

Cisco Application Centric Infrastructure Security: Chain of Trust

What You Will Learn

This document discusses the security of the Cisco® Application Policy Infrastructure Controller (APIC) and fabric switches. Another document discusses tenant security and micro-segmentation.

Zero Trust Model for Multi-Tenant Security

Cisco ACI fabric is inherently secure because it uses a zero-trust model and relies on many layers of security. Many security measures are implemented to protect against and overcome all security attack vectors:

- Unauthorized access
- Man-in-the-middle-attack
- Replay attack
- Data disclosure
- Denial of service

All user system access or API calls require AAA and role-based access control that restricts the read/write access of tenant sub-tree read or write. Northbound interfaces utilize certificates and encryption. Rogue or counterfeit devices cannot access fabric resource because ACI fabric utilizes a HW key store and requires certificate based authentication. Within the fabric, the infrastructure VLAN is (used for APIC to Switch communication) an isolated space and all messages are encrypted. All software images and binaries are signed and verified before they can bootup a device with ACI fabric.

The sections below provide describe Cisco ACI Chain of Trust in more details.

System Access

All management interfaces (representational state transfer [REST], command-line interface [CLI] and GUI) are authenticated in Cisco ACI using authentication, authorization, and accounting (AAA) services (LDAP and Microsoft Active Directory, RADIUS, and TACACS+) and RBAC policies, which map users to roles and domains.

Cisco ACI systems can be accessed in five ways:

- APIC local authentication
- External RADIUS
- External TACACS+
- External LDAP and Active Directory
- Local authentication (not common)

APIC local authentication doesn't require external AAA servers (RADIUS, TACACS+, or LDAP and Active Directory). This type of authentication is common in lab, proof-of-concept, and small customer data center deployments. This user can use the username, password, Secure Shell (SSH) keys, and user certificate. After login, RBAC rules will place this user in the correct security domain and assign read and write privilege levels based on the user role.

External RADIUS can be used for AAA service to access Cisco ACI. The customer administrator must configure the required attributes (**shell:domains**) using a Cisco attribute-value pair. Note (for instance, for RADIUS requests for comments [RFCs]) that the Cisco vendor ID is 9 and the supported option is vendor type 1, which is called **cisco-av-pair**.

```
shell:domains="domainA/write_roles/read_roles, domainB/write_roles/read_roles/,..."
```

Cisco ACI also supports external TACACS+ AAA. This approach has some advantages over RADIUS AAA:

- It provides independent AAA service; for example, the Cisco ACI switch can authorize access without authenticating.
- It uses the TCP transport protocol to send data between the AAA client and the server, providing increased reliability.
- It encrypts the entire protocol payload between the switch and the AAA server to help ensure greater data confidentiality. The RADIUS protocol only encrypts passwords.

Cisco ACI also supports external LDAP and Active Directory AAA. LDAP allows a network element to retrieve AAA credentials, which can be used to authenticate and then authorize the user to perform certain actions. An additional certificate authority configuration can also be used to enable LDAP over SSL trust and prevent man-in-the-middle attacks.

Role-Based Access Control

Cisco ACI implements a role-based access control (RBAC) scheme that restricts not only write access but also the visibility of tenant subtrees to only those users who have read privileges to those tenant subtrees. The RBAC rules are uniform across the APIC controller and leaf-spine switches. A user or API client accessing Cisco APIC on the northbound interface will be valid on the all switches. If the switches need to access external LDAP, RADIUS, or TACACS servers, these servers should be accessible over the infra-network or, alternately, the switches should have management network IP addresses.

The controller forwards requests for data that is stored only on the switches to the switch with the user authentication context contained in the request, and the switch RBAC code validates the request using these credentials prior to returning the requested data or performing the requested action.

Cisco ACI fabric RBAC rules allow tenants to modify configuration and parameters of the ACI fabric that they own and control. Tenants have a self-service model by which they can perform configuration or parameter changes; read statistics; and monitor faults, events, logs, etc. for the entries (that is, managed objects) that apply only to them. Such entries include endpoints, EPGs, contexts, and application network profiles. They cannot access (that is, read, write, or modify) parts of the fabric that RBAC rules restrict.

Cisco ACI represents all managed resources, both physical and virtual, as managed objects. These objects constitute an abstract representation of the entire fabric. Managed objects are instantiated in a logical, hierarchical tree called management information tree (MIT). RBAC rules allow selective read and write to various parts of the MIT, but the rules also prevent read and visibility access to other tenants' data in the tree.

Cisco ACI implements two levels of access control:

- Traditional role-based control: This type of control defines the types of objects that a user is authorized to access. Users are assigned roles (collections of privileges) that govern read and write access to managed objects in the system. All managed object classes have one or more privileges assigned to them.
- Domain-based control: This type of control defines where a user is authorized to access objects. Domain-based access is categorized into **read domains** and **write domains**.

Cisco APIC Northbound Interfaces

The following services use the APIC's northbound interfaces. Each service goes through the full AAA process and RBAC rules before it can access any fabric resource.

- CLI over SSH session
- HTTP/HTTPS for GUI or REST client
- Simple Network Management Protocol Versions 2 and 3 (SNMPv2 and v3) for switches (no APIC)

Northbound Interface: CLI over SSH Session

A CLI over SSH session uses TCP port 22 to log in to the APIC. Two login modes are supported:

- Standard SSH login: Requires username and password
- Password-free SSH login for local user: Requires username and public key

For standard SSH login, the username and password are authenticated using external RADIUS, TACACS+, and LDAP and Active Directory authentication servers. The connection between the SSH client and the APIC session is secure and encrypted using standard SSH encryption ciphers and hash algorithms. After AAA authentication is complete, the controller applies the RBAC rules to the login session and assigns a RBAC security-domain, role, and privilege level to the session. The login session is admitted to the controller and is assigned a ZSH session. The user can now run CLI **show** and **config** commands in the ZSH session. The ZSH session internally uses REST API for all **show** and **config** commands. All SSH login accounting records are aggregated and stored in the controller.

Customers who want to automate CLI tasks with shell scripts and other scripts can use SSH in the script to log in to the APIC. For such users, Cisco ACI supports password-free login authentication, which requires a username and SSH public key. In the APIC, you can configure a local user (username) and upload public keys for such users to allow logins using public and private RSA key pairs. In fact, this approach is more secure than a password-based approach. The APIC local-user AAA is using RADIUS, so neither TACACS+ nor LDAP or Active Directory is required for SSH password-free login. Also note that the session between SSH to the controller is secure and encrypted using TLS 1.2. After login authentication, the APIC will apply the RBAC rules and assign a RBAC security domain, role, and privilege level to the session. The script can now invoke any **show** and **config** commands. All SSH login accounting records are aggregated and stored in the controller. For high-security applications, SSH keys can be changed on the APIC at any time.

Northbound Interface: HTTP/HTTPS for GUI or REST Client

The APIC northbound REST interface is used by REST and Cisco ACI GUI clients. The REST interface can be accessed using HTTP or HTTPS. Inside the controller all HTTP port 80 requests are redirected to HTTPS port 443. The HTTP session is encrypted using TLS 1.2. The APIC provides two methods for accessing the REST interface:

- Webtoken login: Requires username and password
- X.509 certificate login: Requires username and X.509 certificate

The REST web token login method is very common in web browsers. Web token security works like this:

- In the beginning, the REST client sends the username and password to APIC for authentication.
- The APIC applies AAA services (RADIUS, TACACS+, or LDAP and Active Directory) and authenticates the client.
- The APIC then applies RBAC rules and assigns a RBAC security domain, role, and privilege level to the session.
- The APIC generates a web token (or Cisco ACI cookie) and sends it to the REST client.
- From this point on, the REST client uses only the web token to access the APIC (no username or password is required).
- To avoid man-in-the-middle attacks, the web token is encrypted using TLS 1.2.
- When the APIC receives the web token, the token is decrypted and used to enforce RBAC policy rules.
- The web token times out in 600 seconds and must be refreshed through a REST web (HTTP) client.
- The web token is valid only for 24 hours in an APIC for a session. After 24 hours, a new web token must be used.
- The APIC can store 32 SSH public keys per user. The APIC currently supports approximately 8000 users.

The REST X.509 certificate login method is the most secure approach and is the recommended method for REST clients to use to access the Cisco ACI fabric. X.509 certificates do not require login or logoff. Here is how X.509 certificate login works in Cisco ACI:

- The X.509 certificate is a container that has two parts:
 - RSA public key
 - RSA private key
- The user stores the X.509 certificate RSA public key in the APIC.
- The user independently manages the X.509 certificate RSA private key. (The X.509 RSA private key is not stored in the controller nor transmitted over the wire by the REST client. Its safety is the responsibility of the user.)
- Every REST client HTTP request or call to the APIC is signed with the user's RSA private key. X.509 cookies are generated and sent to the APIC in the HTTP header:
 - X.509 certificate fingerprint cookie (using a SHA-1 hash)
 - X.509 certificate distinguished name cookie
 - X.509 certificate RSA signature cookie
 - X.509 certificate algorithm cookie

- The APIC authenticates every REST client request. It retrieves the X.509 certificate fingerprint cookie value and looks up the X.509 public key certificate.
 - If no match is found, the REST request is discarded.
 - If a match is found, the APIC retrieves the signature cookie value and verifies the signature to complete the authentication process for the REST request.
- Authentication is complete. The APIC now applies the RBAC policy to this REST request.
- This process is repeated for every REST request.
- The APIC supports X.509 certificates that are signed using RSA keys of 1024, 1536, and 2048 bits only.
- For scalability, the APIC allows 32 X.509 certificates per user. The controller currently supports 8000 total users.

Northbound Interface: SNMPv2 and v3

SNMP is still widely used in existing (brownfield) and traditional data center environments, and some back-end systems may continue to use SNMP after customers migrate to Cisco ACI: for example, Tier 1 service providers in the United States and Europe.

The APIC GUI allows users to configure SNMP policies, which are pushed down and programmed in Cisco Nexus[®] leaf and spine switches. Then the Cisco ACI leaf and spine switches can directly use SNMPv2 and v3 on the northbound interface to enforce SNMP commands, and the controller is not involved in SNMP enforcement. The APIC only provides an authentication service to SNMPv2 and v3, and after that all SNMP commands (**GET** commands, MIB walks, etc.) are sent directly to the switch's northbound interface for processing.

SNMPv1 and v2 have very weak security models, performing authentication through a clear-text password (SNMP community string) without encryption. However, SNMPv3 provides privacy and security services. The controller can authenticate the SNMPv2 community string (that is, the clear-text password) or SNMPv3 user-based authentication. Table 1 presents a summary of all the models supported by Cisco ACI.

Table 1. SNMPv 2 and v3 Authentication Models Supported by Cisco ACI

Model	Level	Authentication	Encryption	Action
SNMPv2	noAuthNoPriv	Community string	No	Uses a community string match for authentication
SNMPv3	noAuthNoPriv	Username	No	Uses a username match for authentication
SNMPv3	AuthNoPriv	HMAC-MD5 or HMAC-SHA	No	Provides authentication based on the Hash-Based Message Authentication Code (HMAC) Message Digest 5 (MD5) algorithm or the HMAC Secure Hash Algorithm (SHA)
SNMPv3	authPriv	HMAC-MD5 or HMAC-SHA	DES	<ul style="list-style-type: none"> • Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithm • Provides Data Encryption Standard (DES) 56-bit encryption in addition to authentication based on the Cipher Block Chaining (CBC) DES (DES-56) standard

Password Recovery

In a Cisco ACI system, a user can recover the administrator (admin) password on the controller or leaf and spine switches. To recover the admin password, the user must have a direct physical serial console connection to the APIC server or leaf and spine switches. You cannot use a remote-access method to recover the admin password. The use of only special control characters during the bootup sequence helps ensure that the operator has power-cycle access to the APIC servers directly.

The password-recovery procedure is enhanced by the use of a USB dongle that contains a specific file that must be present when the system is booting and the key sequence is issued. The USB dongle must be physically connected to activate the password-recovery procedure.

This is the password recovery procedure:

1. The customer connects to the serial console. (For serial port settings, see Cisco.com documentation.)
2. The customer inserts the USB dongle into the server.
3. The customer reboots the server.
4. The customer enters the key sequence during the N seconds that the preboot prompt is visible on the serial console.
5. The daemon monitoring the serial console during the preboot prompt display is triggered upon key sequence detection and scans for the USB dongle and the data on it.
6. If the USB dongle is not found, nothing happens, and the system continues bootup.
7. If the USB dongle is found, the system displays a privileged CLI prompt, where the session context is **Privilege Bits: 0xffff** and **Domain: All** (basically, superuser mode). The customer can then issue CLI commands to set up the admin password.

Note: This procedure only resets the admin password with a new one. The old password was never stored in plain text. All passwords in the system are stored hashed with a salt (see the well-known LinkedIn issue with unsalted passwords) and thus cannot be recovered.

Cisco ACI Fabric Device Authentication and Security

The Cisco ACI fabric is inherently secure because it uses a zero-trust models and relies on many layers of security. Before any controller or leaf or spine switch becomes a member of the Cisco ACI fabric, it **must** be authenticated and admitted by the fabric administrator. After that, it becomes an operational component of the fabric.

All devices (the APIC and Cisco Nexus 9000 Series leaf and spine switches) use a hardware- based secure key store. The controller uses a Trusted Platform Module (TPM) hardware cryptographic module to store digitally signed certificates. The Cisco Nexus 9000 Series leaf and spine switches in the Cisco ACI fabric use the Anti-Counterfeit Technology 2 (ACT-2) hardware security module (HSM) embedded on the motherboard to store digitally signed certificates.

Cisco ACI security starts at the time of manufacture with the installation of keys in the hardware-based secure key store. The controller's TPM chip is programmed with the Cisco ACI platform's public key. The Cisco Nexus 9000 Series leaf and spine switches use an ACT-2 chip programmed with three keys: a development key, a release key, and a revocation key. All certificates in TPM and ACT-2 chips are unique, digitally signed, and encrypted at the time of manufacture. During Cisco ACI fabric activation or while adding a new device to an existing Cisco ACI fabric, all devices are authenticated based on their digitally signed certificates and identity information. If someone adds a rogue Cisco Nexus 9000 Series switch or APIC device to the Cisco ACI fabric, the certificate authentication will fail, and the infrastructure VLAN on the switch will not even open. As a result, the rogue device will not be added to the fabric and will not have any access to Cisco ACI resources.

The Cisco ACI fabric architecture completely separates and isolates all management VLAN, infrastructure VLAN, and tenant data-plane traffic from each other. The infrastructure VLAN traffic is fully (100 percent) isolated from all tenant (data plane) and management VLAN traffic. The infrastructure VLAN is used for fabric discovery and activation, image management, configuration, monitoring, and operation.

All communication between the APIC and the switch is over the intrafabric management (IFM) channel and is in band over the infrastructure VLAN on the switch. By default, the infrastructure VLAN is closed on every device and can be opened only after a device has been authenticated. All IFM channel communication over the infrastructure VLAN is encrypted using TLS 1.2, and every message that comes to the switch over the IFM channel must be decrypted before it is processed by the switch.

Image Signing and Bootup

All devices that use Cisco NX-OS Software ship with an ACT-2 HSM embedded on the board. The HSM is used to verify the entire chain of trust, first validating the field-programmable gate array (FPGA) software and then the ROM monitor (ROMMON) software, switch preboot image, and the switch's full image. These software bundles are generated by Cisco build servers that then use the Abraxas image signing system to securely sign these images with Cisco private RSA keys.

The image is verified by the ACT-2 HSM using Cisco release keys, which are the public keys that correspond to the Cisco release private key stored securely on the Abraxas servers. The HSM is tamper proof and does not permit this release key to be changed easily. The only way to change the release key in the HSM is to provide a bundle signed by the private part of the revocation key. The ACT-2 HSM verifies this signature using the public part of the revocation key stored inside it, and if this signature is verified, the release key change feature is unlocked.

The APIC downloads a signed NX-OS image to the leaf and spine switches. The switch image is digitally signed at build time. A SHA-512 hash is generated over the entire binary image file, and then the hash is encrypted with a Cisco ACI RSA 2048-bit private key. The switch verifies the signature using the Cisco ACI public key. If the software was not generated by a Cisco ACI build system, the signature verification stage will fail, and the switch will reject the image and stop booting. If the signature verification stage passes using the release key from the ACT-2 key store, then the switch boots the image to the runtime environment.

These are the steps for verifying a signed APIC image at bootup:

1. Download a signed image to the APIC.
2. Extract the signature from the image.
3. Verify the signature by using the APIC platform key TPM chip.
4. Compare the hash with the image's SHA-512 hash.
5. Check the model number for compatibility.
6. Decrypt the root partition to access Cisco ACI binary files and libraries on the Cisco ACI disk.
7. The Cisco ACI binary files and libraries run in a secure runtime environment.

These are the steps for verifying a signed NX-OS image at bootup:

1. Extract a signature from the image.
2. Get the development key and release key from the ACT-2 key store.
3. Compare the hash with the image's SHA-512 hash.
4. Check the model number for compatibility.
5. If the image was correctly verified, unpack the image in the RAM file system.
6. All switch binary files and libraries run in a secure runtime environment in the RAM file system.

Conclusion

Cisco ACI has many layers of security built into its architecture to prevent various types of threats and to deny access to the fabric by rogue devices or devices that have been improperly manipulated.




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 www.cisco.com/go/offices.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)