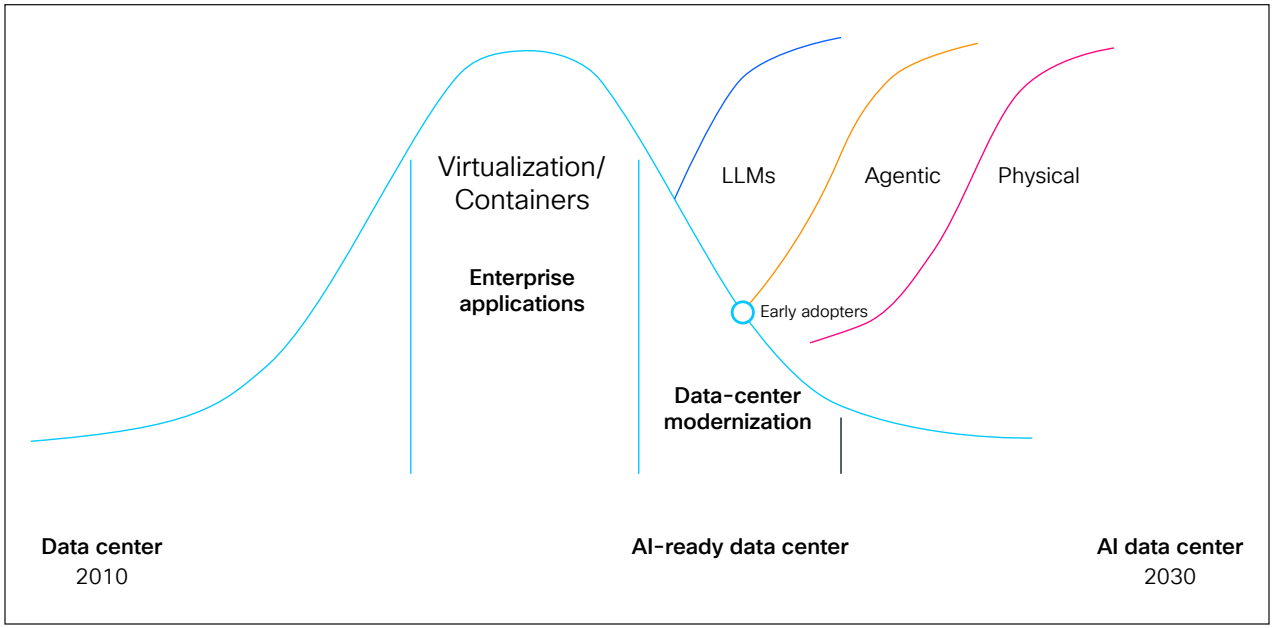

Figure 2. Big shifts redefining the data center

© 2026 Cisco and/or its affiliates. All rights reserved.

“The real challenge in AI infrastructure isn’t just performance—it’s maintaining operational sanity as you scale from dozens to thousands of GPUs. Cisco’s approach with NX-OS and Nexus Dashboard creates a single pane of glass across our entire AI fabric, whether we’re optimizing inference latency in the front-end or maximizing training throughput in the back-end. That operational simplicity translates directly to faster deployments and lower TCO.”

Xiaohe Hu
CEO, Infrawaves

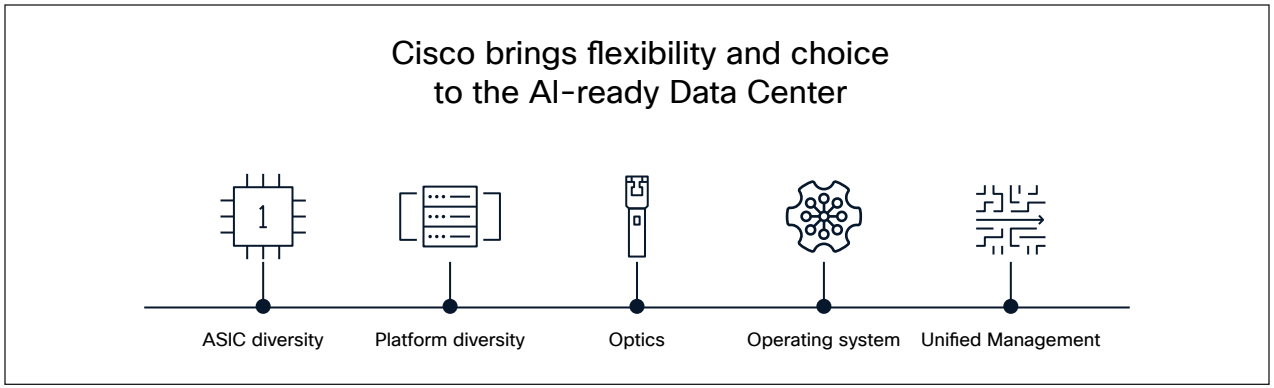

Figure 3. Foundations of Cisco AI Networking with N9000 and Nexus One

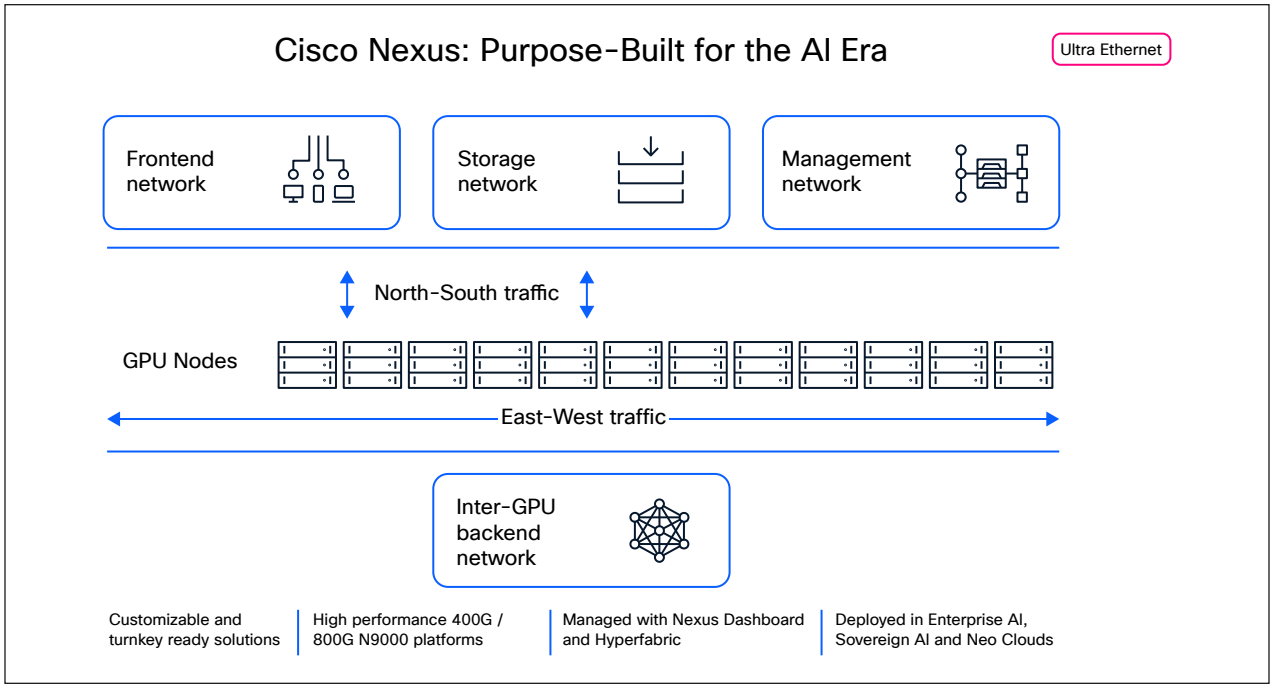

Figure 4. N9000 Switches with Nexus One for AI Networking

© 2026 Cisco and/or its affiliates. All rights reserved.

Cisco delivers a complete unified data center platform built to span the entire AI Networking infrastructure continuum from silicon to software. Our platform comprising Cisco Silicon One and NVIDIA Spectrum-X Ethernet switch ASICs,, OS flexibility, advanced optics, strategically fused security, and unified management are architected specifically for both ever evolving workloads and meeting AI's demanding requirements. Our unified platform approach delivers breakthrough performance that transforms infrastructure complexity into competitive advantage.

Learn more

See how Cisco IT is transforming its AI networks with Cisco N9000 Series Switches and Nexus Dashboard.

Get your AI Networks ready with the Cisco Nexus portfolio. For detailed information, visit https://www.cisco.com/c/en/us/products/collateral/networking/cloud-networking-switches/nexus-9000-switches/nexus-9000-ai-networking-wp.html.

Cisco AI Networking webpage.

© 2026 Cisco and/or its affiliates. All rights reserved. Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

C45-5186151-02 02/26