



Cisco Nexus 9000 Series Switches



The data center networking industry is going through a paradigm shift with the rapid evolution of AI/ML technology. The rise of massive AI/ML workloads together with enterprise workloads is creating scale out requirements that are leading to a tectonic shift in modernizing data centers. AI/ML workloads drive the need for dedicated AI datacenters with high-bandwidth, scalable, low latency and lossless networks, whereas enterprise workloads continue to drive the need for automated, secure, scalable, and efficient data-center fabrics.

Ethernet switches deliver secure high-speed connectivity within the data center for servers, storage, or GPU clusters that host critical applications and workloads that drive the global digital economy. The varying range of performance needs of these applications and workloads require Ethernet switches that offer a choice of speeds, capacity, and scale. Cisco Nexus 9000 Series Switches are designed to cater to all these requirements and offer a complete portfolio of distributed modular, centralized modular, and fixed switches. With the silicon, hardware, software, management tool, and optics, the Cisco Nexus 9000 Series Switches, provide the right bandwidth, latency, congestion management mechanisms, telemetry, visibility, and automation to build desired fabric.

Trends and drivers in modern data center networking

Large enterprise and AI/ML data centers and data halls with storage, servers and networking are undergoing speed transitions in the fabric and access links. Data-center fabrics are being upgraded to 400G and 800G now to cater to the higher speed access connections driven by the adoption of GPUs with ports speeds of 200G or 400G.

High-performance enterprise workloads are accelerating the transition of data center server access connections from 10/25G to 50/100G. This in turn is accelerating the migration of data center fabrics from 100G to 400G/800G speeds.

The increasing adoption of AI/ML is further accelerating the transition to 400G and 800G fabric and access connections, especially as GPUs with 200G or 400G interfaces are being rapidly deployed to run AI/ML workloads.

Besides the need for modernization, factors that drive transition are scale requirements, easy migration, cost/performance balance, and real-time visibility of the network.

Evolution in data-center architectures

To meet customers' needs for data-center network modernization, Cisco Nexus 9000 Series Switches support Cisco® Application Centric Infrastructure (Cisco ACI®), Cisco NX-OS VXLAN EVPN, Cisco IP Fabric for Media, Cisco Nexus Data Broker, and IP-routed or Ethernet-switched Layer-2 fabrics using a comprehensive set of unicast and multicast IPv6/IPv4 and Ethernet protocols.

Software support for in-service software upgrades, graceful insertion and removal, and vPC with fabric peering provide fabrics with enhanced manageability and reliability. A vast ecosystem of technology partners and API integrations are supported.

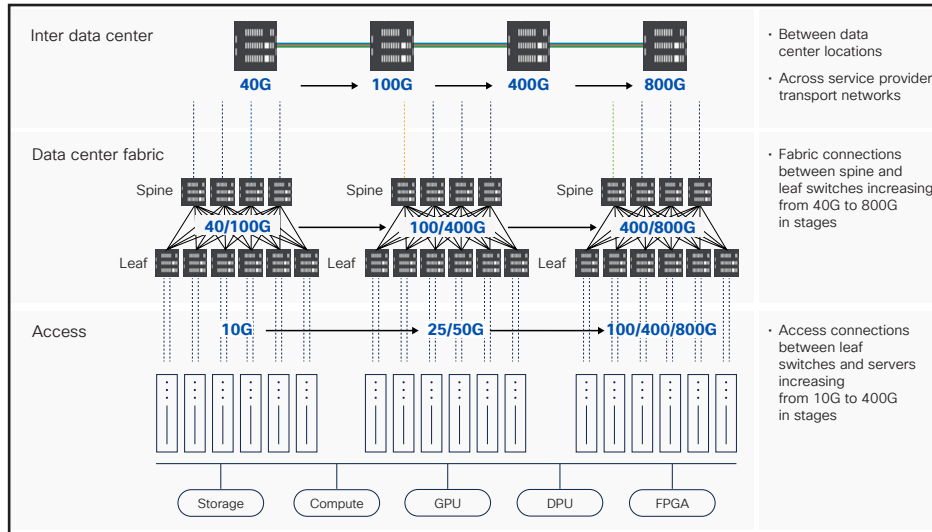


Figure 1. Data Center network connectivity

High-performance enterprise workloads support data-center fabrics transitioning from 40G/100G to 400G/800G; data center server access connections transitioning from 10G/25G to 50G/100/200G.

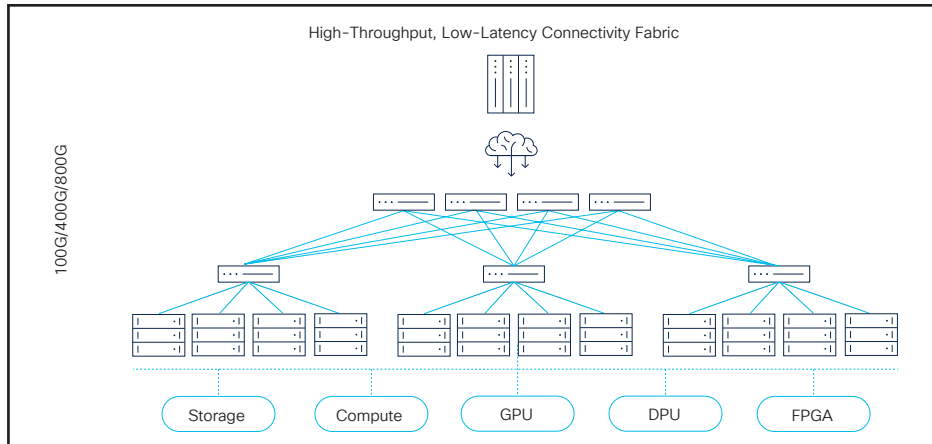


Figure 2. Fabric for High Performance Compute (HPC)

Lossless Ethernet for AI/ML networking

Cisco Nexus 9000 Series Switches support high bandwidth, congestion management and flow-control algorithms along with the right transport protocols, low latency, and telemetry capabilities to meet the design requirements of AI/ML fabrics.

Early Congestion Notification (ECN), Priority Flow Control (PFC), Approximate Fair Drop (AFD), and Weighted Random Early Detection (WRED) capabilities in Cisco ASICs deliver the high performance needed to build lossless Ethernet transport networks. Through RDMA over converged ethernet version 2 (RoCEv2) transport protocol, efficient data movement at a high throughput, low latency can be achieved.

Modernize your data-center switching

Cisco ASICs power Cisco Nexus 9000 Series Switches. It combines innovative congestion management techniques, flow-control algorithms, line-rate telemetry, and visibility along with a range of bandwidth and port-speed options from 1G to 800G.

The Cisco Nexus 9000 Series Switch portfolio is a versatile and streamlined portfolio of distributed modular, centralized modular, and fixed switches that cover all data center requirements – super spine, spine, leaf, and border leaf.

Modern data center fabric spine or super-spine switches require high performance and scale provided by the Cisco Nexus 9000 distributed modular switches and line cards. The centralized modular and fixed form factor switches provide the high capacity and high port density needed for modern data center fabric spines, leafs, and border leafs. In addition, several Cisco Nexus 9300 fixed form factor switches provide high-speed 100G, 400G, or 800G uplinks to connect legacy data-center servers to modern data center spine switches.

Some enterprises want a single network that can connect CPUs, GPUs, and Data-Processing Units (DPUs) which is simple to operate, maintain, and scale. To implement such shared Ethernet networks, the Cisco Nexus 9000 Series switches offer a range of SERDES speeds. Choice of 28G SERDES, 56G SERDES, and 112G SERDES that enable the access connectivity to 10G, 25G, 50G, or 100G servers for enterprise data-center workloads as well as the latest 200G- or 400G-capable servers for AI/ML workloads.



Figure 3. Cisco Nexus 9000 portfolio

High-speed connectivity

Cisco Nexus 9000 Series Switches support various form factor modules for 40G (QSFP+), 100G (QSFP28), 200G (QSFP56), 400G (QSFP-DD), 800G (QSFP-DD800 and OSFP), and future higher speed ports. The choice of using the QSFP form factor for these port speeds provides backwards compatibility and seamless connectivity to legacy lower-speed QSFP ports, enabling a smooth transition to higher speed fabrics.

In addition to the QSFP port form factor, Cisco Nexus 9000 switches and ASICs support 56G and 112G SERDES for high-speed connections that improve performance and efficiency at lower cost and power.




Figure 4. Cisco QSFP and QSFP-DD high-speed optics

Visibility with rich, real-time telemetry

Cisco Nexus Dashboard provides the simplest way to manage and operate global data center and cloud fabrics with visibility and sustainability, built on Cisco Nexus 9000 Series switches. Cisco Nexus Fabric Controller, Insights and Orchestrator are services available on Cisco Nexus Dashboard. By offering these services from a single platform, along with third-party ecosystem partners, Cisco Nexus Dashboard can configure, operate and analyze modern enterprise and AI data-centers and cloud networks through automation and analytics.

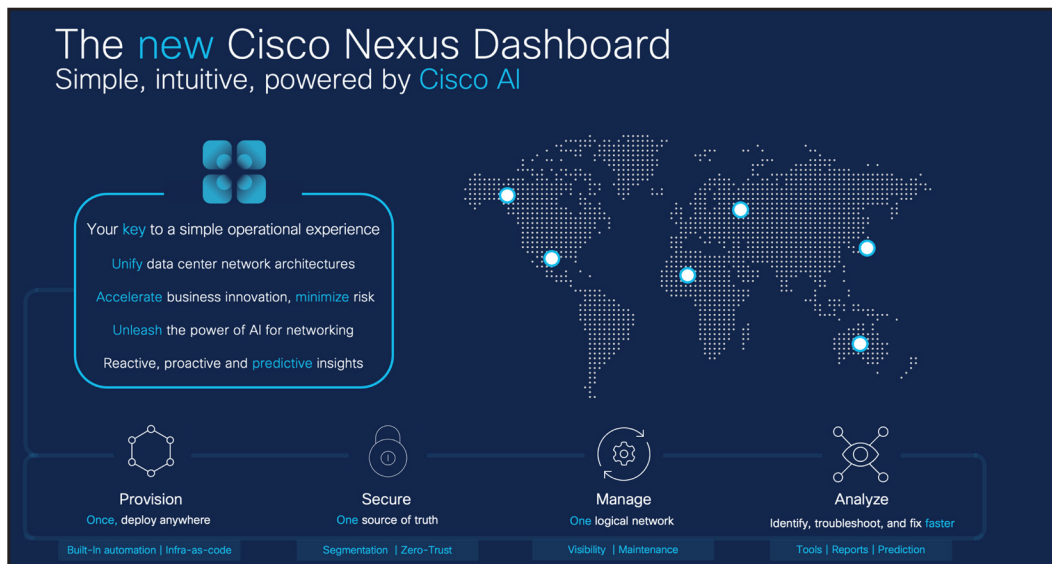


Figure 5. Cisco Nexus Dashboard

Resources

- [Cisco Data Center Networking](#)
- [Cisco Nexus 9000 Switches](#)
- [Cisco 400G and 800G data center switching](#)
- [Cisco Nexus Dashboard](#)

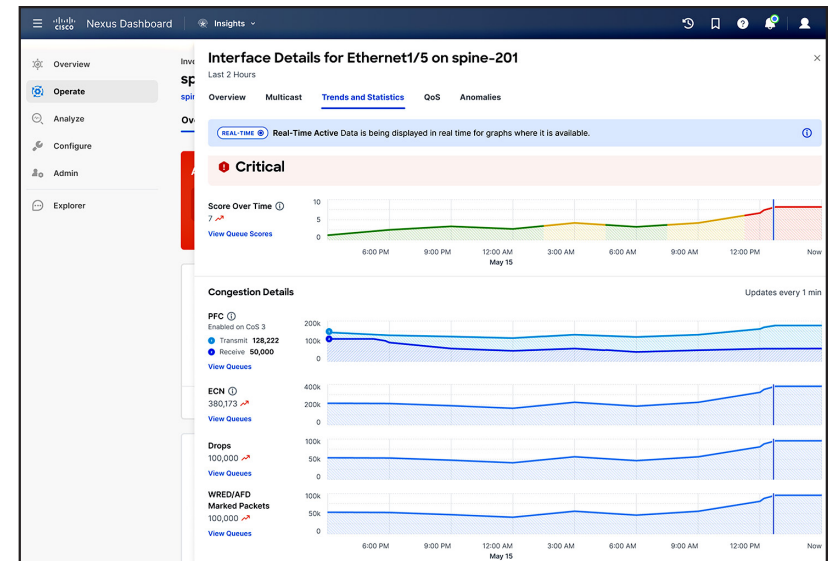


Figure 6. Cisco Nexus Dashboard Insights: Optimizing AI networking with congestion details for lossless Ethernet fabrics