



Infrastructure Management

- [Multi-Site Orchestrator Communication Ports](#), on page 1
- [Defining the Overlay TEP for Cisco APIC Sites Using the Cisco APIC GUI](#), on page 2
- [Configuring Infra Prerequisites and Guidelines](#), on page 2
- [Configuring Infra General Settings](#), on page 3
- [Configuring Sites That Contain Remote Leaf Switches](#), on page 3
- [Adding Sites Using Multi-Site Orchestrator GUI](#), on page 5
- [Deleting Sites Using Multi-Site Orchestrator GUI](#), on page 7
- [Cisco ACI CloudSec Encryption](#), on page 7
- [Multi-Site Cross Launch to Cisco APIC](#), on page 18

Multi-Site Orchestrator Communication Ports

There are three types of network communication to or from the Multi-Site Orchestrator cluster:

- Client traffic to the Multi-Site Orchestrator cluster.

Multi-Site Orchestrator uses TCP port 433 ([https](#)) to allow user access via GUI or REST API for creating, managing, and deploying policy configurations.

- REST API traffic from the Multi-Site Orchestrator to the APIC controllers of the ACI fabrics that are part of the Multi-Site domain

Multi-Site Orchestrator uses TCP port 433 for REST API traffic to deploy policies to each site.

- Intra-cluster communication.

All control-plane and data-plane traffic between Cisco ACI Multi-Site Orchestrator nodes (including intra-cluster communication and container overlay network traffic) is encrypted with IPsec's Encapsulating Security Payload (ESP) using IP protocol number 50 to provide security and allow the cluster deployments over a round-trip time distance of up to 150ms. If there is a firewall between any Orchestrator nodes, proper rules must be added to allow this traffic.

If your Multi-Site Orchestrator cluster is deployed directly in VMware ESX without the Application Services Engine, the following ports are used for Docker communications between the cluster nodes:



Note The following TCP and UDP ports are listed for educational perspective only as no traffic is ever sent in clear text across the network leveraging these ports.

- TCP port 2377 for Cluster Management Communications
- TCP and UDP port 7946 for Inter-Manager Communication
- UDP port 4789 for Docker Overlay Network Traffic

Defining the Overlay TEP for Cisco APIC Sites Using the Cisco APIC GUI

Before connecting a Cisco APIC cluster (fabric) in a Cisco ACI Multi-Site topology, you must configure the Overlay Tunnel Endpoint (TEP) in the **Fabric Ext Connection Policy** for each fabric.

The **Create Intrasite/Intersite Profile** panel in the Cisco APIC GUI is used to add connection details for Cisco APIC multipod, remote leaf switches connecting to the Cisco ACI fabric, and APIC sites managed by Cisco ACI Multi-Site Orchestrator. When the Cisco ACI Multi-Site infrastructure has been configured, the Cisco ACI Multi-Site Orchestrator adds the **Intersite Overlay TEP** to this Cisco APIC policy.

To configure the Overlay TEP in the **Fabric Ext Connection Policy** for each Cisco APIC site to be managed by Cisco ACI Multi-Site Orchestrator, perform the following steps:

-
- Step 1** On the menu bar, click **Tenants > infra**.
 - Step 2** On the navigation pane (prior to Cisco APIC, Release 3.1), expand **Networking** and **Protocol Policies**.
 - Step 3** On the navigation pane (in APIC, Release 3.1 and later), expand **Policies** and **Protocol**.
 - Step 4** Right-click **Fabric Ext Connection Policies** and choose **Create Intrasite/Intersite Profile**.
 - Step 5** Click the + symbol on **Pod Connection Profile**.
 - Step 6** Choose the Pod ID from the list.
 - Step 7** Enter the IP address for overlay traffic to this pod.
 - Step 8** Click **Update** and **Submit**.
-

Configuring Infra Prerequisites and Guidelines

The following sections describe the steps necessary to configure the general as well as site-specific fabric Infra settings.

Before you proceed with Infra configuration, you must have configured and added the sites as described in previous sections, which includes:

- Configuring each site's fabric access policies.
- Configuring direct communication and routable subnets for sites with remote leaf switches.

In addition, keep in mind the following:

- Any infrastructure changes such as adding and removing spine switches or spine node ID changes require a Multi-Site fabric connectivity information refresh.

- The Overlay Unicast TEP, Overlay Multicast TEP, and BGP-EVPN Router-IDs IP addresses assigned on the Orchestrator should not be taken from the address space of the original fabric's `Infra` TEP pool or from the `0.x.x.x` range.

Configuring Infra General Settings

This section describes how to configure general Infra settings for all the sites.

-
- Step 1** Log in to the Cisco ACI Multi-Site Orchestrator GUI.
 - Step 2** In the **Main menu**, click **Sites**.
 - Step 3** In the **Sites** view, click **Configure Infra**.
 - Step 4** In the left pane, under **Settings**, click **General Settings**.
 - Step 5** From the **BGP Peering Type** dropdown, choose either `full-mesh` or `route-reflector`.
 - Step 6** In the **Keepalive Interval (Seconds)** field, enter the keep alive interval seconds.
 - Step 7** In the **Hold Interval (Seconds)** field, enter the hold interval seconds.
 - Step 8** In the **Stale Interval (Seconds)** field, enter stale interval seconds.
 - Step 9** Choose whether you want to turn on the **Graceful Helper** option.
 - Step 10** In the **Maximum AS Limit** field, enter the maximum AS limit.
 - Step 11** In the **BGP TTL Between Peers** field, enter the BGP TTL between peers.
-

Configuring Sites That Contain Remote Leaf Switches

Starting with Release 2.1(2), Multi-Site architecture supports APIC sites with Remote Leaf switches. The following sections describe guidelines, limitations, and configuration steps required to allow Multi-Site Orchestrator to manage these sites.

Multi-Site and Remote Leaf Guidelines and Limitations

If you want to add an APIC site with a Remote Leaf to be managed by the Multi-Site Orchestrator, the following restrictions apply:

- You must upgrade your Cisco APIC to Release 4.1(2) or later.
- You must upgrade your Multi-Site Orchestrator to Release 2.1(2) or later.
- Only physical Remote Leaf switches are supported in this release
- Only -EX and -FX or later switches are supported as Remote Leaf switches for use with Multi-Site:
- Remote Leaf is not supported with back-to-back connected sites without IPN switches
- Remote Leaf switches in one site cannot use another site's L3out
- Stretching a bridge domain between one site and a Remote Leaf in another site is not supported

You must also perform the following tasks before the site can be added to and managed by the Multi-Site Orchestrator:

- You must enable Remote Leaf direct communication and configure routable subnets directly in the site's APIC, as described in the following sections.
- You must add the routable IP addresses of Cisco APIC nodes in the DHCP-Relay configuration applied on the interfaces of the Layer 3 routers connecting to the Remote Leaf switches.

The routable IP address of each APIC node is listed in the **Routable IP** field of the **System > Controllers > <controller-name>** screen of the APIC GUI.

Configuring Routable Subnets for Remote Leaf Switches

Before you can add a site that contains one or more Remote Leaf switches to the Multi-Site Orchestrator, you must configure routable subnets for all Remote Leaf switches in that site

-
- Step 1** Log in directly to the site's APIC GUI.
 - Step 2** From the menu bar, select **Fabric > Inventory**.
 - Step 3** In the Navigation pane, click **Pod Fabric Setup Policy**.
 - Step 4** In the main pane, double-click the pod where you want to configure the subnets.
 - Step 5** In the **Routable Subnets** area, click the + sign to add a subnet.
 - Step 6** Enter the IP address and Reserve Address, if necessary, and set the state to Active or Inactive. Then click **Update** to save the subnet.
 - Step 7** Click **Submit** to save the configuration.
-

Enabling Direct Communication for Remote Leaf Switches

Before you can add a site that contains one or more Remote Leaf switches to the Multi-Site Orchestrator, you must configure direct remote leaf communication for that site. Additional information about remote leaf direct communication feature is available in the *Cisco APIC Layer 3 Networking Configuration Guide*. This section outlines the steps and guidelines specific to the integration with Multi-Site.



Note Once you enable Remote Leaf switch direct communication, the switches will function in the new mode only

- Step 1** Log in directly to the site's APIC.
- Step 2** Enable direct traffic forwarding for Remote Leaf switches.
 - a) From the menu bar, navigate to **System > System Settings**.
 - b) From the left side bar, select **Fabric Wide Setting**.
 - c) Check the **Enable Remote Leaf Direct Traffic Forwarding** checkbox.

Note You cannot disable this option after you enable it.

- d) Click **Submit** to save the changes.

Verifying Remote Leaf Configuration

After you enable direct communication for Remote Leaf switches, you can verify the configuration using the following steps.

Step 1 SSH in to the switch.

Step 2 Verify that direct communication is enabled.

In the following output, verify that `rlDirectMode` is set to `yes`:

```
remote-leaf-switch# cat /mit/sys/summary
# System
[...]
remoteNetworkId      : 0
remoteNode           : no
rlOperPodId          : 1
rlRoutableMode       : yes
rlDirectMode        : yes
[...]
```

Step 3 Verify that the remote leaf switches are in complete routable mode and are talking to Cisco APIC's public IP address.

- a) Verify that `rlRoutableMode` is set to `yes`.

```
remote-leaf-switch# moquery -c topSystem | grep rlRoutableMode
rlRoutableMode      : yes
```

- b) Verify that you can ping the Cisco APIC routable IP address from the remote leaf switch.

```
remote-leaf-switch# iping -v overlay-1 110.0.0.225

PING 110.0.0.225 (110.0.0.225) from 193.0.3.20: 56 data bytes

64 bytes from 110.0.0.225: icmp_seq=0 ttl=61 time=0.401 ms
```

- c) Verify that `dhcpRespMo` in the remote leaf switch is set to the APIC's routable IP address.

```
remote-leaf-switch# moquery -c dhcpResp

serverId           : 110.0.0.225
siAddr             : 110.0.0.225
status              :
subnetMask          : 255.255.255.255
yiAddr              : 191.2.0.72
```

Adding Sites Using Multi-Site Orchestrator GUI

This section describes how to add sites using the Cisco ACI Multi-Site Orchestrator GUI.

Before you begin

You must have completed the site-specific configurations in each site's APIC, as described in previous sections in this chapter.

Step 1 Log in to the Multi-Site GUI, in the **Main menu**, click **Sites**.

If you are logging in for the first time, log in as the **admin** user with the default password **We1come2msc!**, you will then be prompted to change that default password. The new password requirements are:

- At least 12 characters
- At least 1 letter
- At least 1 number
- At least 1 special character apart from * and space

Step 2 In the **Main menu**, click **Sites**.

Step 3 In the **Sites** view, click **Add Site**.

Step 4 In the **Add Site** page, provide the site's details.

- a) In the **Name** field, enter the site name.
- b) In the **Labels** field, choose or create a label.

You can choose to provide multiple labels for the site.

- c) In the **APIC Controller URL** field, enter the Cisco APIC URL.

For the APIC URL, you can use the `http` or `https` protocol and the IP address or the DNS hostname, such as `ashttps://<ip-address>` or `https://<dns-hostname>`.

- d) If you have a cluster of APICs in the fabric, click **+APIC Controller URL** and provide the additional URLs.
- e) In the **Username** field, enter the admin user's username for the site's APIC.
- f) In the **Password** field, enter the user's password.
- g) You can turn on the **Specify Login Domain for Site** switch, if you want to specify a domain to be used for authenticating the user you provided.

If you turn on this option, enter the domain name in the **Domain Name** field.

- h) In the **APIC Site ID** field, enter a unique site ID.

The site ID must be a unique identifier of the Cisco APIC site, ranged between 1 and 127. Once specified, the site ID cannot be changed without factory resetting Cisco APIC.

Step 5 Click **Save** to add the site.

Step 6 If prompted, confirm proxy configuration update.

If you have configured the Orchestrator to use a proxy server and are adding an on-premises site that is not already part of the "no proxy" list, the Orchestrator will inform you of the proxy settings update.

For additional information on proxy configuration, see the "Administrative Operations" chapter in *Cisco ACI Multi-Site Configuration Guide*.

Step 7 Repeat these steps to add any additional sites.

Deleting Sites Using Multi-Site Orchestrator GUI

This section describes how to delete sites using the Multi-Site GUI.

-
- Step 1** Log in to the Multi-Site GUI.
 - Step 2** Ensure you unbind the site from any Schema's before trying to delete the site.
 - Step 3** In the **Main menu**, click **Sites**.
 - Step 4** In the **Sites List** page, hover over the site you want to delete and choose **Action > Delete**.
 - Step 5** Click **YES**.
-

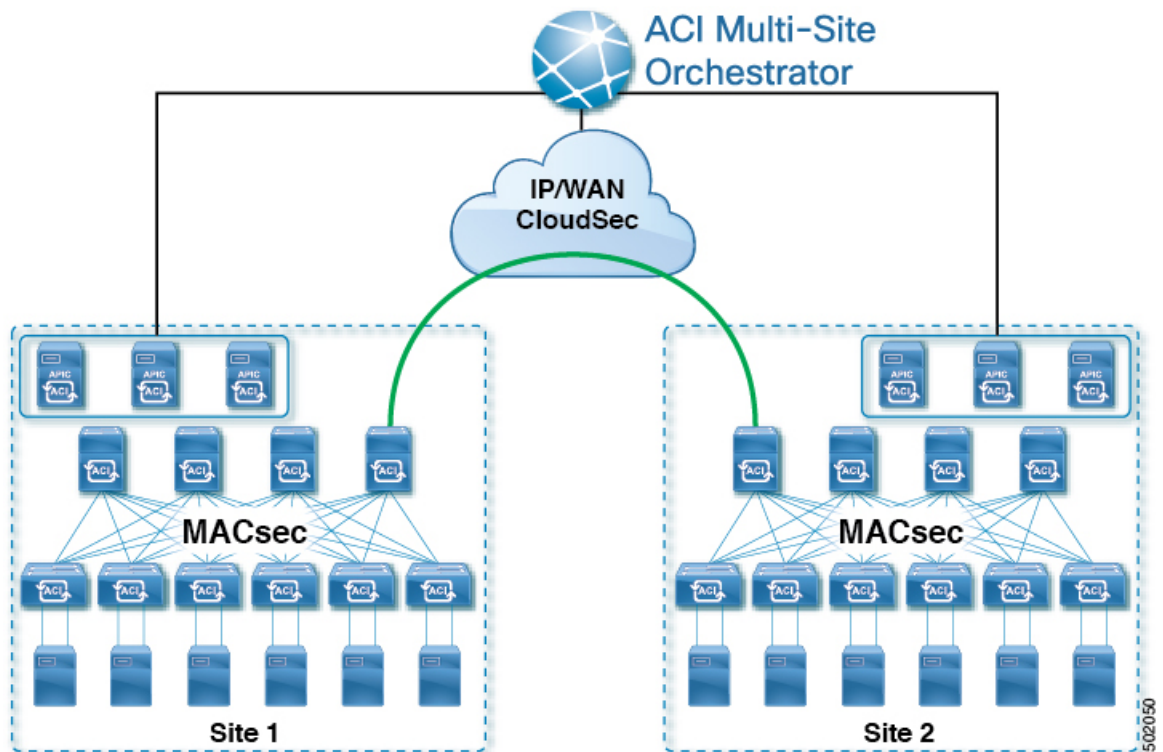
Cisco ACI CloudSec Encryption

As most Cisco ACI deployments are adopting the Cisco ACI Multi-Site architecture to address disaster recovery and scale, the current security implementation using MACsec encryption within local site is becoming insufficient to guarantee data security and integrity across multiple sites connected by insecure external IP networks interconnecting separate fabrics. Cisco ACI Multi-Site Orchestrator Release 2.0(1) introduces the CloudSec Encryption feature designed to provide inter-site encryption of traffic.

Cisco ACI Multi-Site topology uses three tunnel end-point (TEP) IP addresses to provide connectivity between sites. These TEP addresses are configured by the admin on Cisco ACI Multi-Site Orchestrator and pushed down to each site's Cisco APIC, which in turn configures them on the spine switches. These three addresses are used to determine when traffic is destined for a remote site, in which case an encrypted CloudSec tunnel is created between the two spine switches that provide physical connectivity between the two sites through the Inter-Site Network (ISN).

The following figure illustrates the overall encryption approach that combines MACsec for local site traffic and CloudSec for inter-site traffic encryption.

Figure 1: CloudSec Encryption



Hardware Requirement

The current release supports CloudSec Encryption on the following hardware::

- Cisco Nexus 9364 and 9332 spine switches
- FX line cards in Cisco Nexus 9500 switches

CloudSec Encryption Terminology

CloudSec Encryption feature provides a secure upstream symmetric key allocation and distribution method for initial key and rekey requirements between sites. The following terminology is used in this chapter:

- **Upstream device** – The device that adds the CloudSec Encryption header and does the encryption of the VxLAN packet payload on transmission to a remote site using a locally generated symmetric cryptography key.
- **Downstream device** – The device that interprets the CloudSec Encryption header and does the decryption of the VxLAN packet payload on reception using the cryptography key generated by the remote site.
- **Upstream site** – The datacenter fabric that originates the encrypted VxLAN packets.
- **Downstream site** – The datacenter fabric that receives the encrypted packets and decrypts them.
- **TX Key** – The cryptography key used to encrypt the clear VxLAN packet payload. In ACI only one TX key can be active per remote site.

- **RX Key** – The cryptography key used to decrypt the encrypted VxLAN packet payload. In ACI two RX key can be active per remote site.
- **Symmetric Keys** – When the same cryptography key is used to encrypt (TX Key) and decrypt (RX Key) a packet stream by the upstream and downstream devices respectively.
- **Rekey** – The process initiated by the upstream site to replace its old key with a newer key for all downstream sites after the old key expires.
- **Secure Channel Identifier (SCI)** – A 64-bit identifier that represents a security association between the sites. It is transmitted in encrypted packet in CloudSec header and is used to derive the RX key on the downstream device for packet decryption.
- **Association Number (AN)** – A 2-bit number (0, 1, 2, 3) that is sent in the CloudSec header of the encrypted packet and is used to derive the key at the downstream device in conjunction with the SCI for decryption. This allows multiple keys to be active at the downstream device to handle out of order packet arrivals with different keys from the same upstream device following a rekey operation.


In ACI only two association number values (0 and 1) are used for the two active RX keys and only one association number value (0 or 1) is used for the TX Key at any point in time.
- **Pre-shared key (PSK)** – One or more keys must be configured in the Cisco APIC GUI to be used as a random seed for generating the CloudSec TX and RX keys. If multiple PSK are configured, each rekey process will use the next PSK in order of their indexes; if no higher index PSK is available, a PSK with the lowest index will be used. Each PSK must be a hexadecimal string 64 characters long. Cisco APIC supports up to 256 pre-shared keys.

CloudSec Encryption and Decryption Handling

In order to provide a fully integrated, simple, and cost-effective solution that addresses both, data security and integrity, starting with Release 2.0(1), Cisco ACI Multi-Site provides a CloudSec Encryption feature that allows for complete source-to-destination packet encryption between Multi-Site fabrics.

The following figure shows packet diagram before and after CloudSec encapsulation, followed by descriptions of the encryption and decryption processes:

Figure 2: CloudSec Packet



Packet Encryption

The following is a high level overview of how CloudSec handles outgoing traffic packets:

- The packets are filtered using the outer IP header and Layer-4 destination port information and matching packets are marked for encryption.
- The offset to use for encryption is calculated according to the fields of the packet. For example, the offset may vary based on whether there is a 802.1q VLAN or if the packet is an IPv4 or IPv6 packet.
- The encryption keys are programmed in the hardware tables and are looked up from the table using the packet IP header.

Once the packet is marked for encryption, the encryption key is loaded, and the offset from the beginning of the packet where to start the encryption is known, the following additional steps are taken:

- The UDP destination port number is copied from the UDP header into a CloudSec field for recovery when the packet is decrypted.
- The UDP destination port number is overwritten with a Cisco proprietary Layer-4 port number (Port 9999) indicating that it is a CloudSec packet.
- The UDP length field is updated to reflect the additional bytes that are being added.
- The CloudSec header is inserted directly after the UDP header.
- The Integrity Check Value (ICV) is inserted at the end of the packet, between the payload and the CRC.
- The ICV requires construction of a 128-bit initialization vector. For CloudSec, any use of the source MAC address for ICV purposes is replaced by a programmable value per SCI.
- CRC is updated to reflect the change in the contents of the packet.

Packet Decryption


The way CloudSec handles incoming packets is symmetric to the outgoing packets algorithm described above:

- If the received packet is a CloudSec packet, it is decrypted and the ICV is verified.
If ICV verification passed, the extra fields are removed, the UDP destination port number is moved from the CloudSec header to the UDP header, the CRC is updated, and the packet is forwarded to destination after decryption and CloudSec header removal. Otherwise the packet is dropped.
- If the key store returns two or more possible decryption keys, the Association Number (AN) field of the CloudSec header is used to select which key to use.
- If the packet is not a CloudSec packet, the packet is left unchanged.

CloudSec Encryption Key Allocation and Distribution

Initial Key Configuration

Figure 3: CloudSec Key Distribution



The following is a high level overview of the CloudSec encryption key initial allocation and distribution process illustrated by the figure above:

- The upstream site's Cisco APIC generates a local symmetric key intended to be used for data encryption of VxLAN packets transmitted from its site. The same key that is used by the upstream site for encryption is used for decryption of the packets on the downstream remote receiving sites.

Every site is an upstream site for the traffic it transmits to other sites. If multiple sites exist, each site generates its own site-to-site key and use that key for encryption before transmitting to the remote site.

- The generated symmetric key is pushed to the Cisco ACI Multi-Site Orchestrator (MSO) by the upstream site's Cisco APIC for distribution to downstream remote sites.
- The MSO acts as a message broker and collects the generated symmetric key from the upstream site's Cisco APIC, then distributes it to downstream remote sites' Cisco APICs.
- Each downstream site's Cisco APIC configures the received key as RX key on the local spine switches which are intended to receive the traffic from the upstream site that generated the key.
- Each downstream site's Cisco APIC also collects the deployment status of the RX Key from the local spine switches and then pushes it to the MSO.
- The MSO relays the key deployment status from all downstream remote sites back to the upstream site's Cisco APIC.
- The upstream site's Cisco APIC checks if the key deployment status received from all downstream remote sites is successful.
 - If the deployment status received from a downstream device is successful, the upstream site deploys the local symmetric key as its TX key on the spine switches to enable encryption of the VxLAN packets that are sent to the downstream site.
 - If the deployment status received from a downstream device is failed, a fault is raised on the Cisco APIC site where it failed and it is handled based on the "secure mode" setting configured on the MSO. In "must secure" mode the packets are dropped and in the "should secure" mode the packets are sent clear (unencrypted) to the destination site.



Note In current release, the mode is always set to "should secure".

Rekey Process

Each generated TX/RX key expires after a set amount of time, by default key expiry time is set to 15 minutes. When the initial set of TX/RX keys expires, a rekey process takes place.

The same general key allocation and distribution flow applies for the rekey process. The rekey process follows the "make before break" rule, in other words all the RX keys on the downstream sites are deployed before the new TX key is deployed on the upstream site. To achieve that, the upstream site will wait for the new RX key deployment status from the downstream sites before it configures the new TX key on the local upstream site's devices.

If any downstream site reports a failure status in deploying the new RX key, the rekey process will be terminated and the old key will remain active. The downstream sites will also keep the old and the new RX keys deployed after the new key deployment is finished for some duration to ensure that out of order packet deliveries with either key can be properly decrypted.



Note Special precautions must be taken in regards to rekey process during spine switch maintenance, see [Rekey Process During Spine Switch Maintenance, on page 17](#) for details.

Rekey Process Failure

In case of any downstream site failing to deploy the new encryption key generated by the rekey process, the new key is discarded and the upstream device will continue to use the previous valid key as TX key. This approach keeps the upstream sites from having to maintain multiple TX keys per set of downstream sites. However, this approach may also result in the rekey process being delayed if the rekey deployment failures continue to occur with any one of the downstream sites. It is expected that the Multi-Site administrator will take action to fix the issue of the key deployment failure for the rekey to succeed.

Cisco APIC's Role in Key Management

The Cisco APIC is responsible for key allocation (both, initial key and rekey distribution), collection of the key deployment status messages from the spine switches, and notification of the Cisco ACI Multi-Site Orchestrator about each key's status for distribution to other sites.

Cisco ACI Multi-Site Orchestrator's Role in Key Management

The Cisco ACI Multi-Site Orchestrator is responsible for collecting the TX keys (both, initial key and subsequent rekeys) from the upstream site and distributing it to all downstream sites for deployment as RX keys. The MSO also collects the RX key deployment status information from the downstream sites and notifies the upstream site in order for it to update the TX key on successful RX key deployment status.

Upstream Model

In contrast to other technologies, such as MPLS, that use downstream key allocation, CloudSec's upstream model provides the following advantages:

- The model is simple and operationally easier to deploy in the networks.
- The model is preferred for Cisco ACI Multi-Site use cases.
- It provides advantages for multicast traffic as it can use the same key and CloudSec header for each copy of the replicated packet transmitted to multiple destination sites. In downstream model each copy would have to use a different security key for each site during encryption.
- It provides easier troubleshooting in case of failures and better traceability of packets from the source to destination consistently for both, unicast and multicast replicated packets.

CloudSec Supported Hardware and Port Ranges

The following table provides the hardware platforms and the port ranges that are capable of CloudSec encryption:

Hardware Platform	Port Range
N9K-C9364C spine switches	Ports 49-64
N9K-C9332C spine switches	Ports 25-32
N9K-X9736C-FX line cards	Ports 29-36

If CloudSec is enabled for a site, but the encryption is not supported by the ports, a fault is raised with `unsupported-interface` error message.

CloudSec encryption's packet encapsulation is supported if Cisco QSFP-to-SFP Adapters (QSA), such as CVR-QSFP-SFP10G, is used with a supported optic. The full list of supported optics is available from the following link: <https://www.cisco.com/c/en/us/support/interfaces-modules/transceiver-modules/products-device-support-tables-list.html>.

Configuring Cisco APIC for CloudSec Encryption

You must configure one or more Pre-Shared Keys (PSK) to be used by the Cisco APIC for generating the CloudSec encryption and decryption keys. The PSK are used as a random seed during the re-key process. If multiple PSK are configured, each re-key process will use the next PSK in order of their indexes; if no higher index PSK is available, a PSK with the lowest index will be used.

Because PSK is used as a seed for encryption key generation, configuring multiple PSK provides additional security by lowering the over-time vulnerability of the generated encryption keys.



Note If no pre-shared key is configured on the Cisco APIC, CloudSec will not be enabled for that site. In that case, turning on CloudSec setting in Cisco ACI Multi-Site will raise a fault.

If at any time you wish to refresh a previously added PSK with a new one, simply repeat the procedure as if you were adding a new key, but specify an existing index.

You can configure one or more pre-shared keys in one of three ways:

- Using the Cisco APIC GUI, as described in [Configuring Cisco APIC for CloudSec Encryption Using GUI, on page 14](#)
- Using the Cisco APIC NX-OS Style CLI, as described in [Configuring Cisco APIC for CloudSec Encryption Using NX-OS Style CLI, on page 15](#)
- Using the Cisco APIC REST API, as described in [Configuring Cisco APIC for CloudSec Encryption Using REST API, on page 15](#)

Configuring Cisco APIC for CloudSec Encryption Using GUI

This section describes how to configure one or more pre-shared keys (PSK) using the Cisco APIC GUI.

Step 1 Log in to APIC.

Step 2 Navigate to **Tenants > infra > Policies > CloudSec Encryption**

Step 3 Specify the **SA Key Expiry Time**.

This option specifies how long each key is valid (in minutes). Each generated TX/RX key expires after the specified amount of time triggering a re-key process. The expiration time can be between 5 and 1440 minutes.

Step 4 Click the + icon in the **Pre-Shared Keys** table.

Step 5 Specify the **Index** of the pre-shared key you are adding and then the **Pre-Shared Key** itself.

The **Index** field specifies the order in which the pre-shared keys are used. After the last (highest index) key is used, the process will continue with the first (lowest index) key. Cisco APIC supports up to 256 pre-shared keys, so the PSK index value must be between 1 and 256.

Each **Pre-Shared Key** must be a hexadecimal string 64 characters long.

Configuring Cisco APIC for CloudSec Encryption Using NX-OS Style CLI

This section describes how to configure one or more pre-shared keys (PSK) using the Cisco APIC NX-OS Style CLI.

Step 1 Log in to the Cisco APIC NX-OS style CLI.

Step 2 Enter configuration mode.

Example:

```
apicl# configure
apicl (config)#
```

Step 3 Enter configuration mode for the default CloudSec profile.

Example:

```
apicl (config)# template cloudsec default
apicl (config-cloudsec)#
```

Step 4 Specify the Pre-Shared Keys (PSK) expiration time.

This option specifies how long each key is valid (in minutes). Each generated TX/RX key expires after the specified amount of time triggering a re-key process. The expiration time can be between 5 and 1440 minutes.

Example:

```
apicl (config-cloudsec)# sakexpirytime <duration>
```

Step 5 Specify one or more Pre-Shared Keys.

In the following command, specify the index of the PSK you're configuring and the PSK string itself.

Example:

```
apicl (config-cloudsec)# pskindex <psk-index>
apicl (config-cloudsec)# pskstring <psk-string>
```

The <psk-index> parameter specifies the order in which the pre-shared keys are used. After the last (highest index) key is used, the process will continue with the first (lowest index) key. Cisco APIC supports up to 256 pre-shared keys, so the PSK index value must be between 1 and 256.

The <psk-string> parameter specifies the actual PSK, which must be a hexadecimal string 64 characters long.

Step 6 (Optional) View the current PSK configuration.

You can view how many PSK are currently configured and their duration using the following command:

Example:

```
apicl (config-cloudsec)# show cloudsec summary
```

Configuring Cisco APIC for CloudSec Encryption Using REST API

This section describes how to configure one or more pre-shared keys (PSK) using the Cisco APIC REST API.

Configure PSK expiration time, index, and string.

In the following XML POST, replace:

- The value of **sakExpiryTime** with the expiration time of each PSK.

This **sakExpiryTime** parameter specifies how long each key is valid (in minutes). Each generated TX/RX key expires after the specified amount of time triggering a re-key process. The expiration time can be between 5 and 1440 minutes.

- The value of **index** with the index of the PSK you're configuring.

The **index** parameter specifies the order in which the pre-shared keys are used. After the last (highest index) key is used, the process will continue with the first (lowest index) key. Cisco APIC supports up to 256 pre-shared keys, so the PSK index value must be between 1 and 256.

- The value of **pskString** with the index of the PSK you're configuring.

The **pskString** parameter specifies the actual PSK, which must be a hexadecimal string 64 characters long.

Example:

```
<fvTenant annotation="" descr="" dn="uni/tn-infra" name="infra" nameAlias="" ownerKey="" ownerTag="">
  <cloudsecIfPol descr="cloudsecifp" name="default" sakExpiryTime="
10
" stopRekey= "false" status="" >
  <cloudsecPreSharedKey index="
1
" pskString="
123456781234567812345678123456781234567812345678123456781234567812345678
" status=""/>
  </cloudsecIfPol>
</fvTenant>
```

Enabling CloudSec Encryption Using Cisco ACI Multi-Site Orchestrator GUI

The CloudSec encryption can be enabled or disabled for each site individually. However, the communications between two sites will be encrypted only if the feature is enabled on both sites.

Before you begin

Before you enable the CloudSec encryption between two or more sites, you must have completed the following tasks:

- Installed and configured the Cisco APIC clusters in multiple sites, as described in [Cisco APIC Installation, Upgrade, and Downgrade Guide](#)
- Installed and configured Cisco ACI Multi-Site Orchestrator, as described in [Cisco ACI Multi-Site Orchestrator Installation and Upgrade Guide](#) .
- Added each Cisco APIC site to the Cisco ACI Multi-Site Orchestrator, as described in [Cisco ACI Multi-Site Configuration Guide](#) .

-
- Step 1** Log in to the Cisco ACI Multi-Site Orchestrator.
- Step 2** From the left-hand sidebar, select the **Sites** view.
- Step 3** Click on the **Configure Infra** button in the top right of the main window.
- Step 4** From the left-hand sidebar, select the site for which you want to change the CloudSec configuration.
- Step 5** In the right-hand sidebar, toggle the **CloudSec Encryption** setting to enable or disable the CloudSec Encryption feature for the site.
-

Rekey Process During Spine Switch Maintenance

The following is a summary of the CloudSec rekey process during typical maintenance scenarios for the spine switches where the feature is enabled:

- **Normal Decommissioning** – CloudSec rekey process stops automatically whenever a CloudSec-enabled spine switch is decommissioned. Rekey process will not start again until the decommissioned node is commissioned back or the decommissioned node ID is removed from the Cisco APIC
- **Spine Switch Software Upgrade** – CloudSec rekey process stops automatically if a spine switch is reloaded due to software upgrade. Rekey process will resume after the spine switch comes out of reload.
- **Maintenance (GIR mode)** – CloudSec rekey process must be manually stopped using the instructions provided in [Disabling and Re-Enabling Re-Key Process Using NX-OS Style CLI, on page 17](#). Rekey can be enabled back only after the node is ready to forward traffic again.
- **Decommissioning and Removal from Cisco APIC** – CloudSec rekey process must be manually stopped using the instructions provided in [Disabling and Re-Enabling Re-Key Process Using NX-OS Style CLI, on page 17](#). Rekey can be enabled back only after the node is removed from Cisco APIC.

Disabling and Re-Enabling Re-Key Process Using NX-OS Style CLI

It is possible to manually stop and restart the re-key process. You may be required to manually control the re-key process in certain situations, such as switch decommissioning and maintenance. This section describes how to toggle the setting using Cisco APIC NX-OS Style CLI.

-
- Step 1** Log in to the Cisco APIC NX-OS style CLI.
- Step 2** Enter configuration mode.
- Example:**
- ```
apicl# configure
apicl(config)#
```
- Step 3** Enter configuration mode for the default CloudSec profile.
- Example:**
- ```
apicl(config)# template cloudsec default
apicl(config-cloudsec)#
```
- Step 4** Stop or restart the re-key process.
- To stop the re-key process:

Example:

```
apic1(config-cloudsec)# stoprekey yes
```

To restart the re-key process:

Example:

```
apic1(config-cloudsec)# stoprekey no
```

Disabling and Re-Enabling Re-Key Process Using REST API

It is possible to manually stop and restart the re-key process. You may be required to manually control the re-key process in certain situations, such as switch decommissioning and maintenance. This section describes how to toggle the setting using Cisco APIC REST API.

Step 1 You can disable the rekey process using the following XML message.

Example:

```
<fvTenant annotation="" descr="" dn="uni/tn-infra" name="infra" nameAlias="" ownerKey="" ownerTag="">
  <cloudsecIfPol descr="cloudsecifp" name="default" sakExpiryTime="10" stopRekey= "
true
" status="" />
</fvTenant>
```

Step 2 You can enable the rekey process using the following XML message.

Example:

```
<fvTenant annotation="" descr="" dn="uni/tn-infra" name="infra" nameAlias="" ownerKey="" ownerTag="">
  <cloudsecIfPol descr="cloudsecifp" name="default" sakExpiryTime="10" stopRekey= "
false
" status="" />
</fvTenant>
```

Multi-Site Cross Launch to Cisco APIC

Multi-Site currently supports the basic parameters to choose when creating a Tenant and setting up a site. Multi-Site supports most of the Tenant policies, but in addition to that you can configure some advanced parameters.

Use the Multi-Site GUI to manage the basic properties to configure. If you want to configure advanced properties, the capability to cross launch into Cisco APIC GUI directly from the Multi-Site GUI is provided. You can also configure the additional properties directly in Cisco APIC.

There are three different access points in Multi-Site GUI from where you can cross launch into APIC. From these access points in Multi-Site, you can open a new browser tab with access into Cisco APIC. You will log in to Cisco APIC at that point for the first time, and the associated screen is displayed in the Cisco APIC GUI.

Cross-Launch to Cisco APIC from Sites

Before you begin

- At least one site must be configured in Multi-Site.
- The site must contain at least one tenant with entities such as VRF and bridge domain configured.

Step 1 From the left-hand sidebar, open the **Sites** view.

Step 2 From the **Sites** list, hover over the name of the appropriate site, click the **Actions** icon at the end of the row, and choose **Open in APIC User Interface** to access the Cisco APIC GUI.

The APIC GUI login screen is displayed for logging in with APIC GUI credentials.

Cross-Launch to Cisco APIC from Schemas

Before you begin

- At least one site based on a template must be configured in Multi-Site.
- The site must contain at least one tenant with entities such as VRF and bridge domain configured.

Step 1 From the left-hand sidebar, open the **Schemas** view.

Step 2 From the **Schemas** list, click the appropriate *<schema-name>* .

Step 3 From the left-hand sidebar **Sites** list, hover over the name of the appropriate site, click the **Actions** icon at the end of the row, and choose **Open in APIC User Interface** to access the Cisco APIC GUI.

The APIC GUI login screen is displayed for logging in with APIC GUI credentials.

Cross-Launch to Cisco APIC from the Property Pane

Before you begin

- At least one site must be configured in Multi-Site.
- The site must contain at least one tenant with entities such as VRF and bridge domain configured.

Step 1 From the left-hand sidebar, open the **Schemas** view.

Step 2 From the **Schemas** list, click the appropriate *<schema-name>* .

Step 3 From the left-hand sidebar **Sites** list, choose the appropriate site.

Step 4 In the **Canvas**, choose the name of a specific entity.

For example, choose an available VRF, Contract, Bridge Domain, or another entity as appropriate.

The details for the specific entity are displayed in the **Property Pane** on the right.

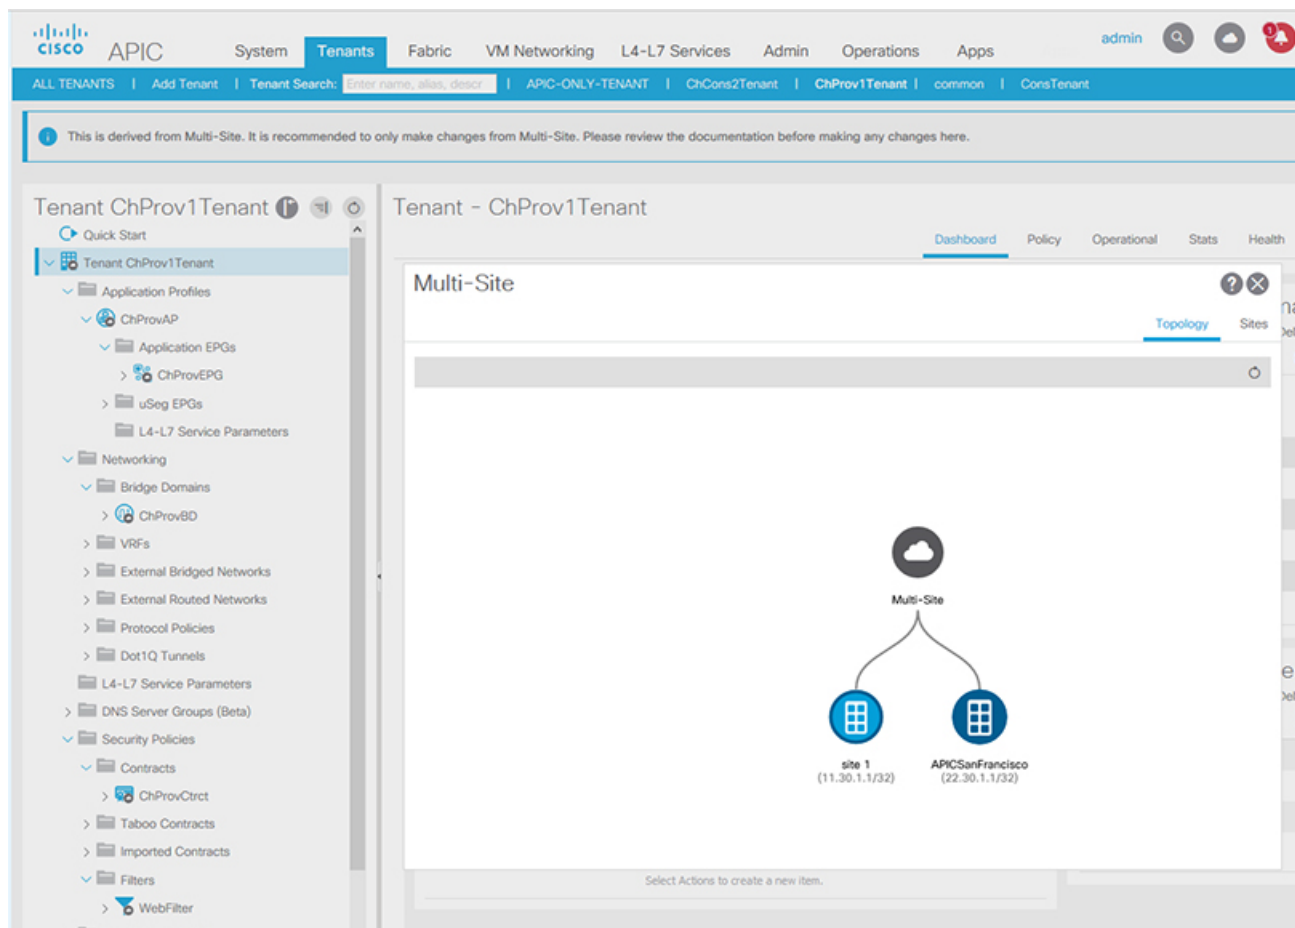
Step 5 In the top right of the **Property Pane**, click the **Open in APIC User Interface** icon to access the Cisco APIC GUI.

The APIC GUI login screen is displayed for logging in with APIC GUI credentials.

Viewing Cisco ACI Multi-Site-Managed Objects Using the Cisco APIC GUI

When an APIC cluster is managed by Multi-Site, cloud icons indicate the relationships with other sites.

Figure 4: Viewing Multi-Site-Managed Objects Using the APIC GUI



Before you begin

The APIC cluster/site must be set up to be managed by Cisco ACI Multi-Site.

Step 1 To view the relationship of the APIC site with other sites, click the cloud icon at the upper right, next to the settings icons. In the diagram, hover over the light blue site icon to see the local site details, and hover over the dark blue icon to see the remote site details.

In the image, T1 and its Application Profile, EPG, BD, VRF, and contracts are marked with cloud icons. This indicates that they are managed by Multi-Site. We recommend that you only make changes to these objects in the Multi-Site GUI.

Step 2 To view the localized or stretched usage of a VRF, bridge domain, or other objects, where there is a **Show Usage** button on the information page, perform the following steps; for example for Bridge Domain and VRF:

- a) On the menu bar, click **Tenants** and double-click on a tenant that is managed by Multi-Site.
- b) Click **Networking > Bridge Domains > BD-name** or **Networking > VRFs > vrf-name**.

Step 3 Click **Show Usage**.

Here you can view the nodes or policies using the object.

Note It is recommended to make changes to managed policies only in the Multi-Site GUI.

Step 4 To set the scope of deployment notification settings for this BD or VRF, click **Change Deployment Settings**. You can enable warnings to be sent for all deletions and modifications of the object on the **Policy** tab.

Step 5 To enable or disable Global warnings, check or uncheck the **(Global) Show Deployment Warning on Delete/Modify** check box.

Step 6 To enable or disable Local warnings, choose **Yes** or **No** on the **(Local) Show Deployment Warning on Delete/Modify** field.

Step 7 To view any past warnings, click the **History** tab **Events** or **Audit Logs**.
