Cisco Nexus Hyperfabric

Overview and definitions

Q: What is Cisco Nexus® One?

A: Cisco Nexus One is the next-generation data center networking architecture that unifies every fabric — AI, enterprise, cloud, storage, media, and edge — into one consistent, interoperable system. Nexus One is about standardization, flexibility, and choice. Nexus One embraces and extends what you run today; nothing is ripped out or replaced.

The architecture spans four layers: silicon (Cisco® Cloud Scale, Cisco Silicon One®, and NVIDIA Spectrum-X Ethernet), systems (Cisco N9000 Series Switches as the universal hardware anchor, with industry-leading Cisco Optics), software (Cisco NX-OS, Cisco ACI®, and SONiC), and a consistent operating model (Cisco Nexus Dashboard on premises or Cisco Nexus Hyperfabric™ cloud SaaS). Security is fused into every layer, and operations are simplified for every operator — human, Network as Code for automation, or AI agent.

Q: What is Cisco Nexus Hyperfabric?

A: Cisco Nexus Hyperfabric is the transformative cloud-managed operating model for Nexus One. It is a 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 N9300, Cisco N9100, and Cisco 6000 series switches, leveraging six years of Cisco SONiC innovation and experience, they connect to the cloud for centralized real-time visibility and control.

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.

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

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:

- A 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 that consist of one or more cloud-managed Cisco 6000 Series Switches, Cisco N9100 Series Switches, or Cisco N9300 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.

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

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: These models of Cisco 6000 Series, Cisco N9100 Series, and Cisco N9300 Series Switches are available with Cisco Nexus Hyperfabric .

- 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
- Cisco N9K-9336C-SE1 -36x QSFP28
- Cisco N9K-93180YC-FX3 48x SFP28 + 6x 100G QSFP28
- Cisco N9K-93108TC-FX3 48x 100M/1G/10G Cu + 6x 100G QSFP28
- N9164E-NS4-O 800G Spectrum 4

The Cisco N9100 series and N9300 series switches are designed to address the need for high-performance, power-efficient, and compact switching in the networking infrastructure. Also, for Cisco Nexus One, they provide multiple operating models including Nexus Hyperfabric for cloud-managed needs and Cisco Nexus Dashboard for on-premises needs.

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

Q: What makes Cisco Nexus Hyperfabric different?

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 a Nexus One 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 data-center network fabric-as-a-service for customers seeking a cloud-managed operating model 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. They are all part of the Nexus One unified management plane that provides open choices with no compromises.

Q: What type of solution does Cisco Nexus Hyperfabric provide for AI?

A: Cisco Nexus Hyperfabric delivers outstanding results for your AI networking needs. You have two clear options to support your AI initiatives:

Turnkey full-stack AI infrastructure: For the fastest time to first token, select a high-performance, cloud-managed, turnkey full-stack AI infrastructure. This option is fully compliant with NVIDIA Enterprise Reference Architecture (ERA) or can scale to a cloud-reference architecture – ideal for rapid deployment and simplified management with visibility to the GPU. See the data sheet.

- **Bring your own AI:** Alternatively, choose to bring your own AI. With this approach, you gain the flexibility to innovate by using a cloud-managed network fabric that supports your preferred compute, GPU, AI software, and storage – empowering you to tailor your AI infrastructure to your unique requirements.

Q: What is Cisco Cloud Control?

A: Cisco Cloud Control unifies the Cisco estate into a single-pane-of-glass environment for collaborative troubleshooting, monitoring, and single-sign-on management across the Cisco portfolio. It gives teams and AI agents shared context for managing infrastructure and resolving issues across products and IT domains. This creates a practical path to agentic operations while helping organizations improve operational resilience, increase productivity, and gain more value from existing technology investments.

AI Assistant supports everyday work, while AI Canvas gives operators and AI agents a shared workspace for complex cross-domain investigation and resolution. Teams spend less time assembling evidence and coordinating tools, and more time understanding problems and restoring service.

Cisco Nexus Hyperfabric integrates seamlessly with Cisco Cloud Control to provide a unified, cloud-managed network fabric. This synergy delivers intelligent design, automated validation, and integrated monitoring, simplifying operations while ensuring high-performance connectivity across your data center.

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

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 now supports specific Cisco N9300 Series Switches and Cisco N9100 Series Switches.

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 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:

- 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/west 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 .

- Full-mesh topologies:
- Up to six Cisco N9000 Series switches providing 288 host ports

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

- Up to 4x HF6100-60L4D leaf switches (with no spine switch) providing up to 300x 10/25/50Gb SFP+ host ports
- Leaf/spine topologies:
- Up to 32x leaf switches
- Up to 16x spine switches

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.

The N9164-NS4-O and HF6100-64ED switches are good fits for AI deployments for backend networks connecting GPU servers. They may also be used in traditional leaf/spine topologies.

Q: What deployment use cases are supported?

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
- Managed service providers managing multiple customer data centers using the solution's inherent multitenant capabilities

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.

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 cloud controller hosting outside the United States?

A: Cisco hosts cloud controller service in the European Union as well as in the United States. Customers can choose in which geography to host their Hyperfabric organization.

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.

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

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, 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, Cisco N9100 Series Switches, and Cisco N9300 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).

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.

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.

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

How to buy?

Q: How do I buy Cisco Nexus Hyperfabric?

A: Cisco Nexus Hyperfabric may be purchased from any Cisco 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.

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.

Cisco Nexus Hyperfabric full-stack AI pod follows the same model. However, in addition to purchasing the Cisco switches with their associated, and new, Premier subscription service entitlements, customers also purchase Cisco UCS® servers that host NVIDIA components (GPUs and 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: 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 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

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 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.

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

Table 1. Subscription entitlements

Features	Essentials entitlement	Premier entitlement
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

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

Features	Essentials entitlement	Premier entitlement
Hardware-based attestation and security	Yes	Yes
API for headless provisioning and monitoring	Yes	Yes
AI use-case support	Yes	Yes
Cisco Cloud Control	Available	Available
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

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.

More information may be found in the data sheet.

Some products and features described are in development and offered on a when-and-if-available basis. Cisco reserves the right to change delivery timelines and will have no liability for any delays or failures to deliver.

© 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)