



Configuring Infra for Cisco NDFC Sites

- [Prerequisites and Guidelines, on page 1](#)
- [Configuring Infra: General Settings, on page 1](#)
- [Refreshing Site Connectivity Information, on page 4](#)
- [Configuring Infra: NDFC Site-Specific Settings, on page 4](#)
- [Deploying Infra Configuration, on page 6](#)

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 added the sites as described in previous sections.

In addition, keep in mind the following:

- Adding or removing border gateway switches requires a Nexus Dashboard Orchestrator fabric connectivity information refresh described in the [Refreshing Site Connectivity Information, on page 4](#) as part of the general Infra configuration procedures.

Configuring Infra: General Settings

This section describes how to configure general Infra settings for your NDFC sites that are onboarded and managed by Nexus Dashboard Orchestrator.

-
- Step 1** Log in to your Nexus Dashboard and open the Nexus Dashboard Orchestrator service.
 - Step 2** In the left navigation menu, select **Infrastructure > Site Connectivity**.
 - Step 3** In the main pane, click **Configure**.
 - Step 4** In the left sidebar, select **General Settings**.
 - Step 5** Provide **Control Plane Configuration**.
 - a) Select the **Control Plane Configuration** tab.
 - b) Choose **BGP Peering Type**.

- `full-mesh`—All border gateway switches in each site will establish peer connectivity with remote sites' border gateway switches.
 - `route-server`—The route-server option allows you to specify one or more control-plane nodes to which each site establishes MP-BGP EVPN sessions. The route-server nodes perform a function similar to traditional BGP route-reflectors, but for EBGP (and not iBGP) sessions. The use of route-server nodes avoids creating MP-BGP EVPN full mesh adjacencies between all the VXLAN EVPN sites managed by NDO.
- c) If you set the **BGP Peering Type** to `route-server`, click **+Add Route Server** to add one or more route servers. In the **Add Route Server** window that opens:
- From the **Site** dropdown, select the site you want to connect to the route server.
 - The **ASN** field will be auto-populated with the site's ASN.
 - From the **Core Router Device** dropdown, select the route server to which you want to connect.
 - From the **Interface** dropdown, select the interface on the core router device.
- You can add up to 4 route servers. If you add multiple route servers, every site will establish MP-BGP EVPAN adjacencies to every route server.
- d) Leave the **Keepalive Interval (Seconds)**, **Hold Interval (Seconds)**, **Stale Interval (Seconds)**, **Graceful Helper**, **Maximum AS Limit**, and **BGP TTL Between Peers** fields at default values as they are relevant for Cisco ACI fabrics only.
- e) Skip the **OSPF Area ID** and **External Subnet Pool** fields at default values as they are relevant for Cloud Network Controller fabrics only.

Step 6 Provide the **On Premises IPsec Devices** information.

If your inter-site connectivity between on-premises and cloud sites is using private connection and you will not enable IPsec, you can skip this step. For connectivity over public Internet, IPsec is always enabled and you must provide the information in this step.

When you configure inter-site underlay connectivity between on-premises and cloud sites as described in later sections, you will need to select an on-premises IPN device which will establish connectivity to the cloud CSRs. These IPN devices must first be defined here before they are available in the on-premises site configuration screen.

- a) Select the **On Premises IPsec Devices** tab.
- b) Click **+Add On-Premises IPsec Device**.
- c) Choose whether the device is **Unmanaged** or **Managed** and provide the device information.

This defines whether or not the device is directly managed by NDFC:

- For **Unmanaged** IPN devices, simply provide the **Name** and the **IP Address** of the device.

The IP address you provide will be used as the tunnel peer address from the cloud CSRs, not the IPN device's management IP address.

- For **Managed** IPN devices, choose the NDFC **Site** that contains the device and then the **Device** from that site.

Then choose the **Interface** on the device that is facing the Internet and provide the **Next Hop** IP address, which is the IP address of the gateway that is connecting to the Internet.

- d) Click the check mark icon to save the device information.
- e) Repeat this step for any additional IPN devices you want to add.

Step 7 Provide the **IPSec Tunnel Subnet Pools** information.

There are two types of subnet pools that you can provide here:

- **External Subnet Pool**—used for connectivity between cloud site CSRs and other sites (cloud or on-premises).

These are large global subnet pools that are managed by Nexus Dashboard Orchestrator. The Orchestrator, creates smaller subnets from these pools and allocates them to sites to be used for inter-site IPsec tunnels and external connectivity IPsec tunnels.

You must provide at least one external subnet pool if you want to enable external connectivity from one or more of your cloud sites.

- **Site-Specific Subnet Pool**—used for connectivity between cloud site CSRs and external devices.

These subnets can be defined when the external connectivity IPsec tunnels must be in a specific range. For example, where a specific subnet is already being used to allocate IP addresses to the external router and you want to continue using those subnets for IPsec tunnels for NDO and cloud sites. These subnets are not managed by the Orchestrator and each subnet is assigned to a site in its entirety to be used locally for external connectivity IPsec tunnels.

If you do not provide any named subnet pools but still configure connectivity between cloud site's CSRs and external devices, the external subnet pool will be used for IP allocation. .

Note The minimum mask length for both subnet pools is /24.

To add one or more **External Subnet Pools**:

- a) Select the **IPSec Tunnel Subnet Pools** tab.
- b) In the **External Subnet Pool** area, click **+Add IP Address** to add one or more external subnet pools.

This subnet will be used to address the IPsec tunnel interfaces and loopbacks of the Cloud Routers used for on-premises connectivity, which you previously configured in the Cloud Network Controller for inter-site connectivity in earlier Nexus Dashboard Orchestrator releases.

The subnets must not overlap with other on-premises TEP pools, should not begin with 0.x.x.x or 0.0.x.x, and should have a network mask between /16 and /24, for example 30.29.0.0/16.

- c) Click the check mark icon to save the subnet information.
- d) Repeat these substeps for any additional subnet pools you want to add.

To add one or more **Site-Specific Subnet Pools**:

- a) Select the **IPSec Tunnel Subnet Pools** tab.
- b) In the **Site-Specific Subnet Pools** area, click **+Add IP Address** to add one or more external subnet pools.

The **Add Named Subnet Pool** dialogue will open.

- c) Provide the subnet **Name**.
You will be able to use the subnet pool's name to choose the pool from which to allocate the IP addresses later on.
- d) Click **+Add IP Address** to add one or more subnet pools.

The subnets must have a network mask between /16 and /24 and not begin with 0.x.x.x or 0.0.x.x, for example 30.29.0.0/16.

- e) Click the check mark icon to save the subnet information.
Repeat the steps if you want to add multiple subnets to the same named subnet pool.
- f) Click **Save** to save the named subnet pool.

- g) Repeat these substeps for any additional named subnet pools you want to add.

Step 8 Configure **NDFC Settings**.

- a) Select the **NDFC Settings** tab.
- b) Provide the **L2 VXLAN VNI Range**.
- c) Provide the **L3 VXLAN VNI Range**.
- d) Provide the **Multi-Site Routing Loopback IP Range**.

This field is used to auto-populate the **Multi-Site TEP** field for each fabric, which is described in [Configuring Infra: NDFC Site-Specific Settings, on page 4](#).

For sites that were previously part of a Multi-Site Domain (MSD) in NDFC, this field will be pre-populated with the previously defined value.

- e) Provide the **Anycast Gateway MAC**.
-

Refreshing Site Connectivity Information

Infrastructure changes, such as adding and removing border gateway switches, require a Nexus Dashboard Orchestrator fabric connectivity refresh. This section describes how to pull up-to-date connectivity information directly from each site's controller.

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

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

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

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

Step 5 In the main window, click the **Refresh** button to pull fabric information from the controller.

Step 6 (Optional) In the **Confirmation** dialog, check the box if you want to remove configuration for decommissioned border gateway switches.

If you choose to enable this checkbox, all configuration info for any currently decommissioned border gateway 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 site's controller.

Configuring Infra: NDFC Site-Specific Settings

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

Step 1 Log in to your Nexus Dashboard and open the Nexus Dashboard Orchestrator service.

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

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

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

Step 5 In the right **<Site> Settings** sidebar, 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 border gateway switches that are part of the same fabric.

Note If the site you are configuring is part of the NDFC Multi-Site Domain (MDS), this field will be pre-populated with the information imported from NDFC. In this case, changing the value and re-deploying the infra configuration, will impact traffic between the sites that are part of the MDS.

You can choose to **Auto Allocate** this field, which will allocate the next available address from the **Multi-Site Routing Loopback IP Range** you defined in previous section.

Step 6 Within the **<fabric-name>** tile, select the border gateway.

Step 7 In the right **<border-gateway>** setting sidebar, specify the **BGP-EVPN ROUTER-ID** and **BGW PIP**.

For border gateways that are part of a vPC domain, you must also specify a **VPC VIP**

Step 8 Click **Add Port** to configure the port that connects to the IPN.

Note This release does not support importing the port configuration from the NDFC. If the site you are configuring is already part of the NDFC Multi-Site Domain (MDS), you must use the same values that are already configured in NDFC.

Update Port
✕

* Ethernet Port ID

* IP Address

* Remote Address

* Remote ASN

* MTU

BGP Authentication

None Simple

Provide the following information specific to your deployment for the port that connects this border gateway to a core switch or another border gateway:

- From the **Ethernet Port ID** dropdown, select the port that connects to the IPN.
- In the **IP Address** field, enter the IP address and netmask.
- In the **Remote Address** field, provide the IP address of the remote device to which the port is connected.
- In the **Remote ASN** field, provide the remote site's ID.
- In the **MTU** field, enter the port's MTU.

MTU of the spine port should match MTU on IPN side.

You can specify either `inherit` or a value between 576 and 9000.

- For **BGP Authentication**, you can pick either `None` or `Simple` (MD5).
If you select `Simple`, you must also provide the **Authentication Key**.

Deploying Infra Configuration

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

Before you begin

You must have the general and site-specific infra configurations completed as described in previous sections of this chapter.

Step 1 Ensure that there are no configuration conflicts or resolve them if necessary.

The **Deploy** button will be disabled and a warning will be displayed if there are any configuration conflicts from the already configured settings in each site. For example, if a VRF or network with the same name exists in multiple sites but uses different VNI in each site.

In case of configuration conflicts:

- Click **Click to View** link in the conflict notification pop-up.



- Note down the specific configurations that are causing the conflicts.

For example, in the following report, there are ID mismatches between VRFs and networks in `fab1` and `fab2` sites.

Error Type	Error Message
IDMismatch	Policy Name MyVRF_50001 Policy ID 50001 Sites [fab2] conflicting with Policy Name MyVRF_50001 Policy ID 60001 Sites [fab1]
IDMismatch	Policy Name MyNetwork_30000 Policy ID 40000 Sites [fab2] conflicting with Policy Name MyNetwork_30000 Policy ID 30000 Sites [fab1]

c) Click the **X** button to close the report, then exit Infra configuration screen.

d) Unmanage the site in NDO, as described in [Removing Sites](#).

You do not need to remove the site from the Nexus Dashboard, simply unmanage it in NDO GUI.


e) Resolve the existing configuration conflicts.

f) Manage the site again, as described in [Adding Cisco NDFC Sites](#).

Since the site is already added in Nexus Dashboard, simply enable it for management in NDO.

g) Verify that all conflicts are resolved and the **Deploy** button is available.

Step 2 Deploy configuration.



a) In the top right of the **Fabric Connectivity Infra** screen, choose the appropriate **Deploy** option to deploy the configuration.

If you are configuring only NDFC sites, simply click **Deploy** to deploy the Infra configuration.

b) Wait for configuration to be deployed.

When you deploy infra configuration, NDO will signal the NDFC to configure the underlay and the EVPN overlay between the border gateways.

When configuration is successfully deployed, you will see a green checkmark next to the site in the **Fabric Connectivity Infra** screen:

