

Cisco Nexus Hyperfabric

Contents

Overview and definitions	3
Deployment guidelines	5
How to buy?	8

Overview and definitions

Q. What is Cisco [Nexus Hyperfabric?](#)

A. Cisco Nexus Hyperfabric is a new cloud-managed network fabric data-center solution delivered as a service. Using a cloud controller managed by Cisco, customers 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.

Q. What are the components of Cisco Nexus Hyperfabric?

A. Cisco Nexus Hyperfabric 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: Silicon One-based platforms leveraging six years of Cisco SONiC innovation and experience, they connect to the cloud for centralized real-time visibility and control.

A. Additionally, Cisco Nexus Hyperfabric AI will extend the component list and include all the infrastructure elements needed for managing AI clusters with Cisco compute, NVIDIA GPUs, BlueField-3 DPU and SuperNIC, Cisco optics, and optional VAST storage – in a single integrated solution that can be designed, ordered, deployed, and operated as a cohesive solution.

Q. What are the components of Cisco Nexus Hyperfabric?

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

- It cloud controller, operated by Cisco, serves as a single point for configuration, monitoring and maintenance of all tenant customer fabrics, using real-time connection to switches deployed either on premises or in colocation facilities.
- Network fabrics consist of one or more cloud-managed Cisco 6000 Series Switches that offer automated zero-touch provisioning, and an onsite mobile web portal which provides step-by-step cabling instructions combined with real-time verification.
- In a shared-responsibility model, automation and operations from Cisco support manage the cloud controller, the fabric underlay and overlay networks, 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 any issue detected is immediately identified.

-
- A built-in designer helps customers construct a validated fabric design based on desired host and port capacity, oversubscription, and environmental considerations including cabling and power, and 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.

A. High-performance fabric pods are constructed from switches using L3, EVPN/VXLAN, L2 VLANs, and IPv4/IPv6 routing.

A. Cisco Nexus Hyperfabric has an API-first design that lets organizations use DevOps tools, including HashiCorp Terraform and Red Hat Ansible, to fully automate the provisioning and operations of their environment.

Q. What switches does Cisco Nexus Hyperfabric support?

A. Three models of Cisco 6000 Series Switches are available with Cisco Nexus Hyperfabric utilizing Silicon One® Q200 Processor platforms.

- Cisco HF6100-60L4D with 60x 10/25/50GbE SFP56, 4x 100/400GbE QSFP56-DD (16x through 100GbE breakout)
- Cisco HF6100-32D with 32x 100/400GbE QSFP56-DD and 128x 100GbE through 400:100 breakout
- Cisco HF6100-64ED with 64 100/200/400/800 GbE OSPF interfaces and breakout options for greater densities of 100G, 200G and 400G

A. These Cisco 6000 Series Switches are dedicated for use exclusively by Cisco Nexus Hyperfabric, including the 800GbE platform designed for AI deployments.

Q. Why did Cisco introduce Cisco Nexus Hyperfabric?

A. The high cost of hosting applications in the cloud coupled with data-proximity and control requirements is driving many customers to investigate reinvestment in on-premises and colocation-hosting options. However, to meet the rapid changes in business priorities, and with an IT skills shortage in the marketplace, businesses require agile data centers that are as easy to design, deploy, and scale as cloud-computing platforms can provide. Cisco Nexus Hyperfabric is designed to meet these needs with an operational model that is anchored in simplicity.

Q. Does Cisco Nexus Hyperfabric replace Cisco ACI® or Cisco Nexus fabrics?

A. No, Cisco Nexus Hyperfabric is not a replacement for Cisco ACI and/or Cisco Nexus fabrics. It is a new data-center network fabric-as-a-service for customers seeking a cloud-managed solution optimized for IT generalists and ease of use. Cisco ACI and Cisco Nexus fabrics provide customers with on-premises managed solutions featuring tremendous flexibility in the fabric design, configuration and operations.

Q. What is Cisco Nexus Hyperfabric AI?

A. Cisco Nexus Hyperfabric AI is a premium offering of Cisco Nexus Hyperfabric specifically catering for AI workloads. It is a complete AI 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 AI clusters with this converged Ethernet solution and can benefit from rapid time-to-value and reduced operational overhead. Cisco Nexus Hyperfabric AI is best suited for customers looking to build out their private cloud AI infrastructure.

Q. Why are you introducing a Cisco Nexus Hyperfabric AI offering now?

A. The adoption of AI is increasingly viewed by organizations as a key enabler to drive innovation: “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 increase, with 23 percent having limited or no capacity to meet the AI demand with current infrastructures. Given the significant business value and growth in AI adoption across enterprises, Cisco is working closely with NVIDIA and VAST to ensure that customers can rapidly and reliably deploy AI wherever it is needed. This is in addition to simplifying the IT lifecycle, thus enabling IT generalists, data-science teams, and DevOps teams to easily design, deploy, and operate both AI and non-AI data center fabrics. The solution is based on a converged Ethernet network, so organizations can easily support and scale by leveraging existing skills and processes.

Deployment guidelines

Q. Is the Cisco 6000 Series Switch the same as the old Cisco Nexus 6000 switch family?

A. No, they are not the same. The old Cisco Nexus 6000 switches have reached end-of-life and end-of-support and are not supported or compatible with Cisco Nexus Hyperfabric.

Q. Can Nexus Hyperfabric manage Cisco Nexus 9000 Series Switches?

A. Cisco Nexus Hyperfabric does not currently support Cisco Nexus 9000 Series Switches, the new Cisco 6000 switches are required. Support for Cisco Nexus 9000 switches is being evaluated.

Q. Can I run Cisco ACI or NX-OS software on a Cisco 6000 switch?

A. No, the Cisco 6000 switches do not run Cisco ACI or NX-OS software. They ship with preinstalled software that cannot be modified by the customer. The switches are designed to deliver secure, predictable, and reliable performance, and customers are not permitted to install their own software. This is ensured through a locked bootloader and hardware-based root of trust during cloud and peer authentication.

Q. At a high level, how does Cisco Nexus Hyperfabric work?

A. Customers log in to Nexus Hyperfabric to begin building a validated fabric design tailored to their desired host and port capacity, oversubscription, and environmental considerations such as cabling and power. Customers then define the Layer-2 and Layer-3 networks the fabric exposes to hosts, maps them to ports, and specifies the routing needed to connect the fabric to the rest of the network. Nexus Hyperfabric integrated with Cisco ordering tools to guarantee there are no errors when converting a design into a bill of materials. When the Cisco 6000 switches arrive on-site and are

deployed, they automatically connect to the cloud to be claimed and provisioned by the cloud controller with a zero-touch plug and play. This process, resulting in a fully operational network fabric, takes just minutes. Assertion-based monitoring of availability and reliability of the fabric and connected resources are continuously verified, and the root cause of any issue detected is immediately identified. Later, if it is necessary to change the capacity or shape of the design, customers can redefine the in-flight design, approve the changes, and follow the entire process again. The product provides guidance for all the physical changes needed to migrate the old topology to the new desired state, including cabling adjustments, and it reconfigures itself automatically.

Q. How can I connect multiple fabrics with Cisco Nexus Hyperfabric?

A. Cisco Nexus Hyperfabric is designed to interoperate with existing Ethernet fabrics and Layer-3 networks. It delivers support for:

- Northbound-routed peering through BGP and static routes for both IPv4 and IPv6 address families. Nexus Hyperfabric supports multi- VRF routing for concurrently peering multiple virtual-routing instances with upstream routed networks.
- East/ westbound peering with existing Layer-2 Ethernet-based networks and fabrics using multichassis links

Q. What type of topologies can I deploy with Nexus Hyperfabric?

A. Nexus Hyperfabric is designed to support EVPN/VXLAN fabrics in scalable leaf/spine topologies as well as in smaller full-mesh topologies (up to 4x leaf-only switches, using the HF6100-60L4D switches).

- Full-mesh topologies:
 - Up to 4x HF6100-60L4D leaf switches (with no spine switch)
 - Up to 300x 10/25/50Gb SFP+ host ports
- Leaf/spine topologies:
 - Up to 32x HF6100-60L4D leaf switches
 - Up to 4x HF6100-32D spine switches
 - Up to 1920x 10/25/50Gb SFP+ host ports

A. Nexus Hyperfabric also allows flexible connection modes, allowing for host connections and northbound router gateway ports on any switch within the fabric, allowing for streamlined connectivity for Layer-3 peering and high-speed host connections (40/100/200/400Gb) on the fabric.

A. The HF6100-64ED switch is a good fit for AI deployments for the backend network connecting the GPU servers. It may also be used in traditional leaf/spine topologies and allows for greater density of leaves than the HF6100-32D.

Q. What deployment use cases does the first release support?

A. The current release is designed to support the following use cases:

- New or refreshed data centers of modest size (less than 1000 servers), because of the solution's ease of use, agility, and self-service capabilities for fabric tenants
- Data centers in colocation facilities, because of the solution's cloud-based management and "Helping Hands" assistance
- Edge or remote data centers, because of the solution's centralized management in the cloud, mesh-topology support, and plug-and-play deployment

A. Providers managing multiple customer data centers, because of the solution's

Q. Where is the Nexus Hyperfabric Cloud Controller hosted?

A. The Nexus Hyperfabric Cloud Controller is maintained by Cisco and is hosted in the public cloud. The cloud controller includes global scalability, as well as multi-region redundancy, without any additional configuration for end-users.

A. The Nexus Hyperfabric Cloud Controller is reachable through a single URL (<https://nexushyperfabric.cisco.com>) and is used for all communications (including switch-management cloud connectivity, primary user interfaces, and RESTful API endpoints).

Q. What are the plans for the cloud controller hosting outside the United States?

A. Cisco plans to offer local cloud controllers hosted natively in Europe and Asia in future releases.

Q. Is Cisco Nexus Hyperfabric strictly cloud-based, or are there any plans for air-gapped environments?

A. The cloud controller is managed by Cisco operations staff. There are no plans to hand off that responsibility to entities not owned or operated by Cisco.

Q. What are the plans for government cloud compliance (under the Federal Risk and Authorization Management Program [FedRAMP])?

A. Cisco is investigating offering local cloud controllers hosted in FedRAMP in a future release and will consider other government cloud-compliance environments.

Q. Will Nexus Hyperfabric have integrations with other products?

A. HashiCorp Terraform and Red Hat Ansible can be used to fully automate provisioning, and telemetry may be sent to a variety of data collectors, including Amazon S3, Cisco Splunk, and ServiceNow. The Nexus Hyperfabric controller is designed to be operated API-first, and includes a native RESTful API designed to allow flexible integrations both for provisioning tools such as Ansible or Terraform and for external network-and incident-management tools.

Q. How can I export logging to my own logging platforms?

A. The Nexus Hyperfabric Cloud Controller supports external cloud-to-cloud integrations for delivering logging information through external Amazon S3 buckets to logging files in syslog format. In addition, the cloud controller allows API integration for querying internal states, alerts, and other telemetry data from the on-premises fabrics.

Q. What are the bandwidth requirements for cloud connectivity?

A. Bandwidth per switch is less than 2Mb/sec at steady state. Certain types of real-time monitoring may increase upstream bandwidth when initiated by end users through the cloud controller UI.

Q. How is real-time telemetry uplifted to the cloud controller?

A. The Cisco 6000 Series Switches include a telemetry agent that establishes an outbound TLS connection to the cloud controller. This outbound TLS (TCP/443) session is designed to work with existing network security controls for standard web clients and provides all connectivity required for configuration, monitoring, and real-time telemetry for each switch managed by Nexus Hyperfabric.

Q. What happens if I lose cloud connectivity?

A. The cloud connection from a deployed fabric to the cloud controller is used for management and telemetry purposes only. All stateful protocols are maintained independently on the switches residing on premises. Any fabric disconnected from the cloud will continue to operate normally (including local underlay and overlay fabric-management protocols, external peering protocols such as BGP, and all standard bridging functions).

A. When disconnected from the cloud, a fabric will not be able to accept configuration updates, and telemetry will not be received by the cloud controller. The on-switch agent will intelligently buffer telemetry if disconnected and will upload this information and resynchronize any configuration updates once the fabric is reconnected to the cloud controller.

Q. How can I interact with the Nexus Hyperfabric API?

A. The Nexus Hyperfabric Cloud Controller implements a RESTful API using the same URL as standard management functions do. This API is extensively documented in an online API programming guide and is designed to be compatible with OpenAPI specifications.

A. In addition, Nexus Hyperfabric supports providers for both Ansible and Terraform to help users exercise the Nexus Hyperfabric API according to their existing management tooling.

How to buy?

Q. How do I buy Cisco Nexus Hyperfabric?

A. Cisco Nexus Hyperfabric may be purchased from a certified Cisco Nexus Hyperfabric reseller. Organizations able to purchase products directly from Cisco may also purchase the solution.

Q. What is the packaging and licensing model?

A. Cisco Nexus Hyperfabric is packaged as a vertical stack subscription offer. For every Cisco 6000 switch purchased, customers are required to buy a subscription service that includes the software, software entitlement, Cisco cloud management, day-2 automation, Cisco® TAC, and hardware support. The minimum subscription term is three years.

A. The Essentials package is an especially convenient solution for fabrics consisting of small numbers of switches. All the components needed are included in the service; there is no management tooling to deploy locally since the fabrics are fully managed from the cloud.

A. Cisco Nexus Hyperfabric AI, when released, will follow the same model. However, in addition to purchasing the Cisco 6000 switches with their associated, and new, Premier subscription service entitlements, customers also purchase Cisco UCS® servers that host NVIDIA components (GPUs and BlueField DPUs), optional UCS servers that host a VAST storage product, and Cisco® optics. The entire solution is expected to be available on Cisco's price list.

Q. When will Cisco Nexus Hyperfabric be available?

A. Cisco Nexus Hyperfabric was released in early 2025 and Cisco Hyperfabric AI was released later in the year.

Q. Can I design a data center fabric before buying it?

A. Yes. Anyone with a Cisco Connection Online (CCO) account or with an Identity Provider (IdP) tied to CCO may request and receive an organization tenancy in the cloud controller. Within the tenancy, customers may design blueprints that contain all the details needed to order, deploy, configure, and operate the fabric, including:

- Physical components including switches, optics, and connectors, air-flow direction, power consumption, and a cabling plan
- A bill of materials for all the Cisco components that is the source-of-truth and integrated with Cisco Commerce Workspace (CCW) to automate the process of obtaining an accurate quote
- The logical network including the entire network overlay and underlay and upstream route peering
- Host-port assignments to server infrastructure teams
- API integration to automate the provisioning of the blueprint

A. Additional accounts may then be added to that tenancy so teams can collaborate. Once the equipment is deployed, it is automatically provisioned according to the blueprint.

Q. What are the subscription entitlement tiers?

A. A subscription entitlement is needed for every Cisco 6000 switch that is deployed and used. Subscription entitlements may be initially purchased for three, five, or seven years of operation, and they may be renewed. The subscription feature entitlement tiers are based on fabric use cases. Currently there are two packages: the first is for standalone fabrics ("Essentials") and the second

(“Premier”) is for AI fabrics. All the switches in a fabric must use the same entitlement tier; however, an organization may concurrently manage multiple fabrics that use different entitlement tiers.

Table 1. Subscription entitlements

Features	Essentials entitlement	Premier entitlement (only for Nexus Hyperfabric AI)
Cisco support 8x5xNBD	Yes	Yes
Cloud controller	Yes	Yes
Designer (no purchase required)	Yes	Yes
Cloud-driven software upgrade	Yes	Yes
BOM generation with optics	Yes	Yes
Onsite mobile deployment assist	Yes	Yes
Plug-and-play deployment	Yes	Yes
Leaf/spine topologies	Yes	Yes
Mesh (spineless) topologies	Yes	Yes
EVPN/VXLAN underlay (opaque)	Yes	Yes
Static and BGP routing	Yes	Yes
MLAG	Yes	Yes
Supports RDMA over Converged Ethernet v2 (RoCEv2)	Yes	Yes
Real-time cloud-accessed telemetry	Yes	Yes
IPv4 and IPv6	Yes	Yes
Assertion-based monitoring	Yes	Yes
Survivable data and local management plane	Yes	Yes
Hardware-based attestation and security	Yes	Yes
API for headless provisioning and monitoring	Yes	Yes
AI use-case support	Yes	Yes
Aligned to NVIDIA Enterprise Reference Architecture	No	Yes
NVIDIA Adaptive Routing on Backend Switches	No	Yes

Deploy AI-validated blueprints built into workflow	No	Yes
Automatically provisions lossless backend AI and storage networks	No	Yes
Internal performance monitoring between switches and servers	No	Yes
Monitors AI and storage server resources	No	Yes
Options for AI servers	Cisco UCS or bring-your-own	Bundled Cisco UCS

A. In a future release, an advanced Advantage entitlement is expected to be offered, with enterprise features such as advanced security and large scale deployment capabilities.

Q. Can Cisco Nexus Hyperfabric be sold only by qualified resellers?

A. Cisco Nexus Hyperfabric has a reseller qualification program and any organization that has completed that program may sell the product. Non-qualified resellers may also sell Cisco Nexus Hyperfabric; however, they must include deployment services from a qualified organization as part of the solution. Where can I get more information?

More information may be found in the [data sheet](#).

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

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: <https://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)