Configuring Infra

This chapter contains the following sections:

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 described in the Refreshing Site Connectivity Information, on page 2 as part of the general Infra configuration procedures.
- 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`.

The `route-reflector` option is effective only when all sites are part of the same BGP Autonomous System (AS).

Step 6 In the **Keepalive Interval (Seconds)** field, enter the keep alive interval seconds.

We recommend keeping the default value.

Step 7 In the **Hold Interval (Seconds)** field, enter the hold interval seconds.

We recommend keeping the default value.

Step 8 In the **Stale Interval (Seconds)** field, enter stale interval seconds.

We recommend keeping the default value.

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.

Refreshing Site Connectivity Information

Any infrastructure changes, such as adding and removing spines or changing spine node IDs, require a Multi-Site fabric connectivity site refresh. This section describes how to pull up-to-date connectivity information directly from each site's APIC.

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

Step 2 In the **Main menu**, select **Infrastructure** > **Infra Configuration**.

Step 3 In the top right of the main **Infra Configuration** view, click the **Configure Infra** button.

Step 4 In the left pane, under **Sites**, select a specific site.

Step 5 In the main window, click the **Reload Site Data** button to pull fabric information from the APIC.

Step 6 (Optional) In the **Confirmation** dialog, check the box if you want to remove configuration for decommissioned spine switch nodes.

If you choose to enable this checkbox, all configuration info for any currently decommissioned spine switches will be removed from the database.

Step 7 Finally, click **Yes** to confirm and load the connectivity information.

This will discover any new or removed spines and all site-related fabric connectivity will be re-imported from the APIC.

Configuring Infra Site-Specific Settings

This section describes how to configure site-specific Infra settings for each site.

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 **Sites**, select a specific site.

Step 5 In the right **<*Site*> Settings** pane, enable the site by setting the **ACI Multi-Site** knob to `On`.

Step 6 (Optional) Turn on CloudSec encryption for the site.

CloudSec Encryption provides inter-site traffic encryption. The "Infrastructure Management" chapter in the *Cisco ACI Multi-Site Configuration Guide* covers this feature in detail.

Step 7 Specify the **Overlay Multicast TEP**.

This address is used for the inter-site L2 BUM and L3 multicast traffic. This IP address is deployed on all spine switches that are part of the same fabric, regardless of whether it is a single Pod or Multi-Pod fabric.

Step 8 Specify the **BGP Autonomous System Number**.

Step 9 Specify the **BGP Password**.

Step 10 Specify the **OSPF Area ID**.

When configuring the Multi-Site infra OSPF details, we recommend that you use OSPF Area `0`. If you use an Area ID other than `0`, in the next step configure it as a `regular` OSPF area type and not a `stub` area type.

Step 11 Select the **OSPF Area Type** from the dropdown menu.

The OSPF area type can be one of the following:

- `nssa`
- `regular`
- `stub`

Step 12 Select the external routed domain from the dropdown menu.

Choose an external router domain that you have created in the APIC GUI.

Step 13 Configure OSPF settings for the site.

You can either click an existing policy (for example, `msc-ospf-policy-default`) to modify it or click **+Add Policy** to add a new OSPF policy. Then in the **Add/Update Policy** window, specify the following:

- In the **Policy Name** field, enter the policy name.
- In the **Network Type** field, choose either `broadcast`, `point-to-point`, or `unspecified`.

  The default is `broadcast`.
- In the **Priority** field, enter the priority number.

  The default is `1`.

- In the **Cost of Interface** field, enter the cost of interface.

  The default is `0`.

- From the **Interface Controls** dropdown menu, choose one of the following:
  - **advertise-subnet**
  - **bfd**
  - **mtu-ignore**
  - **passive-participation**
- In the **Hello Interval (Seconds)** field, enter the hello interval in seconds.

  The default is `10`.

- In the **Dead Interval (Seconds)** field, enter the dead interval in seconds.

  The default is `40`.

- In the **Retransmit Interval (Seconds)** field, enter the retransmit interval in seconds.

  The default is `5`.

- In the **Transmit Delay (Seconds)** field, enter the transmit delay in seconds.

  The default is `1`.

Configuring Infra: Cloud Site Settings

This section describes how to configure site-specific Infra settings for Cloud APIC sites.

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

Step 2 In the **Main menu**, select **Infrastructure** > **Infra Configuration**.

Step 3 In the top right of the main pane, click **Configure Infra**.

Step 4 In the left pane, under **Sites**, select a specific cloud site.

Most of the information in the cloud site area is automatically populated and cannot be changed, with the exception of the BGP password field.

Step 5 In the right **<*Site*> Settings** pane, enable the **ACI Multi-Site** knob to manage the site from the Orchestrator.

Step 6 Specify the **BGP Password**.

Configuring Infra: Pod Settings

This section describes how to configure pod-specific settings in each site.

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 **Sites**, select a specific site.

Step 5 In the main window, select a pod.

Step 6 In the right **POD Properties** pane, add the Overlay Unicast TEP for the POD.

This IP address is deployed on all spine switches that are part of the same pod and used for intersite known unicast traffic.

Step 7 Click **+Add TEP Pool** to add a routable TEP pool.

The routable TEP pools are used for public IP addresses for inter-site connectivity.

Step 8 Repeat the procedure for every pod in the site.

Configuring Infra: Spine Switches

This section describes how to configure spine switches in each site for Cisco ACI Multi-Site.

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 **Sites**, select a specific site.

Step 5 In the main window, select a spine switch within a pod.

Step 6 In the right **<*Spine*> Settings** pane, click **+Add Port**.

Step 7 In the **Add Port** window, enter the following information:

- In the **Ethernet Port ID** field, enter the port ID, for example `1/29`.
- In the **IP Address** field, enter the IP address/netmask.

  The Orchestrator creates a sub-interface with VLAN 4 with the specified IP ADDRESS under the specified PORT.
- In the **MTU** field, enter the MTU. You can specify either `inherit` or a value between `576` and `9000`.

  MTU of the spine port should match MTU on IPN side.
- In the **OSPF Policy** field, choose the OSPF policy for the switch that you have configured in Configuring Infra Site-Specific Settings, on page 3.

  OSPF settings in the OSPF policy you choose should match on IPN side.
- For **OSPF Authentication**, you can pick either `none` or one of the following:
  - `MD5`
  - `Simple`

Step 8 Enable **BGP Peering** knob.

In a single Pod fabric with more than two spine switches, BGP peering should only be enabled on a pair (for redundancy) of spine switches called **BGP Speakers**. All other spine switches should have BGP peering disabled and will function as **BGP Forwarders**.

In a Multi-Pod fabric BGP peering should only be enabled on a couple of BGP speaker spine switches, each deployed in a different Pod. All other spines switches should have BGP peering disabled and function as BGP forwarders.

Step 9 In the **BGP-EVPN Router-ID** field, provide the IP address used for BGP-eVPN session between sites.

Step 10 Repeat the procedure for every spine switch.

Deploying Infra Configuration

This section describes how to deploy the Infra configuration to each APIC site.

In the top right of the main pane, choose the appropriate **Deploy** option to deploy the configuration.

If you are configuring only on-premises or only cloud sites, simply click **Deploy** to deploy the Infra configuration.

However, if you have both, on-premises and cloud site, the following two additional options become available:

- **Deploy & Download IPN Device config files:** Pushes the configuration to both the on-premises APIC site and the cloud site and enables the end-to-end interconnect between the on-premises and the cloud sites.

  In addition, this option downloads a zip file that contains configuration information that you will use to enable connectivity between the Cisco Cloud Services Router (CSR) deployed in your cloud sites and the on-premises IPsec termination device. A followup screen appears that allows you to select all or some of the configuration files to download.

- **Download IPN Device config files only:** Downloads a zip file that contains configuration information that you will use to enable connectivity between the Cisco Cloud Services Router (CSR) without deploying the configuration.

Enabling Connectivity Between On-Premises and Cloud Sites

If you have only on-premises or only cloud sites, you can skip this section.

This section describes how to enable connectivity between on-premises APIC sites and Cloud APIC sites.

By default, the Cisco Cloud APIC will deploy a pair of redundant Cisco Cloud Services Router 1000Vs. The procedures in this section creates two tunnels, one IPsec tunnel from the on-premises IPsec device to each of these Cisco Cloud Services Router 1000Vs.

The following information provides commands for Cisco Cloud Services Router 1000V as your on-premises IPsec termination device. Use similar commands if you are using a different device or platform.

Step 1 Gather the necessary information that you will need to enable connectivity between the CSRs deployed in the cloud site and the on-premises IPsec termination device.

You can get the required configuration details using either the **Deploy & Download IPN Device config files** or the **Download IPN Device config files only** option in ACI Multi-Site Orchestrator as part of the procedures provided in Deploying Infra Configuration, on page 6.

Step 2 Log into the on-premises IPsec device.

Step 3 Configure the tunnel for the *first* CSR.

Details for the first CSR are available in the configuration files for the ISN devices you downloaded from the Multi-Site Orchestrator.

In the following example, replace:

- <*first-CSR-tunnel-ID*> with a unique tunnel ID that you assign to this tunnel.
- <*first-CSR-elastic-IP-address*> with the elastic IP address of the third network interface of the first CSR.
- <*first-CSR-preshared-key*> with the preshared key of the first CSR.
- <*interface*> with the interface that is used for connecting to the Cisco Cloud Services Router 1000V deployed in Amazon Web Services.
- <*peer-tunnel-for-onprem-IPsec-to-first-CSR*> with the peer tunnel IP address for the on-premises IPsec device to the first cloud CSR.
- <*process-id*> with the OSPF process ID.
- <*area-id*> with the OSPF area ID.

```
crypto isakmp policy 1
    encryption  aes
    authentication pre-share
    group  2
    lifetime 86400
    hash sha
exit

crypto keyring infra:overlay-1-<first-CSR-tunnel-ID>
    pre-shared-key address <first-CSR-elastic-IP-address> key <first-CSR-preshared-key>
exit

crypto isakmp profile infra:overlay-1-<first-CSR-tunnel-ID>
    local-address <interface>
    match identity address <first-CSR-elastic-IP-address>
    keyring infra:overlay-1-<first-CSR-tunnel-ID>
exit

crypto ipsec transform-set infra:overlay-1-<first-CSR-tunnel-ID> esp-aes esp-sha-hmac
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-<first-CSR-tunnel-ID>
    set pfs group2
    set security-association lifetime seconds 86400
exit

interface tunnel <first-CSR-tunnel-ID>
    ip address <peer-tunnel-for-onprem-IPsec-to-first-CSR> 255.255.255.252
    ip virtual-reassembly
    tunnel source <interface>
    tunnel destination <first-CSR-elastic-IP-address>
    tunnel mode ipsec ipv4
```

```
    tunnel protection ipsec profile infra:overlay-1-<first-CSR-tunnel-ID>
    ip mtu 1476
    ip tcp adjust-mss 1460
    ip ospf <process-id> area <area-id>
    no shut
exit
```

Example:

```
crypto isakmp policy 1
    encryption  aes
    authentication pre-share
    group  2
    lifetime 86400
    hash sha
exit

crypto keyring infra:overlay-1-1000
    pre-shared-key address 192.0.2.20 key 12345678900987654321234567890
exit

crypto isakmp profile infra:overlay-1-1000
    local-address GigabitEthernet1
    match identity address 192.0.2.20
    keyring infra:overlay-1-1000
exit

crypto ipsec transform-set infra:overlay-1-1000 esp-aes esp-sha-hmac
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-1000
    set pfs group2
    set security-association lifetime seconds 86400
exit

interface tunnel 1000
    ip address 30.29.1.2 255.255.255.252
    ip virtual-reassembly
    tunnel source GigabitEthernet1
    tunnel destination 192.0.2.20
    tunnel mode ipsec ipv4
    tunnel protection ipsec profile infra:overlay-1-1000
    ip mtu 1476
    ip tcp adjust-mss 1460
    ip ospf 1 area 1
    no shut
exit
```

Step 4 Configure the tunnel for the *second* CSR.

Details for the second CSR are also available in the configuration files for the ISN devices you downloaded from the Multi-Site Orchestrator.

```
crypto isakmp policy 1
    encryption  aes
    authentication pre-share
    group  2
    lifetime 86400
    hash sha
exit

crypto keyring infra:overlay-1-<second-CSR-tunnel-ID>
    pre-shared-key address <second-CSR-elastic-IP-address> key <second-CSR-preshared-key>
exit
```

```
crypto isakmp profile infra:overlay-1-<second-CSR-tunnel-ID>
    local-address <interface>
    match identity address <second-CSR-elastic-IP-address>
    keyring infra:overlay-1-<second-CSR-tunnel-ID>
exit

crypto ipsec transform-set infra:overlay-1-<second-CSR-tunnel-ID> esp-aes esp-sha-hmac
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-<second-CSR-tunnel-ID>
    set pfs group2
    set security-association lifetime seconds 86400
exit

interface tunnel <second-CSR-tunnel-ID>
    ip address <peer-tunnel-for-onprem-IPsec-to-second-CSR> 255.255.255.252
    ip virtual-reassembly
    tunnel source <interface>
    tunnel destination <second-CSR-elastic-IP-address>
    tunnel mode ipsec ipv4
    tunnel protection ipsec profile infra:overlay-1-<second-CSR-tunnel-ID>
    ip mtu 1476
    ip tcp adjust-mss 1460
    ip ospf <process-id> area <area-id>
    no shut
exit
```

Example:

```
crypto isakmp policy 1
    encryption  aes
    authentication pre-share
    group  2
    lifetime 86400
    hash sha
exit

crypto keyring infra:overlay-1-1001
    pre-shared-key address 192.0.2.21 key 12345678900987654321234567891
exit

crypto isakmp profile infra:overlay-1-1001
    local-address GigabitEthernet1
    match identity address 192.0.2.21
    keyring infra:overlay-1-1001
exit

crypto ipsec transform-set infra:overlay-1-1001 esp-aes esp-sha-hmac
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-1001
    set pfs group2
    set security-association lifetime seconds 86400
exit

interface tunnel 1001
    ip address 30.29.1.6 255.255.255.252
    ip virtual-reassembly
    tunnel source GigabitEthernet1
    tunnel destination 192.0.2.21
    tunnel mode ipsec ipv4
    tunnel protection ipsec profile infra:overlay-1-1001
```

```
    ip mtu 1476
    ip tcp adjust-mss 1460
    ip ospf 1 area 1
    no shut
exit
```

Step 5 Repeat these steps for any additional CSRs that you need to configure.

Step 6 Verify that the tunnels are up on your on-premises IPsec device.

Use the following command to display the status. If you do not see that both tunnels are shown as up, verify the information that you entered in the steps in this section to determine where you might have an issue. Do not proceed to the next section until you see that both tunnels are shown as up.

```
ISN_CSR# show ip interface brief | include Tunnel
Interface              IP-Address      OK? Method Status                Protocol
Tunnel1000             30.29.1.2         YES manual up                    up
Tunnel1001             30.29.1.4         YES manual up                    up
```