

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 5
- Configuring Infra: Spine Switches, on page 5

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 left navigation menu, select **Config** > **Site To Site Connectivity**.

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

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 check box, 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 discovers any new or removed spines and all site-related fabric connectivity will be reimported 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.

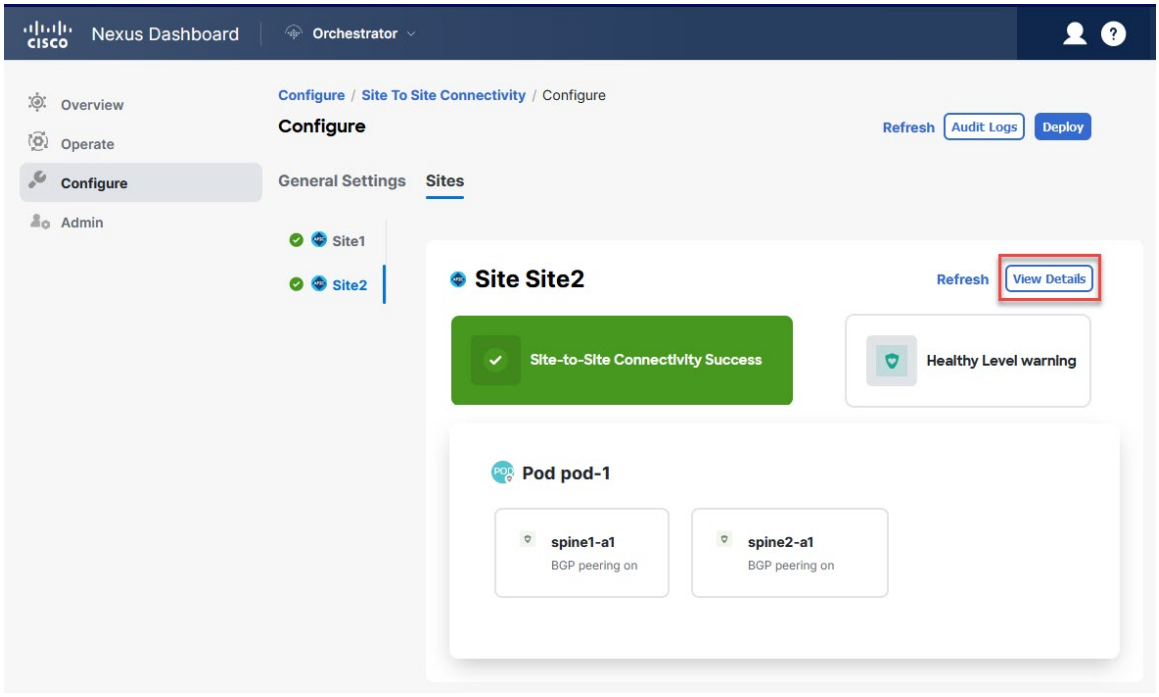
Step 2 In the left navigation menu, select **Configure** > **Site To Site Connectivity**.

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

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

Step 5 Click **View Details** to load site settings.

Figure 1:



Step 6 Provide the **Inter-Site Connectivity** information.

a) 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.

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

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

c) Specify the **Overlay Multicast TEP**.

This address is used for the intersite 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 multipod 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.

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

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

f) Provide the **OSPF Area ID**.

The following settings are required if you are using OSPF protocol for underlay connectivity between the site and the IPN. If you plan to use BGP instead, you can skip this step. BGP underlay configuration is done at the port level, as described in Configuring Infra: Spine Switches, on page 5.

g) Select the **OSPF Area Type** from the drop-down list.

The following settings are required if you are using OSPF protocol for underlay connectivity between the site and the IPN. If you plan to use BGP instead, you can skip this step. BGP underlay configuration is done at the port level, as described in Configuring Infra: Spine Switches, on page 5.

The OSPF area type can be one of the following:

- `nssa`
- `regular`

h) Configure OSPF policies for the site.

The following settings are required if you are using OSPF protocol for underlay connectivity between the site and the IPN. If you plan to use BGP instead, you can skip this step. BGP underlay configuration is done at the port level, as described in Configuring Infra: Spine Switches, on page 5.

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** drop-down list, 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`.

i) (Optional) From the **External Routed Domain** drop-down, select the domain that 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.

j) (Optional) Enable **SDA Connectivity** for the site.

If the site is connected to an SDA network, enable the **SDA Connectivity** knob and provide the **External Routed Domain**, **VLAN Pool**, and **VRF Lite IP Pool Range** information.

If you enable SDA connectivity for the site, you need to configure extra settings as described in the SDA use case chapter of the *Cisco Multi-Site Configuration Guide for ACI Fabrics*.

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

If the site is connected through 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 that are 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 need to configure extra settings as described in the "Sites Connected through SR-MPLS" chapter of the *Cisco Multi-Site Configuration Guide for ACI Fabrics*.

Step 7 Configure intersite connectivity between on-premises and cloud sites.

If you do not need to create intersite 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 must provide an IPN device IP address to which the Cloud Network Controller'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 drop-down, 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 intersite 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 Network Controller Sites.

What to do next

While you have configured all the required intersite connectivity information, it has not been pushed to the sites yet. You must 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 left navigation menu, select **Configure** > **Site To Site Connectivity**.

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

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 sourcing and receiving VXLAN encapsulated traffic for Layer2 and Layer3 unicast communication.

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 IPN to APIC nodes, spine switches, and border leaf nodes. This is required to enable Multi-Site architecture.

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. When you configure the spine switches, you are effectively establishing the underlay connectivity between the sites in your Multi-Site domain by configuring connectivity between the spines in each site and the ISN.

Before Release 3.5(1), underlay connectivity was establishing using OSPF protocol. In this release however, you can choose to use OSPF, BGP (IPv4 only), or a mixture of protocols, with some sites using OSPF and some using BGP for intersite underlay connectivity. We recommend configuring either OSPF or BGP and

not both, however if you configure both protocols, BGP will take precedence and OSPF will not be installed in the route table.

Procedure

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

Step 2 In the left navigation menu, select **Config** > **Site To Site Connectivity**.

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

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

Step 5 In the main pane, 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, provide the underlay connectivity information.

Any port that is already configured directly in APIC for IPN connectivity will be imported and shown in the list. For any new ports you want to configure from NDO, use the following the steps:

a) Provide general information:

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

  This is the interface which will be used to connect to the IPN.

- In the **IP Address** field, enter the IP address/netmask.

  The Orchestrator creates a subinterface 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.

Step 8 Choose the underlay protocol.

a) Enable **OSPF** if you want to use OSPF protocol for underlay connectivity.

If you want to use BGP protocol for underlay connectivity instead, skip this part and provide the information that is required in the next substep.

- Set **OSPF** to `Enabled`.

  The OSPF settings become available.

- From the **OSPF Policy** drop-down, select 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`

- Set **BGP** to `Disabled`.

b) Enable **BGP** if you want to use BGP protocol for underlay connectivity.

If you're using OSPF protocol for underlay connectivity and have already configured it in the previous substep, skip this part.

Note
BGP IPv4 underlay is not supported in the following cases:

- If your Multi-Site domain contains one or more Cloud Network Controller sites, in which case you must use the OSPF protocol for intersite underlay connectivity for both On-Prem to On-Prem and On-Prem to cloud sites.
- If you are using GOLF (Layer 3 EVPN services for fabric WAN) for WAN connectivity in any of your fabrics.

In the above cases, you must use OSPF in the Infra L3Out deployed on the spines.

- Set **OSPF** to `Disabled`.

  We recommend configuring either OSPF or BGP and not both, however if you configure both protocols, BGP will take precedence and OSPF routes will not be installed in the route table because only EBGP adjacencies with the ISN devices are supported.

- Set **BGP** to `Enabled`.

  The BGP settings become available.

- In the **Peer IP** field, provide the IP address of this port's BGP neighbor.

  Only IPv4 IP addresses are supported for BGP underlay connectivity.

- In the **Peer AS Number** field, provide the Autonomous System (AS) number of the BGP neighbor.

  This release supports only EBGP adjacencies with the ISN devices.

- In the **BGP Password** field, provide the BGP peer password.
- Specify any additional options as required:
  - `Bidirectional Forwarding Detection`—Enables Bidirectional Forwarding Detection (BFD) protocol to detect faults on the physical link this port and the IPN device.
  - `Admin State`—Sets the admin state on the port to enabled.

Step 9 Repeat the procedure for every spine switch and port that connects to the IPN.