White Paper

Public

The bridge to possible

Cisco Nexus OS Software Security Primer

Contents

Introduction

This white paper provides recommendations and guidance in configuring Nexus switches to ensure that it meets the highest security standards and is resilient against attacks.

There are three major areas on which administrators need to focus to harden network infrastructure: the management plane, the control plane, and the data plane. All three areas are equally important because they can all be compromised if not correctly hardened, which enables attacks to any of those areas to cause critical damage that, while different in nature, can have a considerable impact.

This white paper addresses typical questions raised about hardening these three areas, the recommended configuration that customers need to perform to harden Nexus 9000 switches, and the features that can and should be used in the given scenarios.

Principles of Secure Operations

Although a major part of this document is devoted to the secure configuration of Nexus 9000 switches, configurations alone do not completely secure a network. The operating procedures in use on the network contribute as much to security as the configuration of the underlying devices.

This section has some operational recommendations that you are advised to implement, as it contributes to maintaining your network securely and minimizing your attack surface and exposure. This section focuses on critical areas of network operations and may not be comprehensive.

Monitor Cisco Security Advisories and Responses

The Cisco Product Security Incident Response Team (PSIRT) creates and maintains publications, commonly referred to as Cisco PSIRT Security Advisories, for security-related concerns in Cisco products.

Cisco security vulnerability publications are available at http://www.cisco.com/go/psirt

.

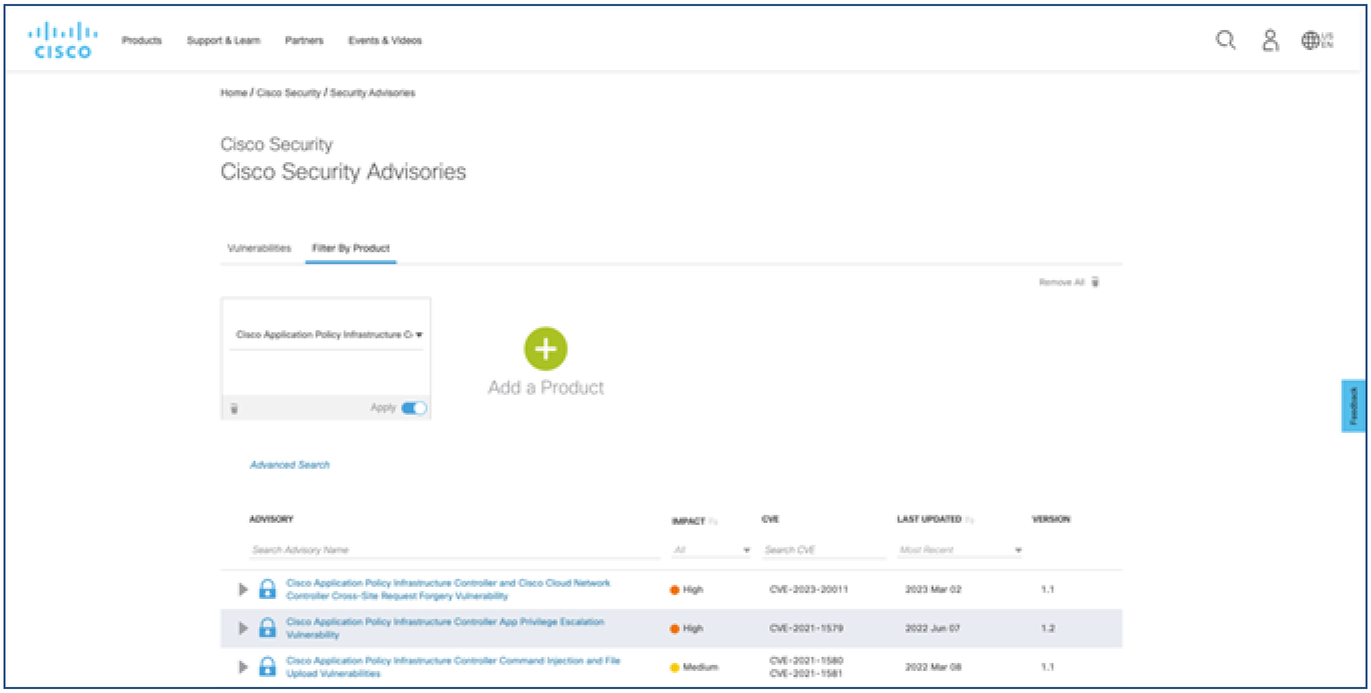

ADVISORY	IMPACT	CVE	LAST UPDATED	VERSION
Cisco Application Policy Infrastructure Controller and Cisco Cloud Network Controller Cross-Site Request Forgery Vulnerability	High	CVE-2023-20011	2023 Mar 02	1.1
Cisco Application Policy Infrastructure Controller App Privilege Escalation Vulnerability	High	CVE-2021-1579	2022 Jun 07	1.2
Cisco Application Policy Infrastructure Controller Command Injection and File Upload Vulnerabilities	Medium	CVE-2021-1580 CVE-2021-1581	2022 Mar 08	1.1

Figure 1 Cisco Security Advisories

Cisco releases bundles of Cisco Security Advisories at 16:00 Greenwich Mean Time (GMT) on a regular schedule twice each year. The specific release dates and schedules are different for every Cisco product.

Specifically, for Cisco ACI and NX-OS, bundles are released every fourth Wednesday of February and August at 16:00 GMT.

Cisco reserves the right to publish individual Security Advisories outside the above schedule.

To maintain a secure network, you must be aware of the Cisco security advisories and responses that have been released. To ease this, Cisco provides several ways to stay connected and receive the latest security vulnerability information from Cisco:

Cisco.com

The Cisco Security portal on Cisco.com provides Cisco security vulnerability documents and Cisco security information, including relevant security products and services.

Email

Cisco Security Advisories provide information about Critical-, High-, and Medium-severity security vulnerabilities. They are clear-signed with the Cisco PSIRT PGP public key and distributed to the external **cust-security-announce@cisco.com** mailing list that you can subscribe to.

To subscribe to the **cust-security-announce** mailing list, email cust-security-announce-join@cisco.com (the content of the message does not matter). You will receive confirmation, instructions, and a list policy statement.

RSS Feeds

Cisco security vulnerability information is also available through RSS feeds from Cisco.com. For information on how to subscribe to the RSS feeds, visit the Cisco Security RSS Feeds page.

Cisco PSIRT openVuln API

The Cisco PSIRT openVuln application programming interface (API) is a RESTful API that allows customers to obtain Cisco security vulnerability information in different machine-consumable formats. To learn about accessing and using the API, visit the PSIRT page on the Cisco DevNet website.

My Notifications

The My Notifications website allows registered Cisco.com users to subscribe to and receive important Cisco product and technology information, including Cisco Security Advisories.
To learn about Cisco security vulnerability disclosure policies and publications, see the Security Vulnerability Policy.

Cisco Nexus Dashboard Insights

Cisco Nexus Dashboard Insights is the Day-2 operations solution from Cisco that helps you manage, operate, and troubleshoot your data center networks by offering you insights, visibility, and analytics.
Among the variety of use cases and features that Nexus Dashboard Insights provides, there is the capability to be proactively notified about Security Advisories and Field Notices impacting your data center fabric.
In contrast with the options listed above, Nexus Dashboard Insights will only inform you about the Security Advisories that are impacting the environment based on the software release, hardware models, and features being used. For Security Advisories impacting the system, a detailed description of recommended actions to be taken is described. Hence, Nexus Dashboard Insights makes monitoring Security Advisories and acting on them easier.

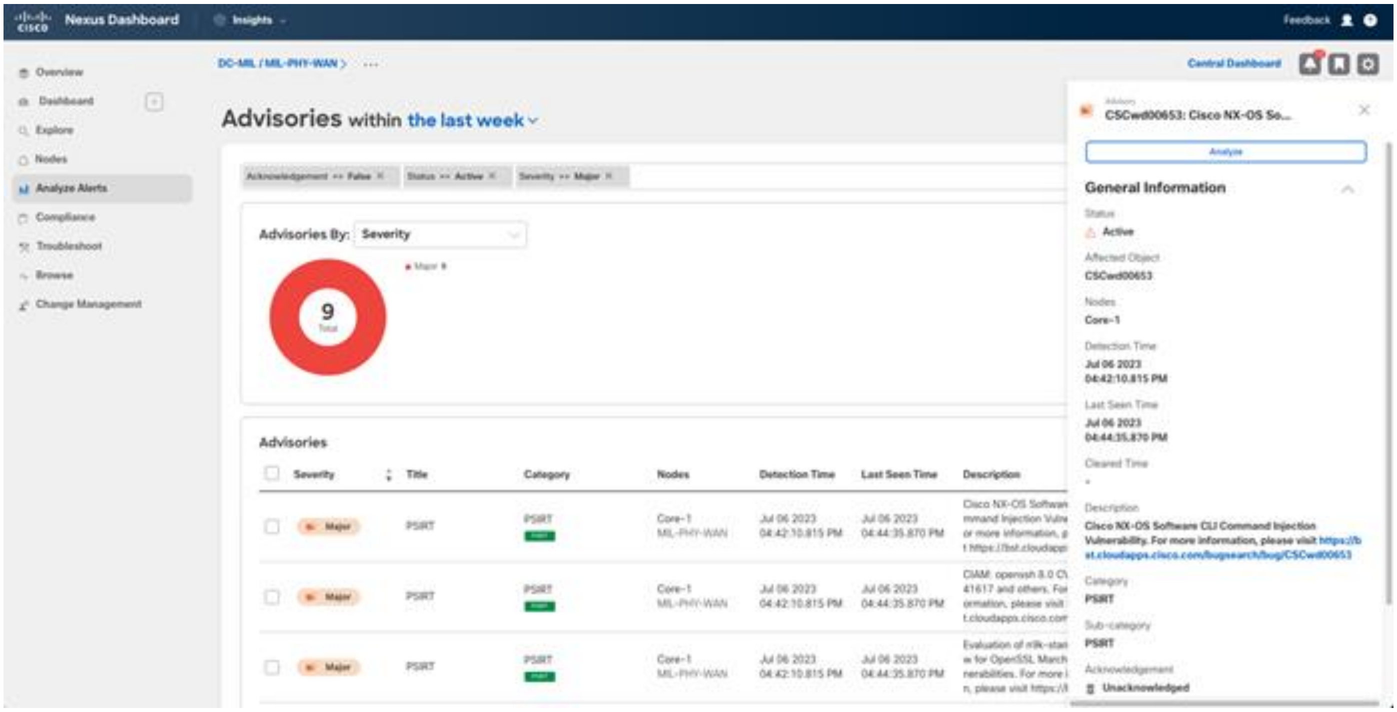

Advisories

Severity	Title	Category	Nodes	Detection Time	Last Seen Time	Description
Major	PSIRT	PSIRT	Core-1 MIL-PHY-WAN	Jul 06 2023 04:42:10.815 PM	Jul 06 2023 04:44:35.870 PM	Cisco NX-OS Softwar mmand Injection Vulr or more information, p t https://bst.cloudapp
Major	PSIRT	PSIRT	Core-1 MIL-PHY-WAN	Jul 06 2023 04:42:10.815 PM	Jul 06 2023 04:44:35.870 PM	CIAM: openssh 8.0 CV 41617 and others. For ormation, please visit t.cloudapps.cisco.com
Major	PSIRT	PSIRT	Core-1 MIL-PHY-WAN	Jul 06 2023 04:42:10.815 PM	Jul 06 2023 04:44:35.870 PM	Evaluation of n9k-stan w for OpenSSL March nerabilities. For more i n, please visit https://l

Advisory

CSCwd00653: Cisco NX-OS So...

General Information

Status: Active

Affected Object: CSCwd00653

Nodes: Core-1

Detection Time: Jul 06 2023 04:42:10.815 PM

Last Seen Time: Jul 06 2023 04:44:35.870 PM

Cleared Time: -

Description: Cisco NX-OS Software CLI Command Injection Vulnerability. For more information, please visit https://bst.cloudapps.cisco.com/bugsearch/bug/CSCwd00653

Category: PSIRT

Sub-category: PSIRT

Acknowledgement: Unacknowledged

Figure 2 Cisco Nexus Dashboard Insights Advisories

Use Authentication, Authorization and Accounting

Authentication, Authorization, and Accounting (AAA) is a well-known security framework used to control access to resources and services in systems and networks.
Authentication refers to the process of verifying the identity of a user or device attempting to access a resource. This is typically done through a username and password.
Authorization refers to the process of determining whether a user or device has the necessary permissions to access a specific resource or perform a particular action within the network.
Accounting refers to the process of logging the activities of users and devices within the network, including the resources they access and the actions they perform. This information can be used for auditing or other purposes.
These three components together form a comprehensive security framework that helps ensure that only authorized users and devices can access network resources, and that their activities can be tracked for security and accountability purposes.

Cisco Nexus OS supports the AAA framework for both local users and remote users.

For more information, see the Cisco NXOS Security Configuration Guide.

Implement the Principle of Least Privilege

The principle of the least privilege is a security concept that suggests that a subject should be granted the minimum level of access necessary to perform their job or function. This means that a subject should only be given the permissions and privileges needed to complete their specific tasks, and no more. In this context, a subject can be either a user, an automation engine (such as Terraform or Ansible), a process, or another system or product.

The principle of the least privilege is important because it helps reduce the risk of unauthorized access or accidental misuse of sensitive data or resources. By limiting access to only what is necessary, the attack surface is reduced and the potential damage that can be caused by a security breach or vulnerability is minimized.

Cisco Nexus OS offers a very strong Role-Based Access Control (RBAC) feature set to help administrators effectively implement this Principle of Least Privilege. For more information, see the Access, Authentication, and Accounting section in the Cisco NX-OS Security Configuration Guide.

Use Centralized Log Collection and Monitoring

Indication removal techniques (IRT) are used to remove traces or indications of a security incident or intrusion to prevent detection or impede further investigation. IRTs can include actions such as deleting or modifying log files, altering system settings or configurations, or manipulating network traffic. From a security perspective, we strongly recommend that you save logs remotely, as they are a critical source of evidence in detecting and investigating security incidents. Logs can provide valuable information such as the date and time of an event, the source and destination of network traffic, and the actions taken by users or processes on a system.

By saving logs remotely, organizations can ensure that they are protected from IRTs and other attempts to modify or delete logs on local systems. Remote logging also allows logs to be stored securely in a centralized location where they can be monitored, analyzed, and correlated more easily.

Cisco Nexus OS provides different mechanisms for exporting events and logs to remote locations. These mechanisms include:

Syslog

Syslog is a standard protocol used for logging system messages and events on network devices and servers. Cisco Nexus OS supports exporting events, faults, and audit logs by using Syslog over UDP, TCL, or SSL.

SNMP

Simple Network Management Protocol (SNMP) is a standard protocol used for network management and monitoring of devices, such as routers, switches, and servers.

SNMP can be used for log collection by allowing network administrators to retrieve logs and other management information from SNMP-enabled devices such as routers, switches, and servers, and send them to a centralized log management system for further analysis and monitoring. SNMP-enabled devices can also send notification messages, called SNMP traps, to the centralized monitoring system when specific events or conditions occur.

Cisco Nexus OS supports both SNMPv2c and SNMPv3. We strongly recommend that you use SNMPv3 from a security perspective as it offers both authentication and encryption.

For more information, see the Cisco Nexus OS Faults, Events, and System Messages Management Guide.

Use Only Secure Protocols

Many of the protocols used for management access and operation purposes carry sensitive network management data that must be properly protected. Therefore, secure protocols should be used for those connections whenever possible. Table 1 lists examples of secure protocols and their insecure counterparts.

Table 1. Secure Protocols and Insecure Protocols

Secure Protocol	Insecure Protocol
HTTPS	HTTP

Secure Protocol	Insecure Protocol
SCP / SFTP	FTP
SNMPv3	SNMPv2
SSH	Telnet

These secure protocols and other management and control-plane protocols that offer integrity, confidentiality, and authenticity, rely on ciphers and crypto algorithms. Besides using protocols that offer those security capabilities, it is also important to use crypto algorithms and ciphers that are considered secure and that have no known weaknesses.

Therefore, those algorithms that are no longer considered secure should be avoided if there is a better alternative available. Examples of those insecure algorithms and ciphers are MD5, SHA1, and TLSv1.0/1.1.

In Nexus OS it is advisable to disable insecure protocols. For example, TLSv1.0 and v1.1 are not recommended to use with TLSv1.2 as the default option (TLSv1.3 is also supported). For more information, see the Cisco Nexus OS Security Configuration Guide.

Perform Configuration Management

Configuration management is the process of managing and controlling the changes made to a system or software throughout its development and operational life cycle. It determines how changes are proposed, reviewed, approved, and deployed.

Within the context of security and hardening, the most relevant part of the Configuration Management is to ensure that configuration backups are periodically collected and archived safely in case they must restore the system. Engineers and administrators can use configuration archives to roll back changes that have been made to network devices and restore the system after a disaster or an incident.

In the context of security, configuration archives can also be used to determine what security changes were made, and when these changes occurred. In conjunction with audit log data, this information can assist in the security auditing of network devices.

Use Strong Passwords

Strong passwords are critical security measures for protecting networks and systems against unauthorized access to sensitive data and systems, and cyber-attacks. Using strong passwords that include a combination of upper- and lower-case letters, numbers, and special characters make it much harder for malicious actors to brute-force the password, even if they are using sophisticated tools.

Additionally, strong passwords increase accountability by making it more difficult for individuals to deny that they accessed a system or network device, given that both their password and their MFA token or device must have been used to do so. Cisco N9K switches allow administrators to enforce, for local users, the use of strong passwords.

Securing the Management Plane

The management plane consists of functions that achieve the management goals of the network. These goals include interactive management sessions using SSH, in addition to statistics gathering with tools and protocols such as SNMP or NetFlow. When considering the security of a network device, you must make sure that the management plane is protected. If a security incident undermines the functions of the management plane, recovering or stabilizing the network will be a challenge.

Managing Passwords

Passwords are a primary mechanism for controlling access to resources and devices. Password protection is accomplished by defining a password or secret that is used to authenticate requests. When a request is received for access to a resource or device, the request is challenged for verification (usually in the form of a request for a password and username). Access then can be granted, denied, or limited based on the authentication result. As a security best practice, passwords should be managed with a TACACS+ or RADIUS authentication server. However, note that a locally configured username and password for privileged access is still needed in the event of a TACACS+ or RADIUS service failure.

In Cisco NX-OS, there is no concept of an enable-secret or enable-password setting like which you see in Cisco IOS devices. Privileges are managed using role-based access control (RBAC). The functional equivalent to enabling mode access on a Cisco IOS Software device is the assignment of an account to the network-admin (global privileged access) or vdc-admin (virtual device context [VDC]-specific privileged access) role in Cisco NX-OS.

In addition, unlike Cisco IOS Software, Cisco NX-OS does not locally store a single enable-secret cross-user shared credential as an individual password item in the configuration. Each user account maintains its own password (stored locally or through AAA), and authorization levels are dictated by the role assigned to a given account. Therefore, you must consider protecting all the passwords used for all accounts assigned privileged access with the network-admin or vdc-admin roles. This task is greatly simplified if password management is centralized using AAA services.

By default, Cisco NX-OS protects all passwords used in the system configuration using irreversible MD5 hashing. There is no option to modify this behavior.

Enforcing Strong Password Selection

Cisco NX-OS has the built-in capability to optionally enforce strong password checking when a password is set or entered. This feature is enabled by default and will prevent the selection of a trivial or weak password by requiring the password to match the following criteria:

- Is at least eight characters long
- Does not contain many consecutive characters (abcde, lmnopq, and so on)
- Does not contain dictionary words (English dictionary)
- Does not contain many repeating characters (aaabbb, tttttyyyy, and so on)
- Does not contain common proper names (John, Mary, Joe, Cisco, and so on)
- Contains both uppercase and lowercase letters.
- Contains numbers

If strong password checking is enabled after passwords have already been set, the system will not retroactively validate any existing passwords.

Although not recommended, password checking can be disabled by using the no password strength-checking command or the system setup script.

Setting Up the Exec Timeout Value

To set the interval that the EXEC command interpreter waits for user input before it terminates a session, run the exec-timeout line configuration command. The exec-timeout command must be used to log out sessions on a vty or physical terminal line (tty) that is left idle (inactive). By default, in Cisco NX-OS, sessions are set to disconnect after 30 minutes of inactivity.

```
!
line console
 exec-timeout <minutes>
line vty
 exec-timeout <minutes>
!
```

Disabling Unused Services

As a general security best practice, disable any unnecessary services. Cisco NX-OS does not run any of the typical Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) small servers often found in Cisco IOS software or other network operating systems by default. As a result, these services do not need to be explicitly disabled. Cisco NX-OS is designed to not run remotely accessible services or protocols, by default, without explicit configuration. SSH, SNMP, and NTP are essential services to run and manage a network. These services are enabled by default. If needed, they can be individually disabled. During the initial setup, Cisco NX-OS offers the option to enable Telnet. Note that this service will not load or run at boot time if it is not enabled during this initial setup. If this service is not enabled when the setup script is run, they can be added manually later, if needed. Cisco recommendation is to use SSH instead of telnet for security reasons.

Cisco Discovery Protocol is a network protocol that is used to discover other devices enabled for Cisco Discovery Protocol for neighbor adjacency and to map a network topology. Cisco Discovery Protocol can be used by network management systems or during troubleshooting. Cisco Discovery Protocol is enabled by default in Cisco NX-OS. Cisco Discovery Protocol must be disabled on all interfaces that are connected to untrusted networks. This disabling is accomplished with the no cdp enable interface command. Alternatively, Cisco Discovery Protocol can be disabled globally with the no cdp enable global configuration command. Note that Cisco Discovery Protocol can be exploited by malicious users, or reconnaissance and network mapping.

Link Layer Discovery Protocol (LLDP) is an IEEE protocol defined in the IEEE 802.1AB standard. LLDP is similar to Cisco Discovery Protocol; however, this protocol allows interoperability between devices that are not supported by the Cisco Discovery Protocol. By default, LLDP is not enabled in Cisco NX-OS. To enable it, the feature set must be enabled by using the feature lldp global configuration command. When enabled, LLDP must be treated in the same manner as Cisco Discovery Protocol and disabled on all interfaces that connect to untrusted networks.

To accomplish this, run the no lldp transmit and no lldp receive interface configuration commands. Run the no feature lldp configuration command to disable LLDP globally. Like Cisco Discovery Protocol, LLDP has the potential to be exploited by malicious users for reconnaissance and network mapping.

Consent Token

A consent token is the centralized mechanism to provide secure, transparent, customer-authorized feature enablement on Cisco Nexus 9000 switches in an auditable and trackable process. A consent token is also a form of a multi-factor authentication system that creates a common client-server infrastructure to allow functionalities such as secure shell access on Cisco Nexus 9000 switches, leveraging the available digital signature verification infrastructure.

The secure shell access primary use case is to provide restrictive, time-bound root shell access to Cisco Nexus 9000 switches with customer consent.

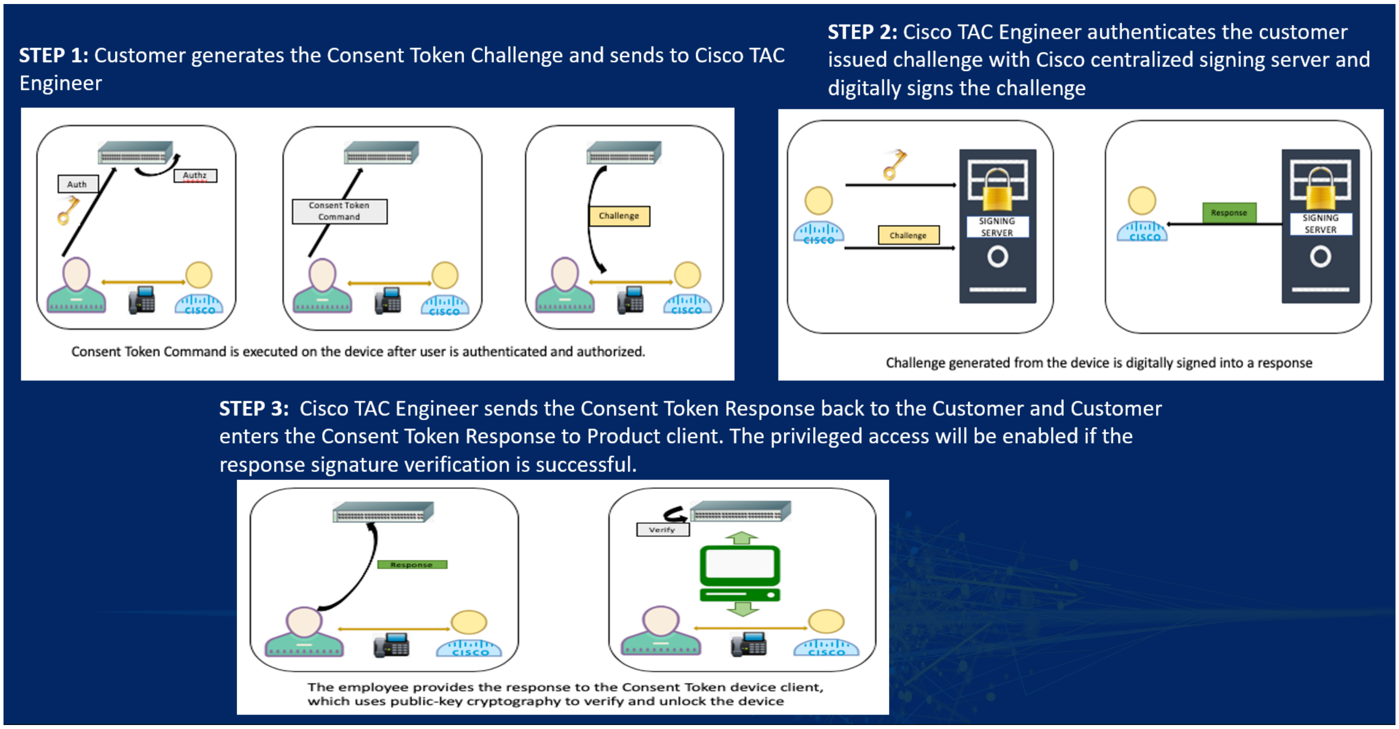

STEP 1: Customer generates the Consent Token Challenge and sends to Cisco TAC Engineer

Consent Token Command is executed on the device after user is authenticated and authorized.

STEP 2: Cisco TAC Engineer authenticates the customer issued challenge with Cisco centralized signing server and digitally signs the challenge

Challenge generated from the device is digitally signed into a response

STEP 3: Cisco TAC Engineer sends the Consent Token Response back to the Customer and Customer enters the Consent Token Response to Product client. The privileged access will be enabled if the response signature verification is successful.

The employee provides the response to the Consent Token device client, which uses public-key cryptography to verify and unlock the device

Figure 3 Consent token for secure shell access.

FIPS Mode

Cisco Nexus OS allows administrators to enable FIPS mode. When FIPS mode is enabled, Nexus 9K switches will change the underlying crypto libraries being used and start using FIPS (FIPS 140-2) approved cryptographic module.

Note: Enabling FIPS mode requires a system-wide reboot to take effect.

After enabling FIPS mode, Cisco N9K switches will be FIPS 140-2 compliant, as stated in the Compliance Letter.

This FIPS-compliant cryptographic module will be used not only for the encrypted SSL sessions, but also for northbound protocols that use encryption algorithms.

The FIPS Object Module is supported for the following protocols:

- TLS v1.2 and v1.3
- Syslog
- Syslog, Radius
- SSHv2
- SNMPv3

Therefore, before enabling FIPS mode, we strongly recommend that you disable non-supported versions of the protocols listed above, which aligns with the recommendations provided in this hardening guide.

For more information about guidelines and limitations and configuration for FIPS, see the Cisco NXOS Security Configuration Guide.

Securing Simple Network Management Protocol (SNMP)

Simple Network Management Protocol (SNMP) has been extensively used by organizations around the globe to monitor their network devices. It is still widely used even though the trend is moving towards monitoring using REST API, gRPC NMI, and other methods. Therefore, Cico Nexus OS supports both SNMPv2c and SNMPv3 for both polling (GET) and notifications (TRAP).

Regarding which version to use, SNMPv2c presents important security limitations that make the protocol less desirable. For example, communities are sent in clear text and hence can be intercepted and exposed. SNMPv3 supports both authentication and encryption, and therefore it is a much secure alternative. Because of this, we strongly recommend that you use SNMPv3 whenever possible.

Securing Interactive Management Sessions

Management sessions for devices allow you to view and collect information about a device and its operations. If this information is disclosed to a malicious user, the device can become the target of an attack, compromised, and commandeered to perform additional attacks. Anyone with privileged access to a device has the capability for full administrative control of that device. Securing management sessions is imperative to prevent information disclosure and unauthorized access.

Encrypting Management Sessions

Because information can be disclosed during an interactive management session, this traffic must be encrypted so that a malicious user cannot gain access to the data being transmitted. Encrypting the traffic allows a secure remote-access connection to the device. If the traffic for a management session is sent over the network in clear text, an attacker can obtain sensitive information about the device and the network.

An administrator can establish an encrypted and secure remote access management connection to a device by using SSH. Cisco NX-OS supports SSH Version 2.0 (SSHv2) only. Note that SSHv1 and v2 are not compatible.

Cisco NX-OS also supports SCP and Secure FTP (SFTP), which allows an encrypted and secure connection for copying device configurations or software images. SCP relies on SSH. This example configuration enables SSH on a Cisco NX-OS device:

```
!
ip domain-name example.com
!
feature ssh
ssh key rsa 2048
!
ssh login-attempts <1-10> (default is 3)
!
This configuration example enables the SCP and SFTP services:
```

```
!
feature scp-server
feature sftp-server
!
```

AAA

The AAA framework is critical to securing interactive access to network devices. The AAA framework provides a highly configurable environment that can be tailored depending on the needs of the network.

TACACS+ Authentication

TACACS+ is an authentication protocol that Cisco NX-OS devices can use for authentication of management users against a remote AAA server. These management users can access the Cisco NX-OS device through SSH or Telnet.

TACACS+ authentication, or more generally AAA authentication, provides the capability to centralize authentication information and authorization policies. It also enables effective centralized accounting of AAA-related transactions for improved auditability.

RADIUS is a protocol similar in purpose to TACACS+; however, RADIUS encrypts only the password sent across the network. In contrast, TACACS+ encrypts the entire TCP payload, including both the username and password. For this reason, TACACS+ is preferred over RADIUS when TACACS+ is supported by the AAA server and network device. Refer to TACACS+ and RADIUS Comparison design technote for a more detailed comparison of these two protocols.

TACACS+ authentication can be enabled on a Cisco NX-OS device using a configuration similar to this example:

```
!
! TACACS+ must be enabled in NX-OS
feature tacacs+
aaa authentication login default group tacacs+
!
tacacs-server host <ip-address-of-tacacs-server>
tacacs-server key [0 | 6 | 7] <key>
!
!
```

The previous configuration can be used as a starting point for an organization-specific AAA authentication template. Refer to the Use Authentication, Authorization, and accounting section of this document for more information about the configuration of AAA.

Authentication Fallback

If all configured AAA servers become unavailable, then a Cisco NX-OS device can rely on secondary authentication methods. Configuration options include the use of local or no authentication if all configured TACACS+ servers are unavailable. You should not use the None option, which in effect would fall back to no authentication if the AAA servers are unreachable. This fallback would potentially allow a DoS attack on the AAA servers to eliminate authentication on the network devices. Instead, authentication fallback should be set to use the local database when AAA servers are unreachable. This approach allows a locally defined user to be created for one or more network administrators. If TACACS+ were to become completely unavailable, each administrator can use his or her local username and password.

Although this action does enhance the accountability of network administrators during TACACS+ outages, it can increase the administrative overhead since local user accounts on all network devices must be maintained. Some of this overhead can be reduced by using the Cisco NX-OS TACACS+ configuration distribution mechanism, which uses the Cisco Fabric Services protocol.

This configuration example builds on the previous TACACS+ authentication example, including fallback authentication to the password that is configured locally with the enable secret command:

```
!
username admin password <password> role network-admin
!
aaa authentication login default group tacacs+
aaa authentication login default fallback error local
!
```

Refer to Configuring Authentication for more information about the use of fallback authentication with AAA.

TACACS+ Command Authorization

Command authorization with TACACS+ and AAA provides a mechanism that permits or denies each command that is entered by an administrative user. When the user enters EXEC or configuration commands, Cisco NX-OS sends each command to the configured AAA server. The AAA server then uses its configured policies to permit or deny the command for that particular user.

This configuration can be added to the previous AAA authentication example to implement command authorization:

```
!
aaa authorization commands default group <server group> [local]
aaa authorization config-commands default group <server group> [local]
!
```

One-Time Password Support

A one-time password (OTP) is a password that is valid for a single login session or a transaction. OTPs avoid multiple disadvantages that are associated with the static passwords. OTPs are not at risk to replay attacks. If an intruder manages to record an OTP that was already used to log into a service or to conduct an operation, it cannot be misused because it is no longer valid.

OTPs are applicable only to the RADIUS and TACACS+ protocol daemons. For a RADIUS protocol daemon, you must ensure that you disable the ASCII authentication mode. For a TACACS+ protocol daemon, you must enable the ASCII authentication mode. To enable the ASCII authentication mode, use the aaa authentication login ascii-authentication command.

Refer to the Configuring AAA section in the Cisco NX-OS Security Configuration Guide for more information about command authorization. For more information about command authorization.

Secure POAP (PowerOn Auto Provisioning)

PowerOn Auto Provisioning (POAP) is a feature designed to automate the initial setup and configuration of Cisco Nexus switches. It simplifies the deployment process by enabling switches to automatically discover and download their startup configuration and software images without manual intervention.

When a device with the POAP feature boots and does not find the startup configuration, the device enters POAP mode, locates a DHCP server, and bootstraps itself with its interface IP address, gateway, and DNS server IP addresses. The device also obtains the IP address of a TFTP server and downloads a configuration script that enables the switch to download and install the appropriate software image and configuration file.

Initially, POAP used common DHCP options to find the boot script, which was then transferred using **HTTP**. This method was **not secure**, leaving the script vulnerable to interception or tampering.

To address this, Cisco introduced **new, secure DHCP options** for both IPv4 and IPv6. These options enable POAP to reach the file server in a **secure manner**, implying the use of encrypted transfer protocols. While the older, less secure options remain supported for compatibility, using the secure options is highly recommended.

Further enhancing security, POAP gained the ability to use **Root-CA bundles** instead of single certificates for trusting provisioning script information. This improves how the switch verifies the authenticity and integrity of the source.

Securing the Control Plane

Control-plane functions consist of the protocols and processes that communicate between network devices to move data from the source to the destination. These include routing protocols such as BGP and protocols such as ICMP.

It is important that events in the management and data plane do not adversely affect the control plane. If a data plane event such as a DoS attack affects the control plane, the entire network can become unstable. The information in this section about Cisco NX-OS features and configurations can help ensure the resilience of the control plane.

General Control-Plane Hardening

Protecting the control plane of a network device is critical because the control plane helps ensure that the management and data planes are maintained and operational. If the control plane were to become unstable during a security incident, it can be impossible to recover the stability of the network.

In many cases, disabling the reception and transmission of certain types of messages on an interface can reduce the CPU load that is required to process unneeded packets.

IP ICMP Redirect Messages

An ICMP redirect message can be generated by a router when a packet is received and transmitted on the same interface. In this situation, the router forwards the packet and sends an ICMP redirect message back to the sender of the original packet. This behavior allows the sender to bypass the router and forward future packets directly to the destination (or to a router closer to the destination). In a properly functioning IP network, a router sends redirect messages only to hosts on its own local subnets. In other words, ICMP redirect messages should never go beyond a Layer 3 boundary.

There are two types of ICMP redirect messages: redirect messages for a host address, and redirect messages for an entire subnet. A malicious user can exploit the capability of the router to send ICMP redirect messages by continually sending packets to the router, forcing the router to respond with ICMP redirect messages, resulting in adverse impact on the CPU and on the performance of the router. To prevent the router from sending ICMP redirect messages, use the no ip redirects interface configuration command.

ICMP Unreachable Messages

Filtering with an interface access list elicits the transmission of ICMP unreachable messages back to the source of the filtered traffic. Generating these messages can increase CPU utilization on the device. You can disable ICMP unreachable message generation using the interface configuration command no ip unreachable.

Note that the default behavior of ICMP unreachable messages may vary depending on the hardware platform, the Cisco NX-OS Software release, and whether the device interface is in Layer 2 mode or Layer 3 mode. Users are encouraged to test specific ICMP unreachable message behavior in their environments.

Proxy Address Resolution Protocol

Proxy Address Resolution Protocol

Proxy Address Resolution Protocol (ARP) is the technique in which one device, usually a router, answers ARP requests that are intended for another device. By “faking” its identity, the router accepts responsibility for routing packets to the real destination. Proxy ARP can help machines on a subnet reach remote subnets without the need to configure routing or a default gateway. Proxy ARP is defined in RFC 1027.

There are several disadvantages to using proxy ARP. Proxy ARP can result in an increase in the amount of ARP traffic on the network segment and resource exhaustion and man-in-the-middle attacks. Proxy ARP presents a resource exhaustion attack vector because each proxied ARP request consumes a small amount of memory. An attacker could attempt to exhaust memory unnecessarily by sending a large number of ARP requests.

Man-in-the-middle attacks enable a host on the network to spoof the MAC address of the router, causing unsuspecting hosts to send traffic to the attacker. Proxy ARP can be disabled by using the interface configuration command no ip proxy-arp.

Understanding Control-Plane Traffic

To properly protect the control plane of the Cisco NX-OS device, you must understand the types of traffic that are processed by the CPU. Process-switched traffic normally consists of two types of traffic. The first type of traffic is directed to the Cisco NX-OS device and must be handled directly by the Cisco NX-OS device CPU. This traffic consists of this category:

- Receive adjacency traffic: This traffic contains an entry in the Cisco Express Forwarding table through which the next router hop is the device itself, which is indicated by the term receive in the show ip cef CLI output. This indication appears for any IP address that requires direct handling by the Cisco NX-OS device CPU, including interface IP addresses, multicast address space, and broadcast address space.

The second type of traffic that is handled by the CPU is data-plane traffic with a destination beyond the Cisco NX-OS device itself that requires special processing by the CPU. This type of behavior tends to be platform specific and dependent on the specific hardware implementation of the specific Cisco NX-OS platform. Some platforms handle more types of data-plane traffic in hardware, thereby requiring less CPU-based intervention. Regardless of the hardware handling capabilities, you should understand potential sources of control-plane traffic that could affect the CPU system. Although not an exhaustive list of data-plane traffic that can affect the CPU, these types of traffic are potentially processed and can therefore affect the operation of the control plane:

- ACL logging: ACL logging traffic consists of any packets that are generated due to a match (permit or deny) of an access control entry on which the log keyword is used.

- Unicast Reverse Path Forwarding (uRPF): uRPF, used in conjunction with an ACL, can result in the process of switching of certain packets.
- IP options: Any IP packets with options included must be processed by the CPU.
- Fragmentation: Any IP packet that requires fragmentation must be passed to the CPU for processing.
- Time-to-live (TTL) expiry: Packets that have a TTL value less than or equal to 1 require ICMP Time Exceeded (ICMP Type 11, Code 0) messages to be sent, which results in CPU processing.
- ICMP unreachable messages: Packets that result in ICMP unreachable messages due to routing, maximum transmission unit (MTU), or filtering are processed by the CPU.
- Traffic requiring an ARP request: Destinations for which an ARP entry does not exist require processing by the CPU.
- Non-IP traffic: All non-IP traffic is processed by the CPU.

The following list details several methods to determine which type of traffic is being processed by the Cisco NX-OS device CPU:

- The show ip cef command provides the next-hop information for each IP prefix that is contained in the Cisco Express Forwarding table. As indicated previously, entries that contain receive as the next hop are considered receive adjacencies and indicate that traffic must be sent directly to the CPU.
- The show interface switching command provides information about the number of packets being process switched by a device.
- The show ip traffic command provides information about the number of IP packets:
- With a local destination (that is, receive adjacency traffic)
- With options
- That require fragmentation.
- That are sent to broadcast address space
- That are sent to multicast address space

Receive adjacency traffic can be identified by using the show ip cache flow command. Any flows that are destined for the Cisco NX-OS device have a destination interface (DstIf) of local.

CoPP can be used to identify the type and rate of traffic that reaches the control plane of the Cisco NX-OS device. CoPP can be performed through the use of granular classification ACLs, logging, and the show policy-map control-plane command.

CoPP

The CoPP feature can also be used to restrict IP packets that are destined for the infrastructure device itself and require control-plane CPU processing. CoPP in Cisco NX-OS can be used to police different classes of traffic to different permitted levels, effectively applying quality of service (QoS) to control-plane-bound traffic.

The configuration of CoPP is similar to data-plane QoS configuration and uses the same Modular QoS CLI (MQC) configuration structures:

- Class maps are defined to match specific types of traffic.
- Policy maps are created to apply policing (rate-limiting) policies to class-map-matched traffic.

- A service policy is used to map the policy map to the control-plane interface.

Cisco NX-OS provides simplified setup for typical network environments by offering predefined class maps and policy maps by using the initial configuration setup script. When you run the setup script, or at bootup, you can select one of the four predefined templates to be applied for CoPP:

- Strict
- Moderate
- Loose
- None

After a CoPP template is selected, it will be applied to the control-plane interface. If the CoPP policy is changed from one of the actively policing templates (strict, moderate, or loose) to none, the system will not remove the existing class maps or policy maps. It will not map the policy map to the control-plane interface with a service policy. This approach leaves the configuration in place but does not apply it to the interface. If a different active policing template is chosen to replace one in place, the template will overwrite the existing class maps and policy maps with the new settings.

In this example, a CoPP configuration is created in which SSH traffic only from trusted hosts is permitted to reach the Cisco NX-OS device CPU. All other control-plane traffic is allowed:

Note: Dropping traffic from unknown or untrusted IP addresses can prevent hosts with dynamically assigned IP addresses from connecting to the Cisco NX-OS device.

```
!
access-list ALLOW_TRUSTED_SSH
   deny tcp <trusted-addresses> <mask> any eq 22
   permit tcp any any eq 22
   deny ip any any
!
class-map type control-plane match-all COPP-KNOWN-UNDESIRABLE
   match access-group name ALLOW_TRUSTED_SSH
!
policy-map type control-plane COPP-INPUT-POLICY
   class COPP-KNOWN-UNDESIRABLE
          police 1 conform drop violate drop
!
control-plane
   service-policy input COPP-INPUT-POLICY
!
```

In the preceding CoPP example, the ACL entries that match the unauthorized packets with the permit action result in a discard of these packets by the policy-map drop function, while packets that match the deny action are not affected by the policy-map drop function.

This example illustrates the theory, structure, and applicability of CoPP. In reality, an effective CoPP policy is more complex than the simplified example shown here and requires adequate planning and testing before being deployed in a live production environment.

Securing the Data Plane

Although the data plane is responsible for moving data from the source to the destination, within the context of security the data plane is the least important of the three planes. For this reason, when securing a network device, you should protect the management and control planes in preference over the data plane. However, within the data plane itself, there are many features and configuration options that can help secure traffic. The following sections detail these features and options so that you can more easily secure your network.

General Data-Plane Hardening

Most data plane traffic flows across the network as determined by the network's routing configuration. However, IP network functions are available to alter the path of packets across the network. Features such as IP options - specifically, the source routing option - create security challenges in today's networks.

Disabling IP Source Routing

IP source routing uses the Loose Source Route and Record Route options in tandem or the Strict Source Route along with the Record Route option to enable the source of the IP datagram to specify the network path that a packet takes. This function can be used in an attempt to route traffic around security controls in the network.

If IP options have not been completely disabled though the IP Options Selective Drop feature, it is important that you disable IP source routing. IP source routing, which is enabled by default in all Cisco NX-OS releases, is disabled through the no ip source-route global configuration command. This configuration example illustrates the use of this command:

```
!
no ip source-route
!
```

Disabling ICMP Redirect messages

ICMP redirect messages are used to inform a network device of a better path to an IP destination. By default, Cisco NX-OS sends a redirect message if it receives a packet that must be routed through the interface from which it was received.

In some situations, an attacker may be able to cause the Cisco NX-OS device to send many ICMP redirect messages, resulting in an elevated CPU load. For this reason, the transmission of ICMP redirect messages should be disabled. ICMP redirect messages are disabled by using the interface configuration command no ip redirects, as shown in the example configuration:

```
!
interface ethernet 1/1
          no ip redirects
!
```

Disabling or Limiting IP Directed Broadcasts

IP directed broadcasts make it possible to send an IP broadcast packet to a remote IP subnet. After the packet reaches the remote network, the forwarding IP device sends the packet as a Layer 2 broadcast to all stations on the subnet. This directed broadcast function has been used as an amplification and reflection aid in several attacks, including the smurf attack.

Filtering ICMP Packets

ICMP was designed as a control protocol for IP. As such, the messages it conveys can have far-reaching ramifications on the TCP and IP protocols in general. ICMP is used by the network troubleshooting tools ping and traceroute, as well as by path MTU discovery; however, external ICMP connectivity is rarely needed for the proper operation of a network.

Cisco NX-OS provides functions to specifically filter ICMP messages by name or type and code. This example ACL allows ICMP from trusted networks while blocking all ICMP packets from other sources:

```
!
ip access-list ACL-TRANSIT-IN
! !--- Permit ICMP packets from trusted networks only !
 permit icmp <trusted-networks>/<mask> any
! !--- Deny all other IP traffic to any network device !
 deny icmp any any
!
```

Filtering IP Fragments

As discussed previously in the Limiting Access to the Network with Infrastructure ACLs section of this document, the filtering of fragmented IP packets can pose a challenge to security devices.

Because of the nonintuitive nature of fragment handling, IP fragments are often inadvertently permitted by ACLs. Fragmentation is also often used in attempts to evade detection by intrusion-detection systems. For these reasons, IP fragments are often used in attacks and should be explicitly filtered at the top of any configured ACLs. The ACL shown here includes comprehensive filtering of IP fragments. The function illustrated in this example must be used in conjunction with the functions shown in the previous examples:

```
!
 ip access-list ACL-TRANSIT-IN
!
!--- Deny IP fragments using protocol-specific ACEs to aid in
!--- classification of attack traffic !
   deny tcp any any fragments
   deny udp any any fragments
   deny icmp any any fragments
   deny ip any any fragments
!
```

Using IP Source Guard

IP source guard is an effective means of spoofing prevention that can be used if you have control over Layer 2 interfaces. IP source guard uses information from Dynamic Host Configuration Protocol (DHCP) snooping to dynamically configure a port ACL (PACL) on the Layer 2 interface, denying any traffic from IP addresses that are not associated in the IP source binding table.

IP source guard can be applied to Layer 2 interfaces belonging to VLANs enabled for DHCP snooping. These commands enable DHCP snooping:

```
!
```

```
feature dhcp
ip dhcp snooping
!
```

After DHCP snooping is enabled, these commands enable IP source guard:

```
!
interface ethernet <slot>/<port>
   ip verify source dhcp-snooping vlan
!
```

Using Port Security

Port security is used to mitigate MAC address spoofing at the access interface. Port security can use dynamically learned (sticky) MAC addresses to facilitate the initial configuration. After port security has determined a MAC address violation, it can use one of the four violation modes: protect, restrict, shutdown, and shutdown VLAN. In instances in which a port provides access only for a single workstation using standard protocols, a maximum value of 1 may be sufficient. Protocols that use virtual MAC addresses such as Hot Standby Router Protocol (HSRP) do not function when the maximum value is set to 1.

```
!
feature port-security
interface <slot>/<port>
   switchport
   switchport port-security [mac address sticky]
!-- Optionally enable sticky MAC address learning
!
```

DAI (Dynamic Arp Inspection)

DAI can be used to mitigate ARP poisoning attacks on local segments. An ARP poisoning attack is a method in which an attacker sends falsified ARP information to a local segment. This information is designed to corrupt the ARP cache of other devices. Often an attacker uses ARP poisoning to perform a man-in-the-middle attack.

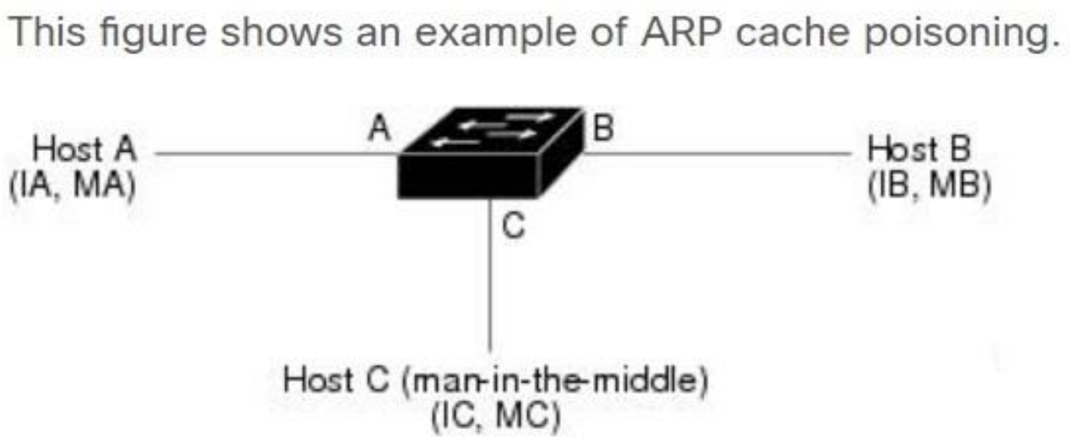
This figure shows an example of ARP cache poisoning.
A
B
Host A
(IA, MA)
Host B
(IB, MB)
C
Host C (man-in-the-middle)
(IC, MC)

Figure 4 ARP Cache Poisoning

DAI intercepts and validates the IP-to-MAC address relationship of all ARP packets on untrusted ports. In DHCP environments, DAI uses the data that is generated by the DHCP snooping feature. ARP packets that are received on trusted interfaces are not validated, and invalid packets on untrusted interfaces are discarded. In non-DHCP environments, the use of ARP ACLs is required.

These commands enable DHCP snooping:

```
!

feature dhcp

ip dhcp snooping

ip dhcp snooping vlan <vlan-range>

!
```

After DHCP snooping has been enabled, these commands enable DAI:

```
!

ip arp inspection vlan <vlan-range>

!
```

Switchport Blocking

Occasionally, unknown multicast or unicast traffic is flooded to a switch port because a MAC address has timed out or has not been learned by the switch. Security issues could arise if unknown multicast and unicast traffic is forwarded to a switch port. You can enable switchport blocking to guarantee that no multicast or unicast traffic is flooded to the port.

Prevents the flooding of unknown multicast or unicast packets on the specified interface.

```
switch# configure terminal
switch(config)#      interface ethernet 1/2
switch(config-if)# switchport block multicast
switch(config-if)# switchport block unicast
switch(config-if)# show running-config interface ethernet 1/2
!Command: show running-config interface Ethernet1/2
```

MACsec

Cisco Nexus 9000 Series switches offer support for Media Access Control Security (MACsec), which is a security technology that provides secure communication for all traffic on Ethernet links. MACsec provides Layer 2 encryption and is standardized by IEEE 802.1AE.

Here is a summary of the main features and aspects of MACsec on Cisco Nexus 9K switches:

Encryption: MACsec encrypts the entire Ethernet frame, providing privacy and ensuring that data cannot be eavesdropped on the link.

Integrity and Authentication: It ensures that frames are not tampered with in transit (integrity) and provides origin authentication, confirming that each frame comes from the entity that claims to send it.

Point-to-point security: MACsec operates on point-to-point links, which means that both ends of a direct connection can be secured. This is typically used between switches or between a switch and a host.

Secure Key Exchange: MACsec uses the MACsec Key Agreement (MKA) protocol for the secure distribution of keys. MKA is part of the IEEE 802.1X-2010 standard.

Scalability: The Nexus 9K switches support MACsec on multiple ports, allowing for a scalable implementation of network encryption.

Line-rate encryption: Cisco's implementation of MACsec can encrypt traffic at a line-rate speed, ensuring no degradation in performance when MACsec is enabled.

Compatibility: MACsec on Nexus 9K is compatible with other Cisco devices that support MACsec, allowing for interoperability within the Cisco ecosystem.

Configurability: MACsec can be configured on the Nexus 9K through the Command Line Interface (CLI), providing granular control over the security settings and policies.

Hardware Support: MACsec support is often hardware-dependent, requiring specific modules or versions of the Nexus 9K for the functionality to be available.

License Requirement: Enabling MACsec on Nexus 9000 Series switches may require an additional security license depending on the model and the software version.

It is important to check the specific Nexus 9000 model and its software release notes for detailed information about MACsec support and capabilities, as they can vary between different hardware generations and software releases. Additionally, MACsec configuration and deployment should be planned carefully to ensure compatibility and to maintain desired network performance.

For more information, see the Cisco NXOS Security Configuration Guide.

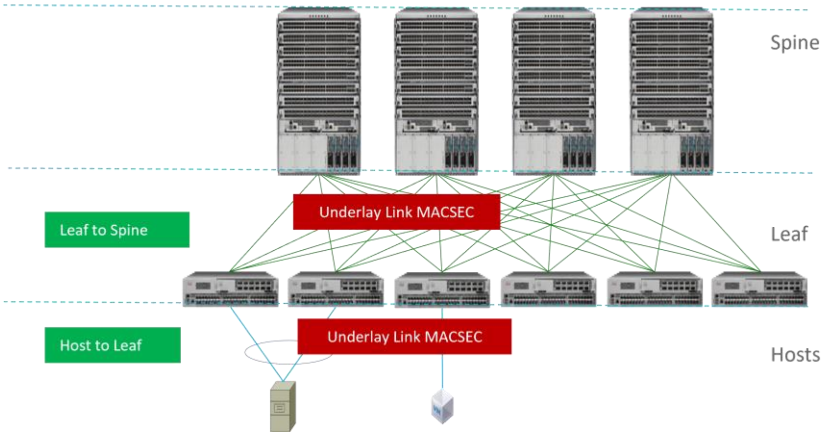

Figure 5 MACsec in DC Fabric

MACsec- QKD Integration with Secure Key Integration Protocol

About Quantum-Safe Encryption

Recent advancements in quantum computing have exposed vulnerabilities in various cryptographic algorithms, making them unsecured for future applications. The RSA (integer factorization) and DHE (discrete logarithms) public key algorithms, which rely on computational complexity, are now at risk of being solved by quantum computers using Shor's or Grover's algorithm.

As a result, establishing a shared secret key between communicating parties has become a significant challenge. To avoid this issue, you can configure quantum-safe algorithms or implement a Quantum Key Distribution (QKD).

Integrating SKIP protocol to the switches empowers to establish communication with external quantum devices. This advancement allows for the utilization of Quantum Key Distribution (QKD) devices in the exchange of MACsec encryption keys between switches.

Beginning with Cisco NX-OS Release 10.4(3)F, Nexus 9300-FX/FX2/GX2/H2R Platform switches support Secure Key Integration Protocol (SKIP) protocol.

QKD operates on the principles of quantum physics, utilizing the quantum state of photons to encode and share information through an optical link. Additionally, an authenticated classical channel is used for sharing measurements. The change in quantum states helps the two end parties of the communication channel to identify any interception of their key.

QKD is a secure key exchange mechanism against quantum attacks even in the future advancements in crypto Analysis or quantum computing. QKD doesn’t require continual updates based on discovered vulnerabilities.

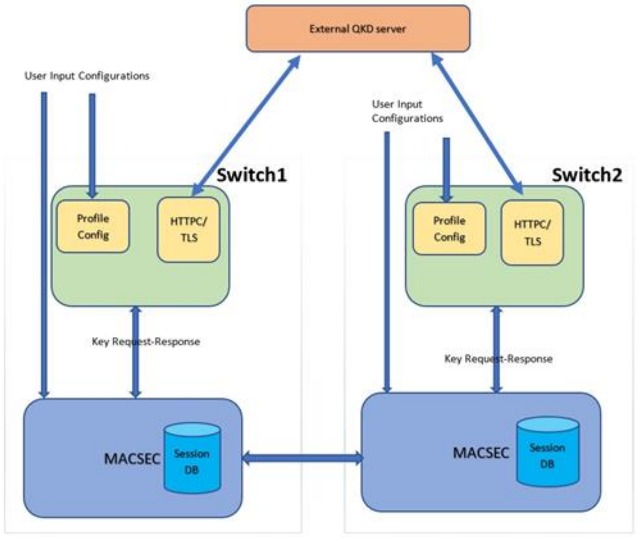

Figure 6 MACsec with QKD server Integration.

Conclusion

In conclusion, in a world where cyber-attacks are becoming increasingly frequent and sophisticated, infrastructure hardening and being aware of the security features on Nexus 9000 switches is more important than ever before. By implementing a range of security measures to strengthen their digital infrastructure, organizations can reduce their risk exposure, protect their critical assets and information, and maintain their operations even in the face of persistent threats.

As outlined the document cisco Nexus OS implements multiple security mechanisms to reduce the attack surface and provide organizations with a solid and secure data center network infrastructure. This level of security by default, together with the recommendations mentioned in this document, make your Datacenter a solid infrastructure that meets the more demanding security requirements.

References

Cisco Nexus 9000 Series NX-OS Security Configuration Guide