

Cisco Data Center Networking AI/ML

AI/ML Networking with Cisco Nexus 9000

Cisco® data center networking with Cisco Nexus® 9000 Series Switches and the Nexus Dashboard operations and automation platform provides the interplay of cutting-edge silicon, hardware, software, optics, and management tools to deliver high bandwidth, lossless, and low latency AI/ML networks.



Overview

The applications of AI in areas of finance, autonomous systems, healthcare, and research have taken center stage in modern data centers. These require AI-capable systems that are capable of analyzing vast amounts of data, training models, and making decisions. Designing an AI-optimized data center requires seamless collaboration of silicon, systems, software, networking, and operational management tools. This rapid pace of technological advancements requires high infrastructure costs, data quality, scalability, and the need for data centers capable of delivering these with operational simplicity.

By combining silicon, hardware, software, management tools, and optics, the Cisco Nexus

data center network delivers optimal bandwidth, latency, congestion management mechanisms, telemetry, visibility, and automation needed to build an AI/ML network fabric. Using the Cisco Nexus 9000 series, AI/ML fabrics can be tailored to build based on the desired technology, workload, and application with scalability ranging from few tens to thousands of GPU clusters, while providing power efficient, lossless, and low job-completion times networks.

Built on Cisco Cloud Scale or Silicon One technology, and an array of Cisco Nexus 9000 series switches with industry-leading Cisco NX-OS software and the Cisco Nexus Dashboard operations and automation platform, Cisco's data center solutions are easy to automate and deploy for AI training and inference models.

Benefits

- Build a non-blocking, high bandwidth, low latency, lossless ethernet AI/ML network with Cisco Nexus 9000 series' portfolio of [high-speed switches](#).
- Simplify Day 0 provisioning and automation of AI/ML networks with the [Cisco Nexus Dashboard Fabric Controller](#).
- Granular visibility to congestion score, real-time telemetry, and troubleshooting through [Nexus Dashboard Insights](#).
- Scale AI/ML fabric from tens to thousands of GPUs with [Cisco validated design for data center networking](#).
- Ensure seamless interoperability and consistent performance with leading GPU vendors. Cisco is a founding member of the Ultra Ethernet Consortium (UEC), and Cisco Nexus 9000 platforms are forward compatible with future specifications.

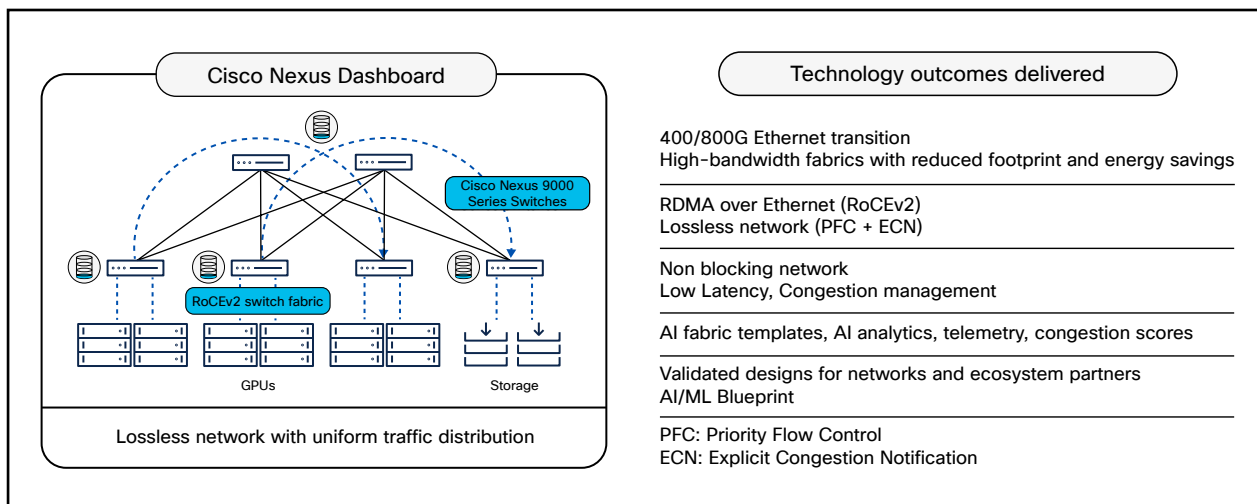


Figure 1. Transform Infrastructure for AI enterprise data centers

Trends and challenges

The applications, workloads, and size of AI/ML clusters are influenced by factors such as the size of datasets, desired job completion times, type of model, and the Total Cost of Ownership (TCO).

These considerations lead to challenges that impact deployment, adoption, and success of AI/ML fabrics:

- **Silos in network fabric:** Within a data center, this could lead to various designs of data center networks for AI-specific and non-AI fabrics. Relying on proprietary protocols such as InfiniBand for AI fabric would require the use of specialized hardware, software, management tools, training of resources, or personnel trained in this new technology.
- **Evolving needs of AI/ML networks:** As datasets grow, scaling of AI models across multiple clusters for compute, networking, and storage becomes necessary. This operational complexity leads to higher costs in power, cooling, and maintenance.
- **Integration with existing data center:** Integrating AI/ML networks with existing software ecosystems, databases, tools, and processes can result in additional development costs.

Using Cisco Nexus 9000 solutions for existing data center and AI/ML fabric enables leveraging the same silicon, hardware, software, and management tools with the ability to scale from tens to thousands of GPUs. In addition, the forward compatibility with UEC standards allows standardization of the AI/ML fabric transport protocols.

How it works

- **Role of GPU:** The complexity of machine learning algorithms, collective communication, and Remote Direct Memory Access (RDMA) operations remain within the end devices hosting the GPUs. The network remains unaware of these details and transfers the UDP/IP packets to the destination using its standard forwarding mechanism. Although this simplifies the networks for AI/ML infrastructure, the continuous synchronization of the GPU states while calculating patterns in vast data sets create unprecedented traffic patterns. For example, the GPUs may send and receive traffic at line rate.
- **Designing a network for AI:** Smaller GPU clusters can use a single-switch network. For example, a 64-GPU cluster can be interconnected using the [Cisco Nexus 9364D-GX2A](#) switch.

For large GPU clusters, a spine-leaf network is the optimal design because of its consistent and predictable performance and scalability. The edge switchports (that connect to the GPU NIC ports) should operate at the highest available speeds, such as 400 GbE on the Cisco Nexus 9364D-GX2A switch. The core switchports between the leaf switches and spine switches should match or exceed the speed at which the GPUs connect to the network. The inter-GPU network should not be oversubscribed. For example, on a 64-port leaf switch, if 32 ports connect to the GPUs, the rest 32 ports should connect to the spine switches.

- **Handling network congestion:** Cisco Nexus Series Switches isolate a faulty GPU Network Interface Card (NIC) port that remains unresponsive and unable to receive traffic continuously for a longer duration. After some time, if the GPU NIC recovers, the switch detects this change and allows this NIC port to communicate via the network by de-isolating it. Both the isolation of an unresponsive NIC port and the de-isolation of a healthy NIC port are automatically performed by the Nexus switches. This feature, called Priority Flow Control Watchdog (PFC Watchdog), eliminates the source of congestion from the network, thereby ensuring other GPUs in the networks are not affected by the issue.

Cisco Nexus Dashboard simplifies configuring the lossless networks by enabling PFC on all network ports using a single check box. It has ready-made templates with optimized buffer thresholds to fine-tune the sending of the Pause frames by the switch ports during congestion. Further optimizations can be applied uniformly across the network by changing the templates with refined thresholds.

Cisco Nexus Dashboard simplifies the interpretation of various congestion symptoms like the number of Pause frames sent or received, the number of Explicit Congestion Notification (ECN) marked packets, and traffic sent and received on a link, by calculating a congestion score and an automatic assignment to mild, moderate, or severe categories. The automatic correlations in Nexus Dashboard eliminate long troubleshooting cycles by accurately detecting the source and cause of congestion. This real-time visibility simplifies the buffer fine-tuning by showing fabric-wide congestion trends and providing the ability to consistently apply the fine-tuned thresholds across the network.

- **Managing job completion time:** QoS features of Cisco Nexus Switches allow prioritizing or using dedicated queues for small-sized response or acknowledgment packets so that they are not delayed while waiting behind the large-sized data packets in the same queue. Multiple no-drop queues or priority queues can also be used based on the traffic profile from the GPU NICs. Cisco Nexus Dashboard automatically detects unusual spikes in network delay and flags them as anomalies. This latency information is provided at a flow granularity and correlated with other relevant events such as bursts and errors. Getting real-time visibility into the network hot spots gives enough insights to a user to take corrective actions like adjusting the QoS thresholds, increasing network capacity, or using an improved load-balancing scheme. To apply these actions uniformly across the network, Nexus Dashboard offers customizable templates and full-featured APIs.

Use cases

Table 1. Use cases

Industry name	Use case description
Autonomous driving	<ul style="list-style-type: none"> ▪ Process vast amount of data from sensors ▪ Low latency to make quick decisions
E-commerce	<ul style="list-style-type: none"> ▪ Chatbots to handle inquiries from customers ▪ Analyze consumer behavior and make recommendations of products or services
Service providers	<ul style="list-style-type: none"> ▪ Reduce congestion and provide faster service ▪ Identify potential network issues and downtime
Finance	<ul style="list-style-type: none"> ▪ Quant based fast training networks ▪ Process transactions and market trend with low latency



Services

Cisco offers a wide range of services to help accelerate your success in deploying and optimizing the Cisco Nexus 9000 solution for your AI/ML data center networks. The innovative Cisco Services offerings are delivered through a unique combination of people, processes, tools, and partners and are focused on helping you increase operational efficiency and improve your data center network. Cisco Advanced Services uses an architecture-led approach to help you align your data center infrastructure with your business goals and achieve long-term value. Cisco SMARTnet™ Service helps you resolve mission-critical problems with direct access at any time to Cisco network experts and award-winning resources.

Learn more

- [Cisco AI/ML Blueprint for data center networks](#)
- [Cisco data center networking](#)
- [Cisco Nexus Dashboard](#)
- [Intel Gaudi Platform | Intel® Gaudi® 2 AI Accelerator](#)

Embracing new solutions and scalability is key to de-risking AI infrastructure. The AI accelerator ecosystem is evolving, with innovations in GPUs, CPUs, and specialized AI hardware. Avoiding proprietary technologies enhances flexibility and reduces costs.

For instance, using Ethernet for GPU interconnections allows adaptation to new advancements. Cisco Nexus 9000, certified to support the Intel® Gaudi 2 AI accelerator, exemplifies this. Cisco tested Gaudi 2 servers with Nexus 9000 Ethernet fabric, ensuring performance and compatibility. The team conducted months of extensive testing with Cisco Nexus 9364D-GX2A switches on the network and Intel Gaudi 2 accelerators on the compute. Documented configurations provide a deployment guide, resulting in high performance and reliability. With this, Cisco is helping customers expedite their return on investments by validating GPU clusters for enterprise data center environments.

Cisco Capital

Financing to help you achieve your objectives

Cisco Capital® can help you acquire the technology you need to achieve your objectives and stay competitive. We can help you reduce CapEx. Accelerate your growth. Optimize your investment dollars and ROI. Cisco Capital financing gives you flexibility in acquiring hardware, software, services, and complementary third-party equipment. And there's just one predictable payment. Cisco Capital is available in more than 100 countries. [Learn more.](#)

The Cisco Advantage

The wide array of Cisco Nexus 9000 fixed and modular series high-speed switches, combined with its industry-leading Cisco NX-OS software and Cisco Nexus Dashboard operations and automation platform, offers the simplest and most efficient way to deploy, scale, and analyze data center fabrics for AI training and inference models. With focus on speed, power efficiency, and operational costs, Cisco Nexus 9000 Series Switches enable organizations to build data centers that meet the evolving requirements of current and future AI workloads.