

Cisco Nexus HyperFabric

Contents

About Cisco Nexus HyperFabric	2
Capabilities of Cisco Nexus HyperFabric	2
Components of Cisco Nexus HyperFabric	3
Switches Supported by Cisco Nexus HyperFabric.....	3
Cisco Nexus HyperFabric Availability	3
Cisco Nexus HyperFabric for AI Clusters	4
Is Cisco Nexus HyperFabric part of Cisco Networking Cloud?	4



About Cisco Nexus HyperFabric

Q: What is Cisco Nexus® HyperFabric?

A: The high cost of hosting applications in the cloud, coupled with data-proximity and control requirements, drives customers to investigate reinvestment in on-premises and colocation hosting options. However, to meet the rapid changes of business priorities, and facing an IT skills shortage in the market, customers require agile data centers that are as easy to design, deploy, and scale as is offered by cloud-computing platforms.

Cisco Nexus HyperFabric is a new cloud-managed network fabric-as-a-service data center offer. Using a cloud controller managed by Cisco, it enables customers to easily design, deploy, and manage any number of fabrics located anywhere, spanning primary data centers, colocation facilities, and distributed data center edge sites. It reinvents the IT operations lifecycle of the data center by simplifying every step of the process and ensures repeatable and predictable outcomes by IT-generalist, application, and DevOps teams. The vertical stack solution consists of purpose-built hardware, software, cloud management, day-2 automation, and Cisco support.

Cisco and NVIDIA announced and are collaborating on an AI solution combining Cisco Nexus HyperFabric for AI clusters with NVIDIA GPU systems and an optional VAST storage platform. The converged Ethernet solution can easily be deployed and operated by research and data science teams, enabling them to focus on AI innovation rather than on IT operations.

Capabilities of Cisco Nexus HyperFabric

Q: What are the capabilities of Cisco Nexus HyperFabric?

A: Cisco Nexus HyperFabric is a plug-and-play cloud-managed data center networking solution that greatly simplifies the IT operations lifecycle of the data center:

- Cisco Nexus HyperFabric has a cloud controller, operated by Cisco, that is the single point of configuration, monitoring, and maintenance of all tenant customer fabrics using real-time connection to switches deployed on premises or in colocation facilities.
- Cisco Nexus HyperFabric fabrics consist of one or more cloud-managed Cisco 6000 Series Switches that immediately connect

to the cloud for automated zero-touch provisioning, and a “Helping Hands assistant” provides step-by-step cabling instructions combined with real-time verification.

- Cisco Nexus HyperFabric automation and operations are responsible for the cloud controller, the fabric underlay network, and the software upgrade process, while customers maintain direct control of all interconnections to their applications, hosts, and the rest of their network.
- Assertion-based monitoring continuously verifies the availability and reliability of the fabric and connected resources, and the root cause of issues detected are quickly identified.
- A built-in designer helps customers construct a validated fabric design based on desired host and port capacity, subscription, and environmental considerations including cabling and power, then creates an accurate Bill of Materials (BoM).
- Self-service fabric tenancies empower host and application teams to monitor and manage the fabric services they have been allocated, removing the need to depend on IT for most support services.



An API-first design lets organizations use DevOps tools including HashiCorp Terraform and Red Hat Ansible to fully automate the provisioning and operations of their environment.

Cisco Nexus HyperFabric for AI Clusters will extend the component list and include all the infrastructure elements needed to for an NVIDIA DGX cluster – NVIDIA GPU, DPU, and BlueField-3 SuperNIC components, Cisco optics, and optional VAST storage – in a single integrated solution that can be designed, ordered, deployed, and operated as a cohesive solution.

Components of Cisco Nexus HyperFabric

Q: What are the components of Cisco Nexus HyperFabric?

A: Cisco Nexus HyperFabric is a data center fabric-as-a-service solution that consists of two components:

- Cloud controller – a scalable, globally distributed multitenant cloud service that is used to design, plan, control, upgrade, and monitor fabrics using a browser or APIs
- Cisco 6000 Series Switches – installed with Cisco Nexus HyperFabric-managed software, the 6000 Series Switches connect to the cloud for centralized real-time visibility and control

High-performance fabric pods are constructed of switches using EVPN/VxLAN L3, L2 VLANs, and IPv4/IPv6 routing. The initial release will support fabrics of up to 2000 host-facing ports of 10/25/50GbE in mesh or spine-leaf topologies, while later releases will augment the scale of a single fabric.

Cisco Nexus HyperFabric for AI Clusters will extend the component list and include all the infrastructure elements needed to for an NVIDIA DGX cluster – NVIDIA GPU, DPU, and BlueField-3 SuperNIC components, Cisco optics, and optional VAST storage – in a single integrated solution that can be designed, ordered, deployed, and operated as a cohesive solution.

Switches Supported by Cisco Nexus HyperFabric

Q: What switches does Cisco Nexus HyperFabric support?

A: Two new Cisco 6000 Series Switches will be offered with the first release of Cisco Nexus HyperFabric, both are Cisco Silicon One™ Q200 1RU platforms.

- Cisco HF6100-60L4D, with 60x 10/25/50GbE SFP56, 4x 100/400GbE QSFP56-DD (16x via 100GbE breakout)

- Cisco HF6100-32D, with 32x 100/400GbE QSFP56-DD and 128x 100GbE via 400:100 breakout

The new Cisco 6000 Series Switches are dedicated for use exclusively by Cisco Nexus HyperFabric.

Additional Cisco 6000 Series Switches will be offered after the first release, including an 800GbE platform designed for AI deployments.

Cisco Nexus HyperFabric Availability

Q: When will Cisco Nexus HyperFabric be available?

A: Cisco Nexus HyperFabric is planned to be released at the end of 2024. Shortly after that date, a subsequent full-stack AI release of the product that is integrated with NVIDIA and VAST, optimized for AI use-cases, will be released.

Q: Does Cisco Nexus HyperFabric replace Cisco ACI or Nexus?

A: No, Cisco Nexus HyperFabric is a new data center network fabric-as-a-service for customers seeking a cloud-managed solution optimized for ease of use. Cisco ACI® and Cisco Nexus Dashboard provide customers with on-premises managed solutions featuring tremendous flexibility in the fabric design, configuration, and operations.

Cisco Nexus HyperFabric for AI Clusters

Q: What is Cisco Nexus HyperFabric for AI clusters?

A: Cisco Nexus HyperFabric for AI clusters is a premium offering of Cisco Nexus HyperFabric specifically catered to AI workloads. It is a complete AI networking solution that includes the latest in GPU/DPU technology from NVIDIA along with the latest Ethernet networking and management solutions from Cisco to create an easily deployable, scalable, and manageable AI/ML solution. Customers can easily and reliably deploy Cisco Nexus HyperFabric for AI clusters with this converged Ethernet solution and benefit from reduced operational overhead. Cisco Nexus HyperFabric for AI clusters is best suited for customers looking to build out their private cloud AI infrastructure.

Q: Why are you introducing the Cisco Nexus HyperFabric for AI clusters offering now?

A: The adoption of AI is increasingly common. “By 2027, 40% of enterprises will deploy GenAI network fabrics to enable cost and performance optimized support for AI workloads in their own data centers” (IDC Perspective, March 2024). Our findings show us that 95 percent of customers are aware that AI will increase workloads, but only 17 percent are equipped to handle this complexity, with 23 percent having limited or no capacity to meet the AI demand with current infrastructures. With the industry trending toward AI, Cisco is working closely with NVIDIA and our partners to ensure customers can rapidly and reliably deploy AI wherever it is needed by the business, and to simplify the IT lifecycle so IT-generalist, data-science, and DevOps teams can easily design, deploy, and operate the solution. The solution is based on a converged Ethernet network, so organizations can leverage their existing skills and processes.

Is Cisco Nexus HyperFabric part of Cisco Networking Cloud?

Q: Is Cisco Nexus HyperFabric part of Cisco Networking Cloud?

A: The vision of Cisco Networking Cloud is to create a simpler network management platform experience to help customers easily access and manage all Cisco networking products from one place. Cisco Nexus HyperFabric is part of that vision, by providing a new cloud-managed plug-and-play data center fabric-as-a-service solution. It comprises purpose-built hardware, software, automation, and support, simplifying data center networking for IT-generalist, application, and DevOps teams.