White paper
Cisco public

Nexus Hyperfabric Security

Overview of the Solution’s Security Architecture and Procedures

Contents

Overview and introduction

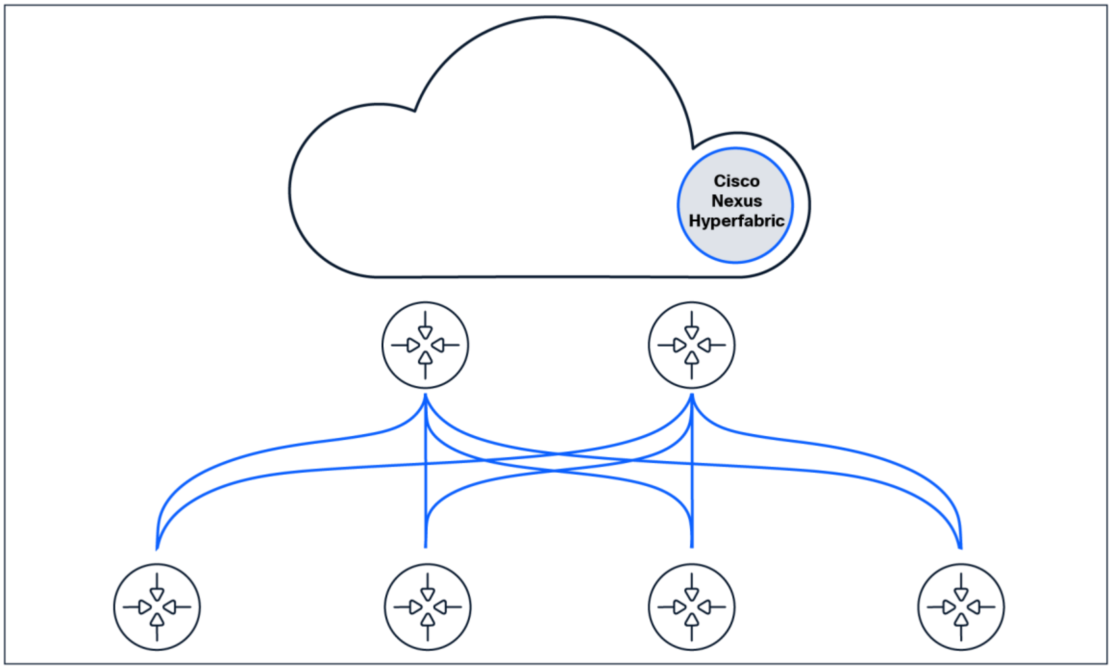

The Cisco Nexus® Hyperfabric cloud-management and switching platform offers a purpose-built, fully cloud-managed fabric-as-a-service networking solution, enabling true zero-touch provisioning of data-center EVPN-VXLAN fabrics. From the cloud-management platform to the switch, the architecture is crafted to establish a security-first operating model that protects your information and intellectual property while offering a best-in-class cloud management solution.

The following document explains how Nexus Hyperfabric is designed to deliver a new way to deploy and manage data-center fabrics in a secure and scalable manner. Traditionally, data-center fabrics were relegated to mainly large, centralized data-center or cold environments – which are secured by both physical and logical network-access controls. However, because many workloads repatriate from cloud-edge computing, data-center fabric requirements are expanding to account for a more distributed, highly available service infrastructure, expanding from a handful of large sites to many smaller locations as well. In this new world, a data-center fabric not only must support remote management but may not be able to rely on the same level of existing physical and network security controls.

Nexus Hyperfabric was designed from the ground up, to not only manage data-center fabrics in remote locations, but also to provide an unmatched level of hardening and integrity protection.

Nexus Hyperfabric security focuses on multiple layers:

- **Cloud service security** – secure and resilient operations of the cloud services, including measures taken to ensure customer devices, configuration, and data remain confidential and safe
- **Management access security** – role-based access control for both programmatic and user interfaces in the cloud platform
- **Managed device security** – purpose-built, physically hardened, tamper-resistant hardware platforms and secure software architecture

- **Fabric operational security** – controlled threat surface for the control and data plane of switches in operation

  For more information, including user guides, please see the following link: Cisco Nexus Hyperfabric Documentation.

Terminology

This document uses the following terms:

- **BGP** – Border Gateway Protocol is used for dynamic routing and endpoint location sharing.
- **EVPN** – Ethernet VPN is a BGP multi-protocol control-plane addition that provides MAC&IP routes for a VXLAN topology to route unicast traffic to the proper peer in the topology.
- **VXLAN** – Virtual eXtensible Local Area Network that overlays a Layer-2 (L2) network over a Layer-3 (L3) underlay infrastructure. It uses MAC-in-UDP encapsulation to extend Layer-2 segments across a Layer-3 network.
- **Fabric blueprint** – The fabric blueprint, in the Nexus Hyperfabric dashboard, includes the fabric design, L2 and L3 logical networks, switch binding, switchport configurations, and assurance data. It provides a framework to maintain consistent design and operation of a data center fabric.
- **Assertion** – Events and state in the Nexus Hyperfabric environment
- **Organization** – An organization is the customer's representation in the Nexus Hyperfabric platform composed of the fabric blueprints and Nexus Hyperfabric devices owned and operated by the customer. There is no data-sharing between organizations.
- **API** – Application programmable interface utilized for configuration and data collection through programmatic tools
- **Fabric exit-node** – A switch in the fabric providing connectivity to the cloud for other members proxied through the out-of-band management interface

Secure cloud management

The Nexus Hyperfabric solution, as a cloud-managed service, is designed to provide remotely managed data center fabrics ranging from controlled data center environments to remote sites with limited staff and security infrastructure. The solution prioritizes secure operations and management with the highest level of importance. The cloud-management platform implements comprehensive security measures to ensure not only high-availability through a resilient globally distributed architecture, but also robust processes that protect against potential threats that could lead to data loss or other serious consequences.

To help ensure simplicity of operations and predictable network behavior, the Nexus Hyperfabric dashboard is architected as the single-source of truth for all aspects of management; for everything from telemetry ingestion and processing to monitoring of the fabric blueprints that contain switch-fabric membership, the switching software, and the configurations related to all operations of the cloud-managed data-center EVPN-VXLAN fabrics.

The cloud-management console is secured exclusively through HTTPS/TLS encrypted channels, whether using the web GUI interface or the respective APIs that power it, and, by design, no form of unencrypted communication is permitted or utilized in any manner. The web interface leverages the same publicly available APIs used for programmatic monitoring and configuration, ensuring that all administrative surfaces are secured through a common framework.

On the ground, Nexus Hyperfabric uses Cisco 6000 Series Switches, which are designed to execute an outbound TLS connection to the cloud, making the solution a call-home architecture. This concept allows devices to establish a control channel to the cloud, enabling the cloud-service footprint to be reduced to a standard set of endpoints necessary for communication. The control channel is mutually attested by both the cloud and device, securing against misconfiguration as well as malicious activities, such as hijacking DNS or BGP. This concept provides a flexible and robust management architecture, where northbound security appliance rules only need to provide outbound HTTPS/TLS sessions. This is one of many aspects of the solution that will be explained in greater detail later in this document, starting with how the network devices are hardened through purpose-built hardware and software.

Network-device hardening

Hardening network devices can come in many forms, from software controls providing limited access to the platform, to deeply integrated security functions delivered through purpose-built hardware. Cisco 6000 Series Switches are built to provide the most robust protections against tampering, starting from the hardware carrying secure context and controls into the software environment in a seamless multifactored design. For remotely deployed fabrics, predictability is imperative to success. This purpose-driven design sets the stage for predictable boot behavior and validation of software as a required function to providing a trustworthy cloud-managed architecture.

Cisco 6000 Series Switches all contain Cisco's Trust Anchor Module (TAM), also referred to as a Hardware Security Module (HSM), to provide secure-boot services ensuring that the hardware and software loaded on the platform is vetted against expected boot behaviors. This process is known as secure boot, with measured boot attestation outlined in the following sections.

Secure boot with measured boot attestation

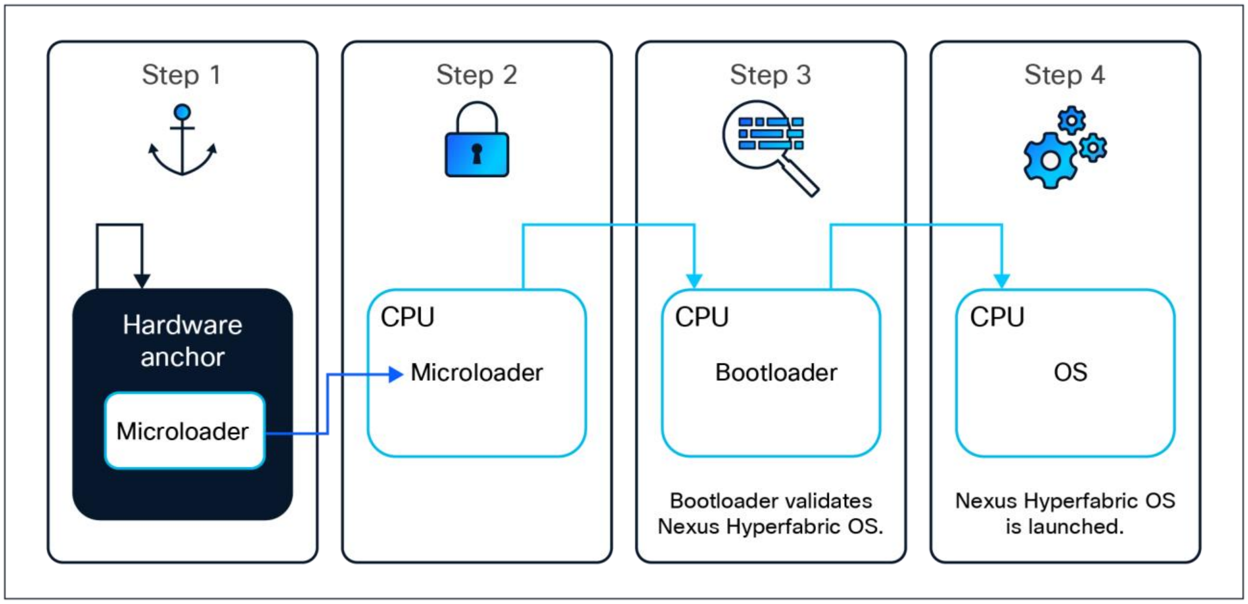

Figure 1.
Hardware anchored secure-boot.

Cisco 6000 Series Switches use secure boot with measured boot attestation as a framework to ensure the integrity and authenticity of the software running on the device. It begins with a hardware-anchored root of trust (Trust Anchor Module), which is a tamper-resistant hardware security module that validates the initial code (the microloader) before it is allowed to execute. This process establishes a chain of trust, where each subsequent component, such as the firmware and bootloader, is verified through digital signatures before execution.

Measured boot attestation extends this concept by recording the measurements of each component during the boot process, allowing for attestation and verification against known good values. These values are stored as Platform Configuration Registers (PCRs), and, as part of the software update procedure, PCRs are updated and then validated during boot prior to execution. This ensures that only unaltered software, signed by Cisco, is loaded, preventing the execution of malicious code. By integrating secure boot with measured boot attestation, Cisco 6000 Series Switches provide robust protection against unauthorized modifications in software and maintain the integrity of the system throughout its software lifecycle. In the event a modification was made to any aspect of the software, the switch will stop the boot process in a fail-closed behavior, protecting the topology and data within from modification or exfiltration.

For a review of Cisco's methodology and foundational design used for Nexus Hyperfabric hardware and software, please see the Cisco Trustworthy Technologies Data Sheet.

Hardening through an Immutable filesystem

The Nexus Hyperfabric solution not only performs attestation of the hardware and software on the switch but also utilizes an immutable filesystem as a security measure, ensuring that the core filesystem is not modified and is consistent in operation. All software packages are signed by Cisco, verified by the platform and cloud on every boot and at runtime, and provide no ability to run third-party vendor software packages on the platform. This architecture guarantees against various forms of "rootkit" infections to the operating system and maintains the closed-loop verification process necessary to deliver predictable behavior, which helps to optimize the experience of the product.

Securing device-to-cloud communication

The utilization of hardware-backed secure boot builds the framework for providing a secure communication channel that can be deployed over the public cloud.

One of the key facets to a secure cloud-management platform is ensuring that the communications between the controller and the devices it manages are not only encrypted but validated in multiple forms to protect against any potential threats to the networks, network devices, and services hosted in the cloud. This architecture is deeply rooted in the use of hardware-security modules and the treatment of all communications as untrusted (implemented as a zero-trust policy) by both the cloud platform and network devices. Validation is required on both sides of the connection before any connectivity is established, and mutual attestation is a constant throughout all aspects of the service lifecycle. This process is performed routinely in operations, protecting against any potential attempts at masquerading as (for example) a Nexus Hyperfabric switch or any aspect of the cloud service.

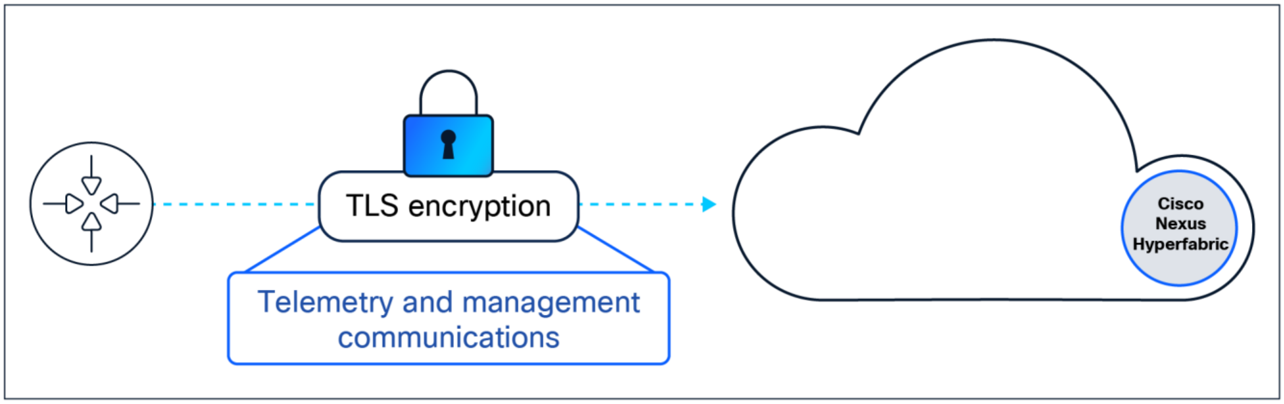

Figure 2.
Management channel secure communication.

Each Nexus Hyperfabric instance is equipped with at least one secure channel to the cloud, utilizing a standard TLS-encrypted outbound session to guarantee data privacy in transit. This secure channel serves as the sole network-device connection to the cloud, used for configuration management, telemetry, and the secure transmission of software and firmware upgrades. In addition to encryption in transit, measures are in place to maintain data privacy even in the event of man-in-the-middle decryption attempts against the management session.

Switch association to a fabric is orchestrated through the cloud controller, and no local administrative surfaces allow for admission of a switch into the fabric, negating any potential for a rogue switch to be locally added to a fabric, whether malicious or accidental. The communication channel undergoes mutual attestation between the ground (switch) and cloud (Nexus Hyperfabric controller), providing protected and streamlined, zero-touch provisioning and management of Nexus Hyperfabric instances.

Platform and cloud mutual attestation

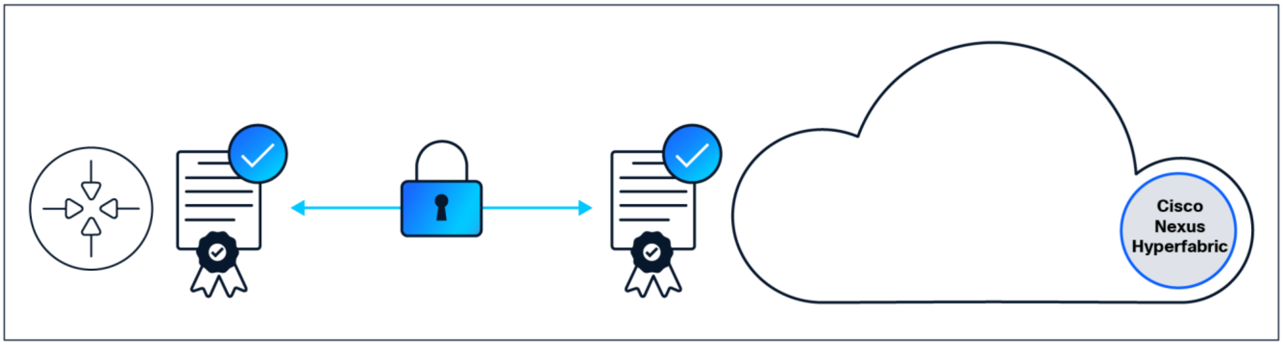

Figure 3.
Device and cloud mutual attestation.

The switches establish secure communication through a TLS-encrypted session to the Nexus Hyperfabric cloud, ensuring data integrity and confidentiality. The service architecture facilitates mandatory mutual attestation between the cloud service and switches, preventing interception and unauthorized access to any data in transit. The cloud challenges and verifies the authenticity of the switch by using a combination of the Secure Unique Device Identifier (SUDI) certificate and software fingerprint, confirming the switch's identity and software integrity. Simultaneously, the switch challenges and verifies that the cloud controller and connections are legitimate by performing certificate-chain validation through the switch's management agent, leveraging pinned certificates stored in the agent's certificate repository as well as identity verification through a second form of cryptographic validation of the cloud services. If either side rejects the challenge responses, the communication channel will not be formed.

The cloud controller certificates are rotated annually as a standard security practice. The switch's management agent performs attestation of the cloud endpoints through multiple forms of identity verification. The cloud service utilizes multiple distinct certificate authorities to ensure separation of risk surfaces. This design ensures that no single service can be compromised to establish any aspect of control or connectivity. The switch's repository of cloud certificates and identity fingerprints are updated through the management agent (which is regularly updated by Cisco in operation), ensuring that security can be maintained without requiring a new switch OS image installation to the fabric.

This robust security framework guarantees that both the cloud controller and switches are authenticated and authorized continuously, building the basis for a secure and reliable network infrastructure.

Nexus Hyperfabric services connectivity requirements

Nexus Hyperfabric is a publicly accessible service and is hosted under Cisco's base domain of **Cisco.com**. Hyperfabric's management across both the user interface, API, and management connectivity all leverage the domain: "**hyperfabric.cisco.com**."

The following protocols and ports are necessary for the service to operate:

- DNS – UDP/53
  - Customer-defined DNS servers utilized for the switch's out-of-band management connectivity DNS resolution
- HTTPS/TLS – TCP/443
  - Access to*.**hyperfabric.cisco.com**
    - As an example, all firmware is served from a sub-domain of **cdn.hyperfabric.cisco.com**

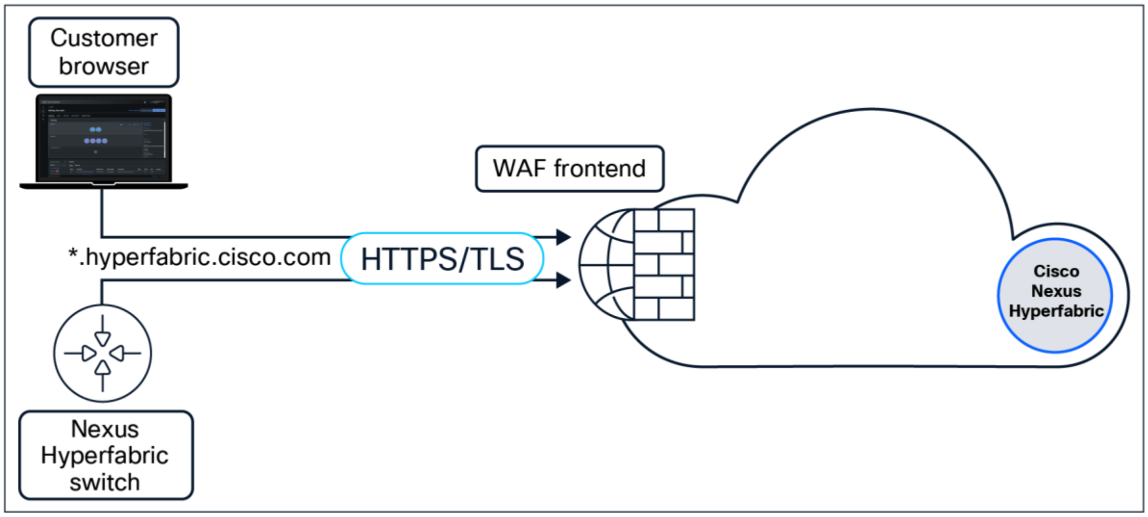

Figure 4.
All traffic is processed through the same ingress interface.

Cloud service IP addresses are subject to change over time and will differ based on regional usage. It is recommended to allow all communications through DNS/URL-based policies.

Out-of-band switch management

Every Cisco 6000 Series Switch is equipped with an out-of-band management interface used for establishing connectivity to the cloud. Every switch in the topology configured with an out-of-band connection to the cloud can provide access to other members of the fabric, which is explained in more detail in the section below, titled "**Fabric exit-node connectivity**." The out-of-band management interface is only used as a cloud connectivity "client" port, and the switch does not perform any dataplane routing functions over this interface. The interface only permits outbound connectivity from the switch to the cloud, Dynamic Host Configuration Protocol (DHCP), as well as Internet Control Message Protocol (ICMP) for local reachability validation. Configuration of the out-of-band management interface can be managed through the Nexus Hyperfabric dashboard, as well as the local serial console port through a menu-driven terminal user interface, described later in this document in the section titled "**Local administrative threat-surface hardening**."

Fabric exit-node cloud connectivity

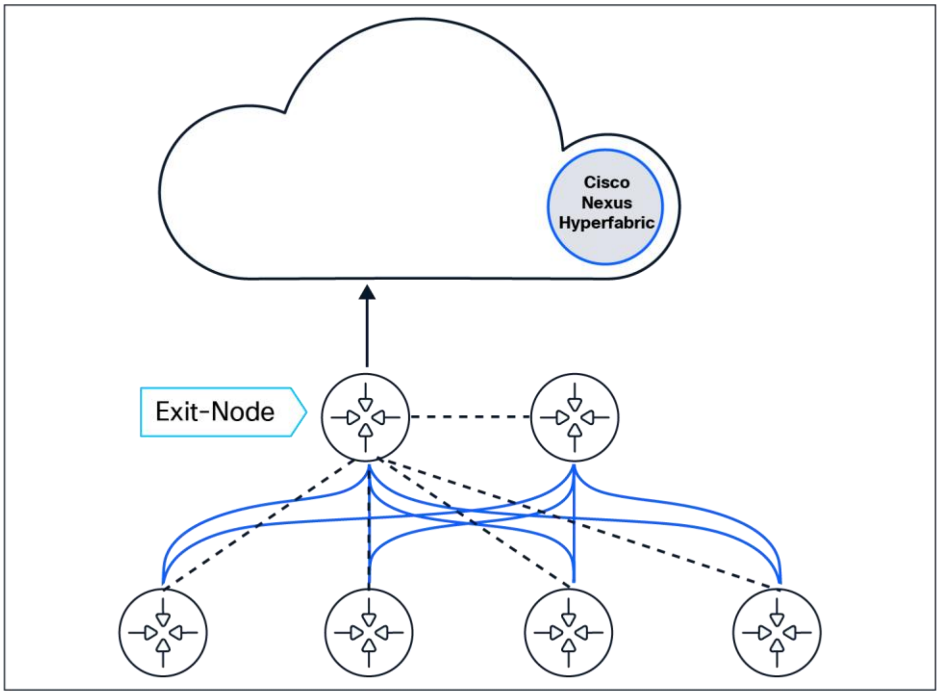

Figure 5.
Management traffic proxying through the exit-node.

In scenarios where not all switches can be directly connected to the cloud with a unique management IP address, the administrator is only required to provide cloud connectivity to the initial switch in the fabric, known as the exit-node. Each subsequent switch that is connected undergoes cryptographic validation (using the SUDI certificate embedded in the hardware security module) by its peer before communications to the cloud are proxied through the exit-node. Cisco 6000 Series Switches that are not bound to the fabric in the fabric blueprint are only permitted to communicate to the cloud; unbound switches are not permitted to join the underlay or participate in the fabric in any manner . This is designed to streamline secure on-boarding and results in substantially reducing the administrative overhead in a secure, cryptographically significant manner. In production environments it is recommended to provide at least two switches (preferably every spine in the topology) with out-of-band connections for resiliency in the event that any individual exit-node loses management port cloud connectivity or reboots during a scheduled upgrade.

Secure fabric operations

Remotely deployed fabrics may be used in locations that are not equipped to provide the physical security that is available in a dedicated data center. Addressing this issue has been a driving factor in the development of every aspect of the behavior of network devices, including how the fabric is operated. Building on a section above, titled "**Platform and cloud mutual attestation**," the fabric contains several security measures to safeguard against rogue switches being added to the fabric, local configuration modifications, or data exfiltration in the event a rogue switch is connected. Ownership of fabric membership is driven through the cloud controller using explicit configurations implemented by a Nexus Hyperfabric administrator. This configuration is rendered through the fabric blueprint in Nexus Hyperfabric's dashboard, which contains all fabric operation and configuration details and, by design, is never automated by the cloud. The configurations received by the fabric nodes from the cloud controller only permit switches bound to the fabric to form adjacencies and become members of the topology.

As explained in the previous section, when a new switch is connected to the topology, the peer discovers the newly connected switch through Link Layer Discovery Protocol (LLDP) and then performs SUDI-certificate verification before providing any cloud connectivity. Even then, the switch is not admitted to the fabric underlay, and the fabric's configuration is only provided to the device after the attestation process and fabric membership validation has concluded. Providing the parameters necessary to join a switch to the fabric topology is not possible to circumvent. If a switch is connected and not bound to the fabric, the switch is considered an untrusted entity, and traffic is limited to onboarding flows only.

Border Gateway Protocol (BGP) is the primary communication mechanism within the fabric, providing underlay and overlay control-plane functions. Once a switch is bound, BGP communications between peers are permitted through configuration. This is then backed by a secondary authentication process leveraging a cloud-generated and -managed BGP TCP-shared secret that can be supplied to a switch only by way of the cloud controller through a configuration update.

Purposeful protocol and port enablement

In remote deployments without robust physical security controls, access to network devices may be attainable by a potential "bad actor." With this circumstance in mind, most network protocols in Nexus Hyperfabric are only enabled by administrator configuration, ensuring that the services running on the platform in operation are limited to what is necessary for the deployment. This reduces the potential threat surfaces available to a malicious endpoint/user connected to the topology.

The following protocols are only enabled on ports with the fabric role once the peer is validated as a bound member of the fabric:

- **IPv6** - used for all switch-to-switch IP communications
- **BGP** - switching fabric control plane protocol
- **VXLAN** - dataplane encapsulation of endpoint communications
- **LLDP** - switch-to-switch and switch-to-endpoint discovery (fabric as well as other port roles)

The following protocols and traffic are enabled based on administrator-driven configurations:

- **DHCPv4 relay** - endpoint dynamic addressing
- **eBGP** - when an external peer is configured

Each switch ships with all interfaces configured as the role “unused,” and in this state the port will listen for a potential fabric peer to help automating interconnectivity; however, it will not forward any dataplane traffic. This default “closed” state helps avoid unnecessary availability on unused ports that could provide access to critical traffic in the fabric.

Multi-tenancy and secured management access

Keeping the network devices and cloud communications secured only prevents malicious attempts to compromise the devices and traffic in transit. Customer data isolation is another area of focus for the solution. Nexus Hyperfabric is designed inherently as a multi-tenant service, divided into organizations, fabrics, and users. Organizations are designed to represent a business entity, such as a company or business unit, and all data is securely partitioned and inaccessible to other organizations in the Nexus Hyperfabric cloud. Within an organization, fabrics represent a physical presence within the data center, composed of one or many switches operating as a logical EVPN fabric. Users are attached to one or many organizations (tenants) and assigned a role providing full access (administrator), read-write, or read-only, on an organization-by-organization basis. The Nexus Hyperfabric service maintains separation and isolation of data based on the organization and securely stores each customer’s data, protecting against any data access between organizations.

For information on handling Personally Identifiable Information (PII) and cloud-hosting details, please refer to the privacy data sheet at: Cisco Nexus Hyperfabric Controller.

Customer data storage

Nexus Hyperfabric as a management platform stores all necessary information for the design, deployment, and operation of data-center fabrics. All tenant information is stored in a segmented and secured model, thereby preventing access between organizations. The data storage is minimized to configuration and operational telemetry including:

- **Fabric blueprint configurations**
  - L2 and L3 logical networks and IP addressing
  - BGP peering
  - Static routing
  - DHCP relays
  - Switchport interfaces
  - Inventory
  - Historical changes
- **Fabric operational data**
  - Device status
  - Management connectivity details
    - IP and DNS
    - Proxy servers
    - Intermediary SSL Decrypt Proxy Certificate Chains

    - Port and traffic statistics
    - Environmental statistics
    - Endpoint LLDP and MAC/IP bindings
- **Administration**
    - User accounts and roles
        - Accounts are referenced only to the CCO ID service, and no credentials are stored in the Nexus Hyperfabric dashboard.
    - API tokens and scopes

Cloud-management administrative access controls

The management interface is another attack vector that the solution provides robust security tooling to protect and audit. Securing access to the Nexus Hyperfabric dashboard starts with Cisco's Connected Online ID (CCO ID) for seamless single sign-on across Cisco services. CCO single sign-on enhances security by incorporating tools such as Multifactor Authentication (MFA), ensuring that accounts and access are protected with robust MFA features. Instead of developing an in-house authentication service, user authentication is delegated to Cisco's CCO service to maintain secure login processes across Cisco® offerings. A user account can be associated to one or many organizations, and in each organization a separate role may be assigned providing full, or varying degrees of, access based on the user's assigned role in the organization.

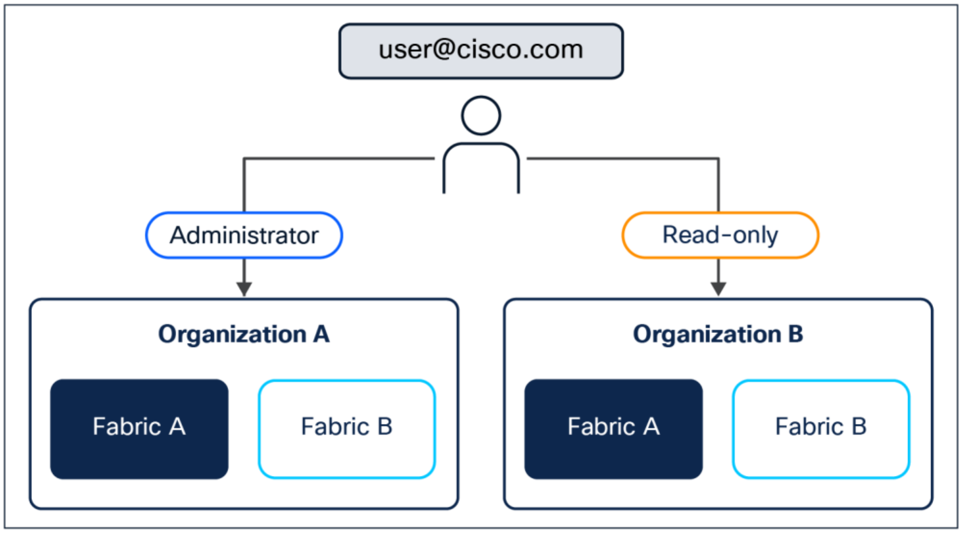

Figure 6.
Users can be attached to multiple organizations with separate roles.

For more information on CCO ID and privacy policy, see the [Cisco Online Privacy Statement].

Nexus Hyperfabric dashboard administrator roles

In a highly distributed deployment, expansion of visibility into the network's operation by local staff is typically necessary to avoid shipping an engineer on site for simple tasks such as port-state and connectivity troubleshooting. Nexus Hyperfabric administrative roles have been created that allow control of access to the management portal and telemetry. These roles encompass not only administration but also visibility through purpose-built tools such as the "On-Site" interface. The On-Site interface is designed to guide fabric deployment by partners and contractors providing the visibility to verify cabling and interface status but not modify configurations affecting the organization's deployments. The following roles are listed from lowest to highest privilege.

Read-only

A user with a read-only privilege is allowed to see fabrics, configurations, and any information regarding the deployment and operations; however, it is not permitted to make any modifications to the environment nor create any API tokens for API access.

Read-write

A user with a read-write privilege can not only view the environment but can make administrative changes to the configurations within the fabrics created in the organization. This includes creating, deleting, and modifying fabrics and their configurations. The read-write role may also create API tokens, including the read-write and read-only scope, depending on the use case.

Administrator

The administrator role has full access to the organization and can perform any operation within the environment. This includes all functions available in the read-write role, as well as the ability to add, delete, and disable user accounts attached to the organization. The administrator role may also create any form of API token, including an admin-scoped token for programmatic management of an organization.

Role-based API access

Role-based access to Nexus Hyperfabric APIs is managed through user-created bearer tokens (token creation is limited to administrator and read-write user roles), which are housed inside each organization with the scopes of read-only, read-write, and administrator. These scopes define the level of access a token has to the APIs, ensuring that operations are performed securely and within the required scope for the use case. The tokens themselves require a lifetime duration to be set, ensuring tokens are rotated in 1-, 6-, 12-, or 24-month intervals.

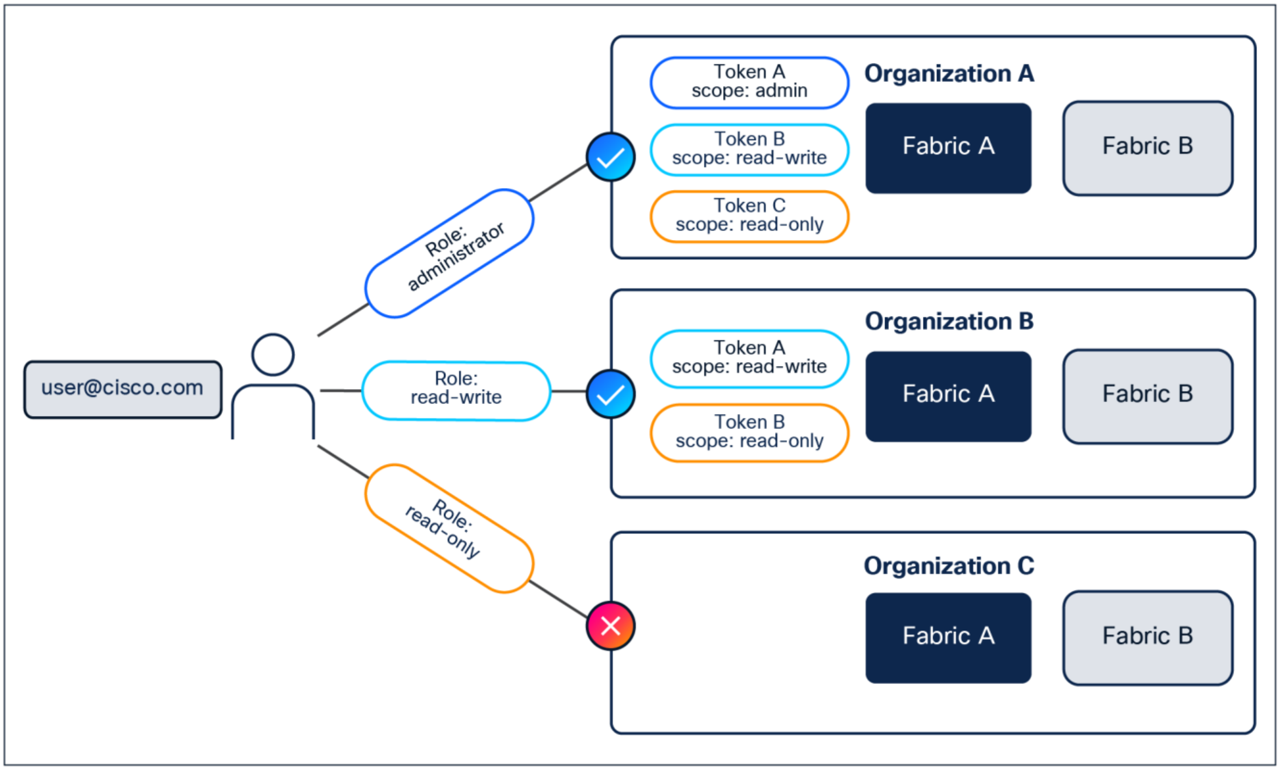

Figure 7.
A user account’s role impacts the ability to create tokens and their scope.

The read-only scope allows API-based view of the telemetry and configuration of fabrics within the organization. The read-write role permits modifying configurations that the read-only scope can only see. The administrator role provides full access, including the ability to manage user creation, deletion, disablement, and permission modifications. The scopes should match the same permissions provided to user account roles. Each scope is a super-set of the lesser scope in that read-write includes access to the API endpoints of read-only, and administrator includes the API endpoints of read-only and read-write as well as user management.

This Role-Based Access Control (RBAC) system is crucial for maintaining security and operational integrity within the Nexus Hyperfabric environment. The architecture ensures that API time-based tokens can be created based on specific use-cases and can only perform actions that are appropriate for their assigned roles, thereby enhancing security and efficiency in managing Nexus Hyperfabric API availability, programmatic management, and monitoring of data-center fabrics.

Local administrative threat surface hardening

Nexus Hyperfabric management is only available through the cloud; local administration is not available for any network forwarding configurations. Every decision made in the architecture assumes that unauthorized access to the network devices is expected. With that assumption in mind, administrative functions of the physical network devices are limited to a protected menu-driven serial port, enabling only troubleshooting and management-interface configuration changes necessary for providing or restoring cloud connectivity. This limited interface capability is constructed so that Nexus Hyperfabric switches can be deployed and managed not only in customer-owned data centers, but especially in unsecured remote locations.

Serial console administration

As outlined above, every Nexus Hyperfabric switch is equipped with a serial console port. This interface does not allow for any administration of the fabric by design, safeguarding the fabric from either unintentional or malicious attempts at modifying the fabric functions.

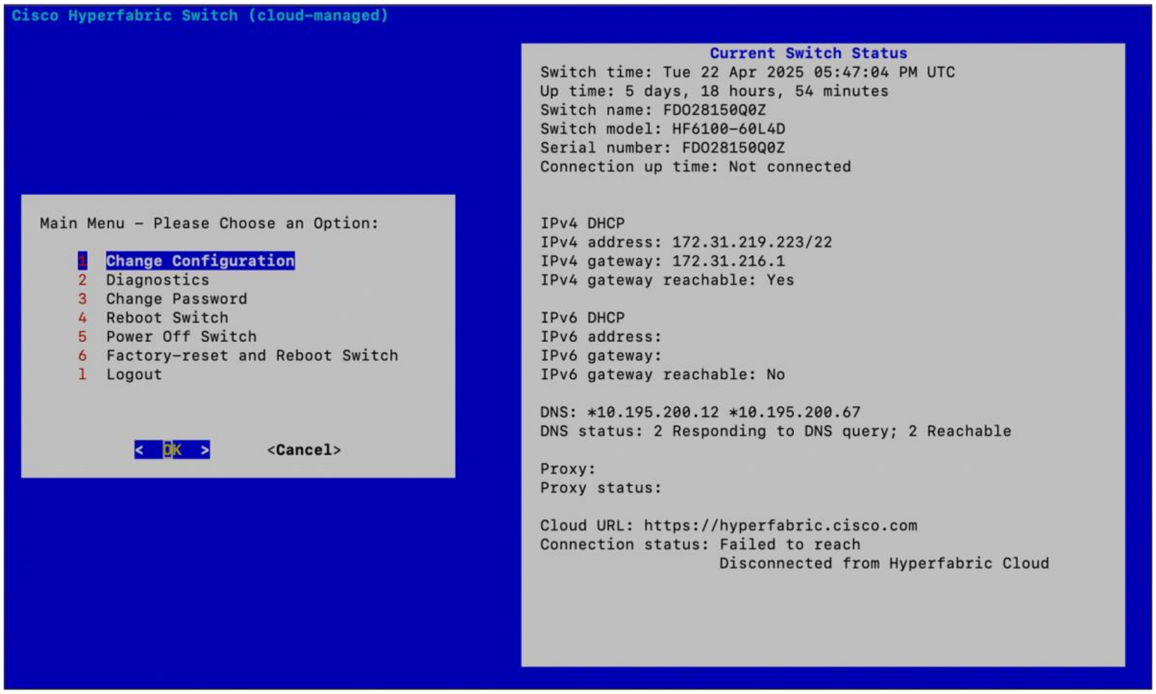

```
Cisco Hyperfabric Switch (cloud-managed)

Main Menu - Please Choose an Option:

    1  Change Configuration
    2  Diagnostics
    3  Change Password
    4  Reboot Switch
    5  Power Off Switch
    6  Factory-reset and Reboot Switch
    1  Logout

          < OK >       <Cancel>
```

```
                  Current Switch Status
Switch time: Tue 22 Apr 2025 05:47:04 PM UTC
Up time: 5 days, 18 hours, 54 minutes
Switch name: FDO28150Q0Z
Switch model: HF6100-60L4D
Serial number: FDO28150Q0Z
Connection up time: Not connected


IPv4 DHCP
IPv4 address: 172.31.219.223/22
IPv4 gateway: 172.31.216.1
IPv4 gateway reachable: Yes

IPv6 DHCP
IPv6 address:
IPv6 gateway:
IPv6 gateway reachable: No

DNS: *10.195.200.12 *10.195.200.67
DNS status: 2 Responding to DNS query; 2 Reachable

Proxy:
Proxy status:

Cloud URL: https://hyperfabric.cisco.com
Connection status: Failed to reach
                   Disconnected from Hyperfabric Cloud
```

Figure 8.
Local serial port menu.

Serial port change configuration scope

Under the serial console menu, a limited set of configurations can be modified including:

- Management IPv4/v6 address assignment (static / DHCP)
- Default gateway
- DNS servers
- Corporate proxy configurations
    - Proxy URL and port
    - Username
    - Password

These configurations are solely for the purpose of providing cloud connectivity and have no effect on the current operating state of the fabric control or dataplane.

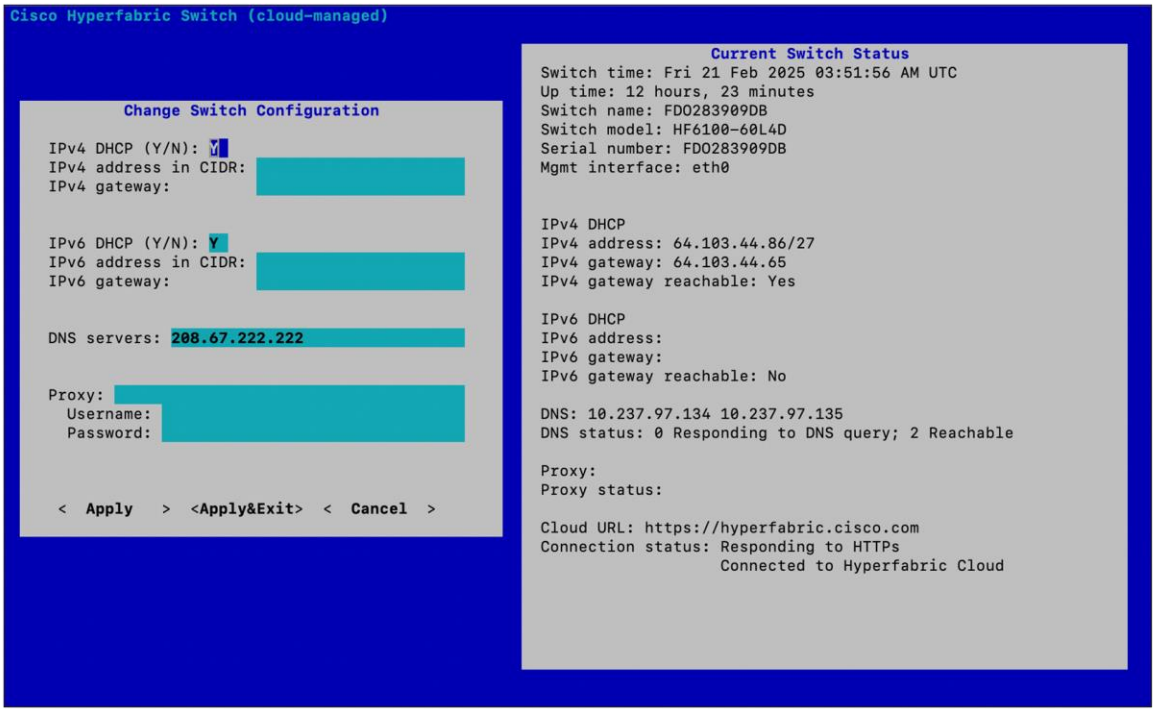

```
Cisco Hyperfabric Switch (cloud-managed)

        Change Switch Configuration

   IPv4 DHCP (Y/N): Y
   IPv4 address in CIDR:
   IPv4 gateway:

   IPv6 DHCP (Y/N): Y
   IPv6 address in CIDR:
   IPv6 gateway:

   DNS servers: 208.67.222.222

   Proxy:
     Username:
     Password:

   <  Apply    >  <Apply&Exit>  <  Cancel  >
```

```
                    Current Switch Status
Switch time: Fri 21 Feb 2025 03:51:56 AM UTC
Up time: 12 hours, 23 minutes
Switch name: FDO283909DB
Switch model: HF6100-60L4D
Serial number: FDO283909DB
Mgmt interface: eth0


IPv4 DHCP
IPv4 address: 64.103.44.86/27
IPv4 gateway: 64.103.44.65
IPv4 gateway reachable: Yes

IPv6 DHCP
IPv6 address:
IPv6 gateway:
IPv6 gateway reachable: No

DNS: 10.237.97.134 10.237.97.135
DNS status: 0 Responding to DNS query; 2 Reachable

Proxy:
Proxy status:

Cloud URL: https://hyperfabric.cisco.com
Connection status: Responding to HTTPs
                   Connected to Hyperfabric Cloud
```

Figure 9.
Serial management configuration settings.

Serial port diagnostics menu

Since the service is dependent on the cloud for visibility and configuration updates, the console ports' diagnostic tools are provided for local troubleshooting of connectivity to the cloud. This menu provides a common set of useful tools, including tests providing:

- Gateway, DNS, proxy, and cloud reachability
- Neighbor switch, LLDP, ACL, ARP, and route tables
- Status information for the management interface and Linux services
- Local device metadata (device inventory)
- Management agent details

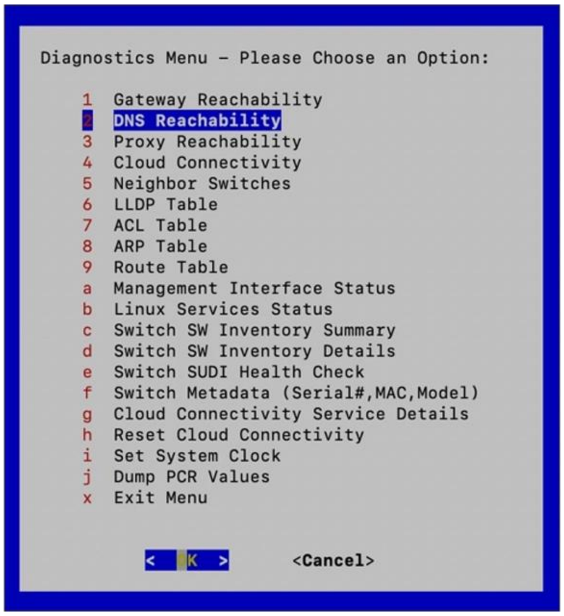

```
Diagnostics Menu - Please Choose an Option:

     1  Gateway Reachability
     2  DNS Reachability
     3  Proxy Reachability
     4  Cloud Connectivity
     5  Neighbor Switches
     6  LLDP Table
     7  ACL Table
     8  ARP Table
     9  Route Table
     a  Management Interface Status
     b  Linux Services Status
     c  Switch SW Inventory Summary
     d  Switch SW Inventory Details
     e  Switch SUDI Health Check
     f  Switch Metadata (Serial#,MAC,Model)
     g  Cloud Connectivity Service Details
     h  Reset Cloud Connectivity
     i  Set System Clock
     j  Dump PCR Values
     x  Exit Menu


           <  OK  >        <Cancel>
```

Figure 10.
Serial port diagnostics menu.

Secure telemetry

In the Nexus Hyperfabric service architecture, customer data is segmented within the cloud based on organizational tenancy, and all communications are protected through an encrypted channel. For Cisco to also protect endpoint data in the fabric, the dataplane’s payloads are not collected by the cloud (unless a customer administrator leverages tools that capture payload data). The telemetry that is captured by default is purely focused on gathering statistics necessary to ensure optimal operation and validation of the health of the fabric. Telemetry is transmitted over the management session through purpose-built payloads, designed for efficient usage of bandwidth.

Examples of telemetry collected:

- Port state, counters, and traffic statistics
- Device health metrics
- Environmental status including:
  - Fan speed and state
  - Power status
  - Component temperature
- Digital optical monitoring statistics
- LLDP peer details
- Network information including:
  - IP/MAC addresses
  - External routes
  - BGP peers

Data-loss prevention and man-in-the-middle decryption

Cloud service attestation is performed by every Nexus Hyperfabric switch that is connected to the cloud, ensuring that the cloud service and certificates presented are not modified. In the event there are requirements to support inline decryption of management traffic, Nexus Hyperfabric switches automatically detect man-in-the-middle proxies and report this information to the cloud-management controller. The dashboard then presents a menu to the administrator, providing details to accept or deny trust of the certificates presented. During this time, the inner channel is encrypted to protect communications (referred to as TLS-in-TLS encryption). If the new certificate chain is not accepted, the secondary encryption channel is maintained, and a red assertion is raised.

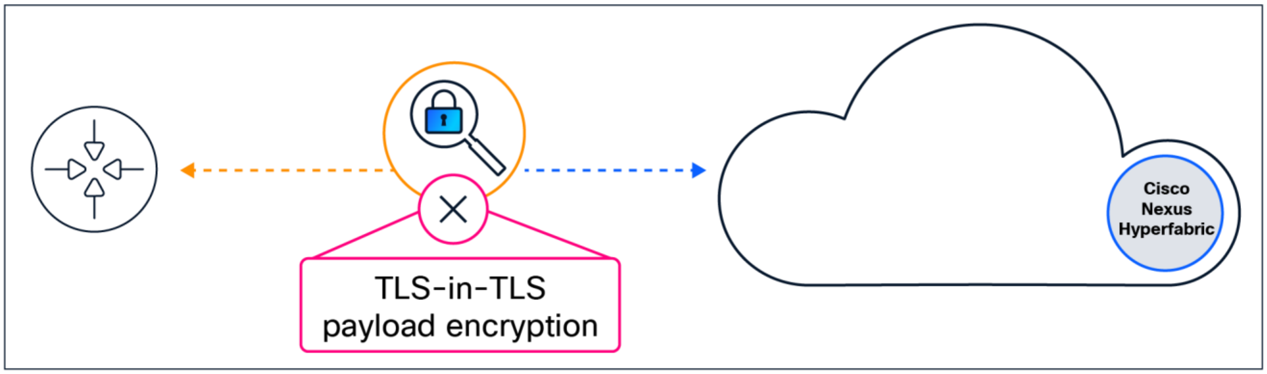

Figure 11.
TLS-in-TLS encryption when the intermediary certificate is not accepted.

If the certificate is accepted, the platform will no longer alert and will tear down the TLS-in-TLS encryption, allowing for situations in which inspecting management traffic to the cloud is included as part of preventing the loss of corporate data.

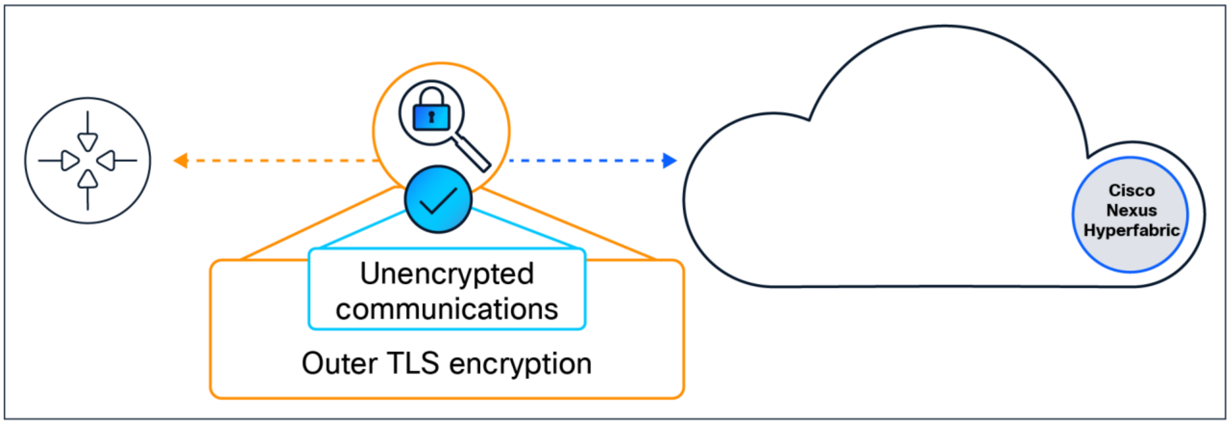

Figure 12.
Management-traffic decryption when the intermediary certificate is accepted.

Day-2 cloud-managed operations

The following sections outline how switch software and services are maintained, and how device software threat surfaces are handled to continuously maintain secure operation of distributed cloud-managed fabrics.

Cloud-accelerated vulnerability mitigation

From an operational security perspective, one of the most valuable aspects of cloud- management is that all devices, services, and states are known and constantly monitored. If there is a security vulnerability in the cloud service, the customer is not expected no required to patch the system; this is instead handled by the provider (Cisco). This is a common aspect of any (I/S)aaS (infrastructure/software as a service) offering; it shifts software maintenance of one of the most important aspects of day-2 operations away from a painful customer-driven exercise and places the responsibility on dedicated Cisco engineering teams. Cisco's Advanced Security Initiatives Group, as well as other teams within Cisco, provide a combined suite of vulnerability assessment and mitigation services, ensuring that the Nexus Hyperfabric cloud is subject to stringent and continuous review.

Software image management

Cisco controls the entire lifecycle of software development and signing following Cisco's Secure Development Lifecycle (CSDL) on the Trust Center webpage. Nexus Hyperfabric customers are provided with tooling to orchestrate when and where an update is scheduled for deployment and which release is utilized; however, no third-party or custom software images are permitted for installation on any Nexus Hyperfabric switch platform. In the event an attempt is made to run a software image that has not been provided by Cisco, due to the hardware-anchored secure-boot processes the switch will fail to boot as a safety measure to protect against compromise.

Every switch maintains two partitions in storage. The primary partition is utilized for booting the current image, and the secondary partition is utilized to store the previously installed image. This design is in place to ensure that automated recovery can be performed in the event of an upgrade issue, reducing the potential need for end-customer administrative intervention.

On-switch software updates and vulnerability mitigation

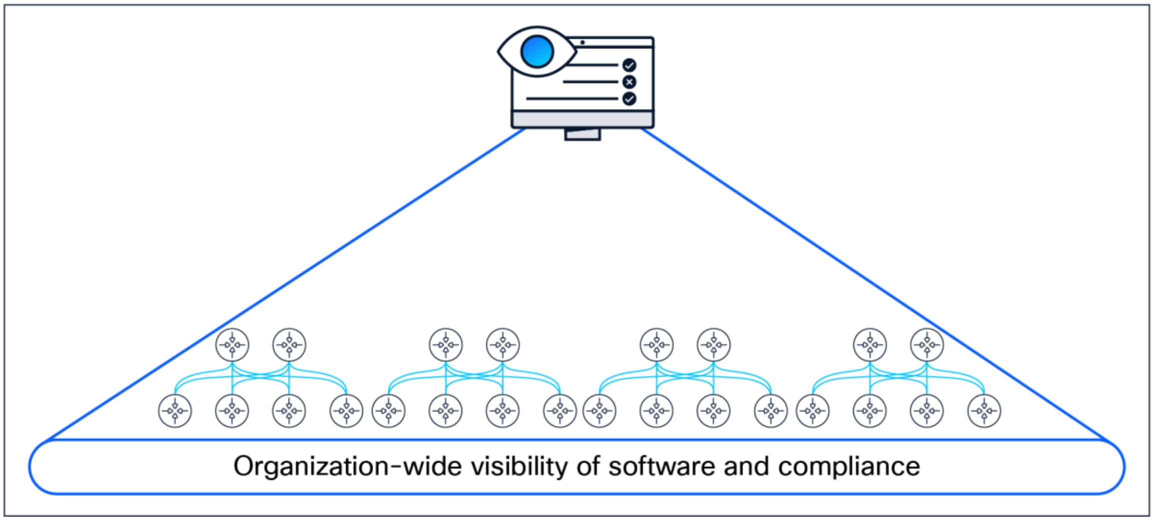

Figure 13.
Administrators have visibility across the organization's deployment.

Cisco's Nexus Hyperfabric solution not only proactively ensures that access to cloud services is secured, but that administrators are provided tools to keep the switches and their software up to date. These updates not only include new features, but also any patches for vulnerabilities. If vulnerability mitigation requires an OS update, the software is made available without delay. In the event an OS update is required for the switch, customers are provided with orchestration tools to upgrade one or many fabrics without the need to host the file locally or have an engineer execute switch-by-switch OS upgrades, as most solutions have historically, in some manner, required.

Nexus Hyperfabric's switch software updates can be broken down into two components:

- The management agent
- The switch-operating system (switch OS)

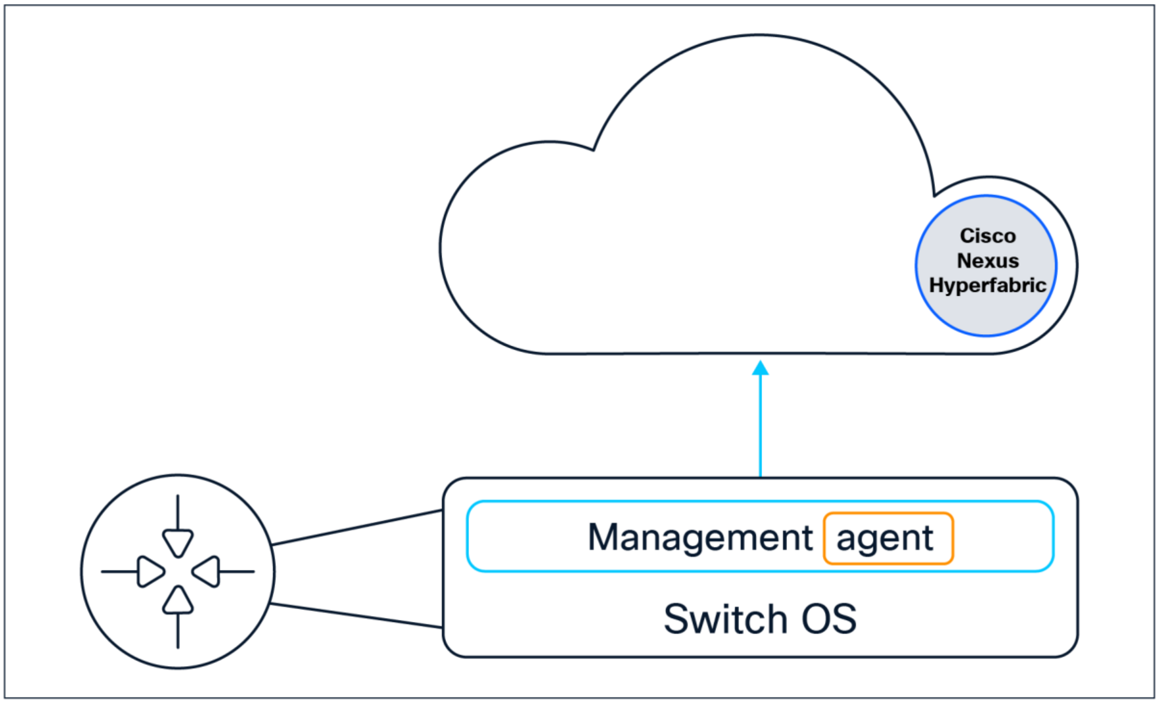

Figure 14.
The management agent is independent from the switch OS.

Management agent

The management agent is routinely updated transparently to the customer as a service update, keeping the communications and management through the cloud as secure as possible. This agent upgrade includes updates to the cloud-controller certificates and any other nonservice-affecting patches to the platform in operation. During the agent upgrade, no fabric-forwarding functions are affected, and at most a minimal interruption in communication to the cloud may be witnessed if the "Cloud Connectivity" table in the Nexus Hyperfabric dashboard is being actively monitored.

Switch-operating system

Switch-operating system updates are orchestrated through the Nexus Hyperfabric dashboard, providing a simple and intuitive process of maintaining the target OS version deployed on a fabric-by-fabric basis. Role-based controls limit which users can initiate an upgrade, and all upgrade activities are logged in the web console. Cisco consistently provides updates, typically incorporating security patches, bug fixes, and new features (typically in major release upgrades). If a vulnerability is found in the platform, proactive alerts are sent by the Cisco Product Security Incident Response Team (PSIRT), outlining the issue and affected firmware and/or feature. In this scenario the Nexus Hyperfabric team will have published the patched OS update, prior to announcement, enabling rapid remediation by the end customer.

For more information on Cisco's Product Security Incident Response Team, follow this link: Cisco PSIRT Infographic.

Switch disaster recovery

For disaster-recovery scenarios, in the event a switch is unable to boot, a procedure has been published for loading a signed image onto a USB flash drive, ensuring the switch can be recovered for normal operation. This would only come into account in the event the switch has been tampered with, as the dual-partition design in the platform ensures that image upgrades can be automatically reverted if there are problems with the primary image post-upgrade. This function is in development and will be made available at a later date.

Security practices in Cisco's Nexus Hyperfabric team

The Nexus Hyperfabric team takes a systematic approach to data protection, privacy, and security. The team believes that robustly secure access control and development practices require active involvement across the development and site-reliability engineering operations teams (SREs) to ensure that customer data, production services, and access control to critical resources are managed through a zero-trust methodology.

There are three separate groups in the Nexus Hyperfabric team that enforce separation of duties: the software-development, QA, and operations (SRE) groups. The collective team participates in a continuous delivery pipeline, and each step in the pipeline includes procedural controls enforced by each group. As code is developed by the development team, it is passed to the QA team to validate and only passed on to the operations team for production deployment when the quality metrics required are met.

Critical infrastructure and development controls

The services providing the basis of the Nexus Hyperfabric controller and management infrastructure are only accessible to the SRE team, through bastion hosts that are continuously logged and audited on access and usage. Developers are not permitted access to production services, and SRE access controls are operated in a least-privilege, zero-trust model. This model only permits access to the necessary aspects of the engineer's responsibility. All development in both cloud infrastructure and switching software follows Cisco's Secure Development Lifecycle (CSDL) on the Trust Center webpage. As part of the CSDL process, Cisco's Security and Trust Organization (STO) provides oversight to ensure that proper procedures are followed in every aspect of the solution's lifecycle. This organization encompasses various groups that focus on delivering secure products and responding to security concerns as they arise. The STO works closely with Cisco IT and Information Security (InfoSec) to ensure that the products built and the infrastructure operated are secure, supporting business productivity while protecting systems and data from internal and external threats. Additionally, the STO is involved in embedding security throughout Cisco's infrastructure and enforcing security-focused policies. STO operates Cisco's Trust Center, which is accessible through the following link: Cisco Trust Center.

This procedural approach ensures that no developer can commit code directly to the production environment and that the code undergoes a thorough vetting process before making its way to customers. While the Nexus Hyperfabric solution delivers rapid updates for critical fixes and new features, the organization will never circumvent the separation of duties and testing requirements in place to deliver a best-in-class user experience along with the code quality necessary to trust critical deployment use of the platform.

The Nexus Hyperfabric solution also utilizes another Cisco service organization, referred to as Cisco's Advanced Security Initiatives Group (ASIG) (read more about this on the Trustworthy Systems: A Peek Behind the Curtain blog), to ensure that all platforms and services are routinely and consistently penetration-tested. ASIG is tightly coupled with the STO team to be proactive in resolving potential threats.

Customer-data access control

Access to customer data is highly controlled and protected; accessibility to this data is limited and is based on appropriate business needs and the functional roles of the personnel accessing it. Multiple security protocols and procedures are leveraged to ensure that personnel are strictly vetted prior to being provided access to systems in which customer data is processed, and they are given access only to their role's entitlement. This includes rigid policies and procedures incorporating role-based multifactor authentication, audit trails, and historical log retention policies. Cisco provides a master data agreement that can reviewed here for more detail: Cisco Master Data Protection Agreement (MDPA).

Quality assurance

The Nexus Hyperfabric team conducts system and unit testing on all code that is intended for production. The QA team tests and validates all code, from end to end, in a multistage process, validating each stage, before the code makes it to the production environment. This results in end-to-end validation of the user experience and assurance that the behavior of the platform and switches meets the requirements and highest expectations set for them.

Cloud-service resiliency and threat response

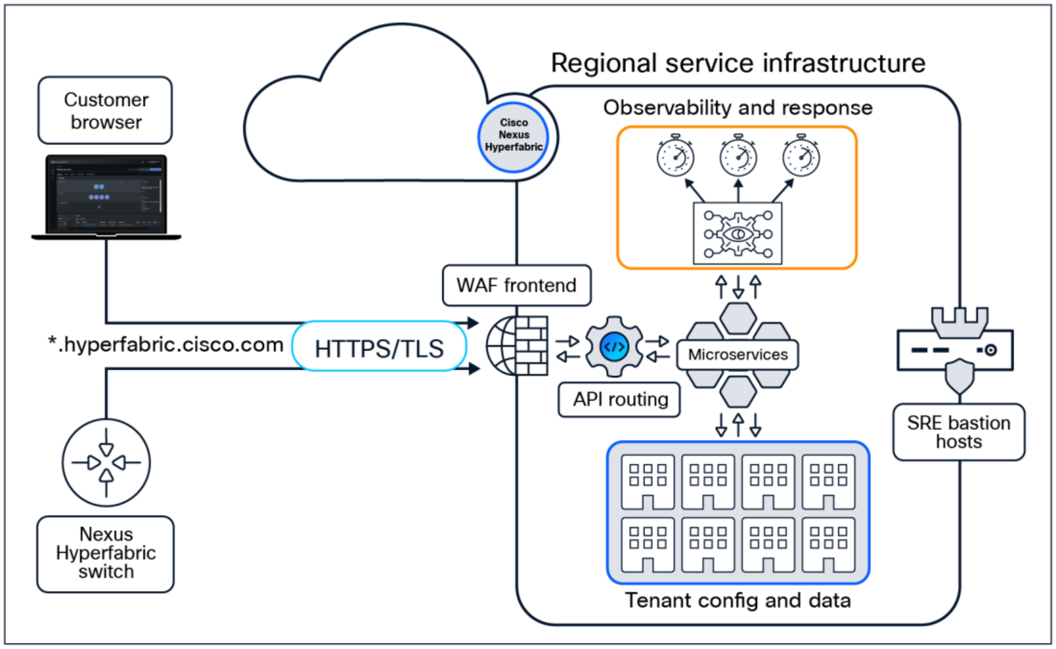

Figure 15.
Regional service infrastructure high-level overview.

Nexus Hyperfabric is hosted as a publicly available service providing access to anyone with internet connectivity. The cloud-service architecture was designed to maintain resiliency and availability with the expectation that potential attacks could be performed against the platform's publicly accessible endpoints. The Nexus Hyperfabric cloud deploys multiple forms of protections against potential attacks such as Distributed Denial of Service (DDoS), leveraging combinations of regional service routing (content delivery network), and front-end service firewalling (web application firewall). In the event an attack is executed against the services, automated threat responses are orchestrated to minimize the effect to the service's availability and customer accessibility to the platforms. All service availability and security-related events are continuously tracked and monitored by the SRE team, and procedures are in place to proactively mitigate any threatening activity. In the unlikely event cloud services are affected, due to the fabric's operational state-machine, no local traffic and routing on site will be affected, and full fabric operation is maintained.

External vulnerability reporting

Cisco provides public-facing bug-reporting tools to enable vulnerability finders, researchers, and ethical hackers to report any discovered vulnerabilities, and in some cases provides a bounty system to encourage proactive reporting. Every discovered security vulnerability is published by Cisco through the Cisco Security Advisories portal.

Where to go next?

For more information on Cisco Nexus Hyperfabric and product details, follow this link to learn more: Cisco's Nexus Hyperfabric.

If you would like to try it out yourself, follow this link to create an organization and design your first fabric: Nexus Hyperfabric dashboard. (requires a CCO account)

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**.

Printed in USA C11-5210754-00 08/25