

# Configuring Infra for Cisco APIC Sites


- Refreshing Site Connectivity Information, on page 1
- Configuring Infra: On-Premises Site Settings, on page 2
- Configuring Infra: Pod Settings, on page 4
- Configuring Infra: Spine Switches, on page 4


# 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.

### Procedure

**Step 1** Log in to the Cisco Nexus Dashboard 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 **Refresh** button to pull fabric information from the APIC.

**Step 6** (Optional) For on-premises sites, 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: On-Premises Site Settings

This section describes how to configure site-specific Infra settings for on-premises sites.

### Procedure

**Step 1** Log in to the Cisco Nexus Dashboard Orchestrator GUI.

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

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

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

**Step 5** Provide the **Overlay Configuration**.

a) In the right **<*Site*> Settings** pane, select the **Overlay Configuration** tab.

b) In the right **<*Site*> Settings** pane, enable the **Multi-Site** knob.

This defines whether the overlay connectivity is established between this site and other sites.

c) (Optional) Enable the **CloudSec Encryption** knob encryption for the site.

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

d) 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.

This address should not be taken from the address space of the original fabric's `Infra` TEP pool or from the `0.x.x.x` range.

e) (Optional) From the **External Routed Domain** dropdown, select the domain you want to use.

Choose an external router domain that you have created in the Cisco APIC GUI. For more information, see the *Cisco APIC Layer 3 Networking Configuration Guide* specific to your APIC release.

f) Specify the **BGP Autonomous System Number**.

g) (Optional) Specify the **BGP Password**.

h) (Optional) Enable **SR-MPLS Connectivity** for the site.

If the site is connected via an MPLS network, enable the **SR-MPLS Connectivity** knob and provide the Segment Routing global block (SRGB) range.

The Segment Routing Global Block (SRGB) is the range of label values reserved for Segment Routing (SR) in the Label Switching Database (LSD). These values are assigned as segment identifiers (SIDs) to SR-enabled nodes and have global significance throughout the domain.

The default range is `16000-23999`.

If you enable MPLS connectivity for the site, you will need to configure additional settings as described in the "Sites Connected via SR-MPLS" chapter of the *Cisco Multi-Site Configuration Guide for ACI Fabrics*.

**Step 6** Provide the **Underlay Configuration**.

a) In the right **<*Site*> Settings** pane, select the **Underlay Configuration** tab.

b) Select the **OSPF Area Type** from the dropdown menu.

The OSPF area type can be one of the following:

- `nssa`
- `regular`

c) 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`.

**Step 7** Configure inter-site connectivity between on-premises and cloud sites.

If you do not need to create inter-site connectivity between on-premises and cloud sites, for example if your deployment contains only cloud or only on-premises sites, skip this step.

When you configure underlay connectivity between on-premises and cloud sites, you need to provide an IPN device IP address to which the Cloud APIC's CSRs establish a tunnel and then configure the cloud site's infra settings.

a) Click **+Add IPN Device** to specify an IPN device.

b) From the dropdown, select one of the IPN devices you defined previously.

The IPN devices must be already defined in the **General Settings** > **IPN Devices** list, as described in Configuring Infra: General Settings

c) Configure inter-site connectivity for cloud sites.

Any previously configured connectivity from the cloud sites to this on-premises site will be displayed here, but any additional configuration must be done from the cloud site's side as described in Configuring Infra for Cisco Cloud APIC Sites.

### What to do next

While you have configured all the required inter-site connectivity information, it has not been pushed to the sites yet. You need to deploy the configuration as described in Deploying Infra Configuration

# Configuring Infra: Pod Settings

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

### Procedure

**Step 1** Log in to the Cisco Nexus Dashboard 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 an external routable TEP pool.

The external routable TEP pools are used to assign a set of IP addresses that are routable across the ISN to APIC nodes, spine switches, and border leaf nodes. This is required to enable the intersite L3Out functionality.

External TEP pools previously assigned to the fabric on APIC are automatically inherited by NDO and displayed in the GUI when the fabric is added to the Multi-Site domain.

**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 Multi-Site.

### Procedure

**Step 1** Log in to the Cisco Nexus Dashboard 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.

  NDO 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`, which would configure an MTU of 9150B, or choose 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: On-Premises Site Settings, on page 2.

  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.