



Nexus Dashboard Orchestrator Infrastructure Configuration for NDFC Fabrics, Release 4.4.x

Table of Contents

Configuring Infra for Cisco NDFC Fabrics	1
Prerequisites and Guidelines	1
Configuring Infra: General Settings	1
Refreshing Fabric Connectivity Information	4
Configuring Infra: NDFC Fabric-Specific Settings	5
Deploying Infra Configuration	7
Adding and Deleting Fabrics	9
Adding Cisco NDFC Fabrics	9
Removing Fabrics	11
Cross Launch to Fabric Controllers	12

Configuring Infra for Cisco NDFC Fabrics

Prerequisites and Guidelines

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

Before you proceed with Infra configuration, you must have added the fabrics 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 Fabric Connectivity Information](#) as part of the general Infra configuration procedures.
- With NDFC integration with Nexus Dashboard Orchestrator, after the RMA of the switches into NDFC, you need to perform a manual import of your networks and VRFs on Nexus Dashboard Orchestrator to view the diff of the schema on NDFC. After performing the RMA on NDFC, Nexus Dashboard Orchestrator sees a change of the serial number after a refresh on the Nexus Dashboard Orchestrator fabric. NDFC requires a reimport of the switch for the serial number to be in sync.

Configuring Infra: General Settings

This section describes how to configure general settings for your NDFC fabrics that are on board and managed by Cisco Nexus Dashboard Orchestrator.

1. Log in to your Cisco Nexus Dashboard and open the Cisco Nexus Dashboard Orchestrator service.
2. In the left navigation menu, choose **Configure > Fabric to Fabric Connectivity**.
3. In the main pane, choose the **Configure** button.
4. Choose the left tab, which is for **General Settings**.
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 fabric establishes peer connectivity with remote fabrics' border gateway switches.
 - **route-server**—The route-server option allows you to specify one or more control-plane nodes to which each fabric establishes MP-BGP EVPN sessions. The route-server nodes perform a function similar to traditional BGP route-reflectors, but for External Border Gateway Protocol (and not Internal Border Gateway Protocol) sessions. The use of route-server nodes avoids creating MP-BGP EVPN full mesh adjacencies between all the VXLAN EVPN fabrics that are 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 **Fabric** drop-down, select the fabric that you want to connect to the route server.
- The **ASN** field will be autopopulated with the fabric's ASN.
- From the **Core Router Device** drop-down, select the route server to which you want to connect.
- From the **Interface** drop-down, select the interface on the core router device.

You can add up to 4 route servers. If you add multiple route servers, every fabric establishes MP-BGP EVPN adjacencies to every route server.

- d. Leave the **Keepalive Interval (Seconds)**, **Hold Interval (Seconds)**, **Stale Interval (Seconds)**, **Graceful Restart**, **Maximum AS Limit**, and **BGP TTL Between Peers** fields at default values as they are relevant for Cisco ACI fabrics only.

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

If your inter-fabric connectivity between on-premises and cloud fabrics 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-fabric underlay connectivity between on-premises and cloud fabrics as described in later sections, you must select an on-premises IPN device which establishes connectivity to the cloud CSRs. These IPN devices must first be defined here before they are available in the on-premises fabric 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 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 that 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 **Fabric** that contains the device and then the **Device** from that fabric.

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 that you want to add.

7. Provide the **IPsec Tunnel Subnet Pools** information.

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

- **External Subnet Pool**-Used for connectivity between cloud fabric CSRs and other fabrics (cloud or on-premises).

These are large global subnet pools that are managed by Cisco Nexus Dashboard Orchestrator. The Orchestrator creates smaller subnets from these Pools and allocates them to fabrics to be used for inter-fabric 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 fabrics.

- o **Fabric-Specific Subnet Pool**—Used for connectivity between cloud fabric 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 fabrics. These subnets are not managed by the Orchestrator and each subnet is assigned to a fabric 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 the cloud fabric's CSRs and external devices, the external subnet pool will be used for IP allocation. .



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 that are used for on-premises connectivity, which you previously configured in the Cloud Network Controller for inter-fabric connectivity in earlier Cisco 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 that you want to add.

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

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

The **Add Named Subnet Pool** dialogue opens.

- c. Provide the subnet **Name**.

You can 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 that you want to add.

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-Fabric Routing Loopback IP Range**.

This field is used to autopopulate the **Multi-Fabric TEP** field for each fabric, which is described in [Configuring Infra: NDFC Fabric-Specific Settings](#).

For fabrics that were previously part of a Multi-Fabric Domain (**MSD**) in NDFC, this field will be prepopulated with the previously defined value.

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

Refreshing Fabric Connectivity Information

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

1. Log in to the Cisco Nexus Dashboard Orchestrator GUI.
2. In the left navigation menu, select **Configure > Fabric to Fabric Connectivity**.
3. In the top right of the main pane, click **Configure**.
4. In the left sidebar, under **Fabrics**, select a specific fabric.
5. In the main window, click the **Refresh** button to pull fabric information from the controller.
6. (Optional) In the **Confirmation** dialog, check the box if you want to remove the configuration for decommissioned border gateway switches.

If you choose to enable this check box, all configuration info for any currently decommissioned border gateway switches will be removed from the database.

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

This discovers any new or removed border gateways and all fabric-related fabric connectivity will be reimported from the fabric's controller.

Configuring Infra: NDFC Fabric-Specific Settings

This section describes how to configure fabric-specific Infra settings for on-premises fabrics.

1. Log in to your Cisco Nexus Dashboard and open the Cisco Nexus Dashboard Orchestrator service.
2. In the left navigation menu, choose **Configure > Fabric to Fabric Connectivity**.
3. In the main pane, click **Configure**.
4. In the left pane, under **Fabrics**, choose a specific NDFC.
5. In the right **<Fabric> Settings** sidebar, specify the **Multi-Fabric VIP**.

This address is used for the inter-fabric L2 BUM and L3 multicast traffic. The IP address is deployed in all border gateway switches that are part of the same fabric.



If the fabric you are configuring is part of the NDFC Multi-Fabric Domain (MSD), this field will be prepopulated with the information imported from NDFC. In this case, changing the value and redeploying the configuration, impacts traffic between the fabrics that are part of the MSD.

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

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

You can also choose to deploy a shared border configuration, this configuration allows you to share these services through these devices, while also providing a means to cross the "border" from the internal fabrics to the Internet. For more information on see, **Configuring Shared Border in Cisco Nexus Dashboard Fabric Controller**.

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



This release does not support importing the port configuration from the NDFC. If you are configuring a fabric which is already part of the NDFC Multi-Fabric Domain (MSD), you must use the same values that are configured in NDFC. You can use the **Inherit BGP Authentication and BFD** radio button to inherit settings across, fabrics and fabrics.

Add Port

Description

Remote Address *

Remote ASN *

MTU *

9216

Inherit BGP Authentication and BFD ⓘ

☒

BGP Authentication

☒ None ☐ Simple ☐ Cisco

Towards Cloud Router ⓘ

☐

BFD Enabled

☒

Log Neighbor

☐

BGP Send Community

☐

Route Tag

66666

Enable Redistribute Direct

☐

RS Route Tag

Figure 1. NDFC Settings - Add Port

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

- o From the **Ethernet Port ID** drop-down, choose the port that connects to the IPN.
- o In the **IP Address** field, enter the IP address and netmask.

- o In the **Remote Address** field, provide the IP address of the remote device to which the port is connected.
- o In the **Remote ASN** field, provide the remote fabric's **Autonomous System Number**.
- o In the **MTU** field, enter the port's maximum transmission unit.

Maximum transmission unit of the spine port must match MTU on the IPN side.

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

- o For **BGP Authentication**, you can pick either **None** or **Simple** (MD5) or **Cisco**.

Provide the **Authentication Key** if you choose **Simple** or **Cisco** authentication methods.

- o Check the **BFD Enabled**, **Log Neighbor**, and **BGP Send Community** radio buttons to inherit all these features to all the multifabric underlay interfaces.
- o **Route Tag** is used to configure the data path across all the switches and is propagated to all the nodes in the fabric. Loopback 0, 1 and 100. Select **Enable Redistribute Direct** if the route tag is specified.

Deploying Infra Configuration

Before you begin:

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

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

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 fabric. For example, if a VRF or network with the same name exists in multiple fabrics but uses different VNI in each fabric.

If configuration conflicts:

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



Figure 2. Click to View

- b. 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** fabrics.

Sites Report

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]

Figure 3. Fabrics Report

- c. Click the **X** button to close the report, then exit Infra configuration screen.
- d. Unmanage the fabric in NDO, as described in [Removing Fabrics](#).

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

- e. Resolve the existing configuration conflicts.
- f. Manage the fabric again, as described in [Adding Cisco NDFC Fabrics](#).

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

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

2. Deploy configuration.

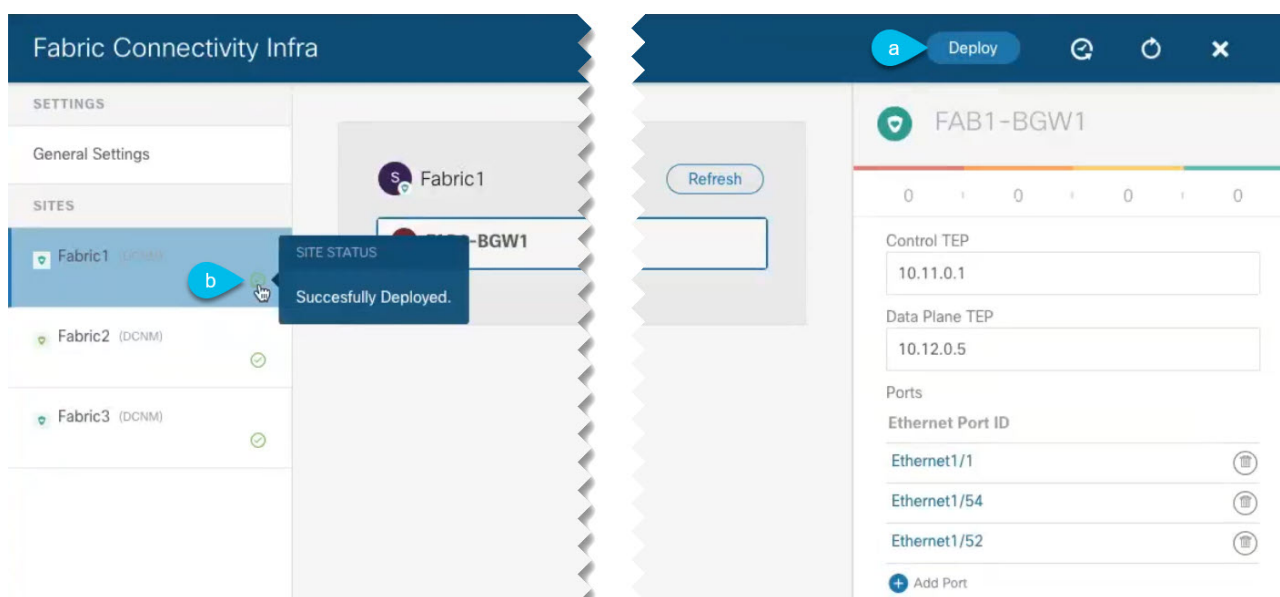


Figure 4. Deploy Configuration

- 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 fabrics, simply click **Deploy** to deploy the Infra configuration.

- b. Wait for configuration to be deployed.

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

When configuration is successfully deployed, you see a green check mark next to the fabric in the **Fabric Connectivity Infra** screen:

Adding and Deleting Fabrics

Adding Cisco NDFC Fabrics

Before you begin:

- You must ensure that the fabric(s) you are adding are running Cisco NDFC, Release 11.5(1) or later.

This section describes how to add a NDFC fabric using the Nexus Dashboard GUI and then enable that fabric to be managed by Nexus Dashboard Orchestrator.

1. Log in to your Nexus Dashboard and open the **Admin Console**.
2. From the left navigation menu, choose **Fabrics** and click **Add Fabric..**
3. Provide fabric information.
 - a. For **Fabric Type**, select **NDFC or NDFC**.
 - b. Provide the NDFC controller information.

You need to provide the **Host Name/IP Address** of the in-band (**eth2**) interface, **User Name**, and **Password**. for the NDFC controller currently managing your NDFC fabrics.

- c. Click **Select Fabrics** to select the specific fabrics managed by the controller.

In the fabric selection window that opens, select the fabrics you want to add to the Nexus Dashboard and click **Select**:

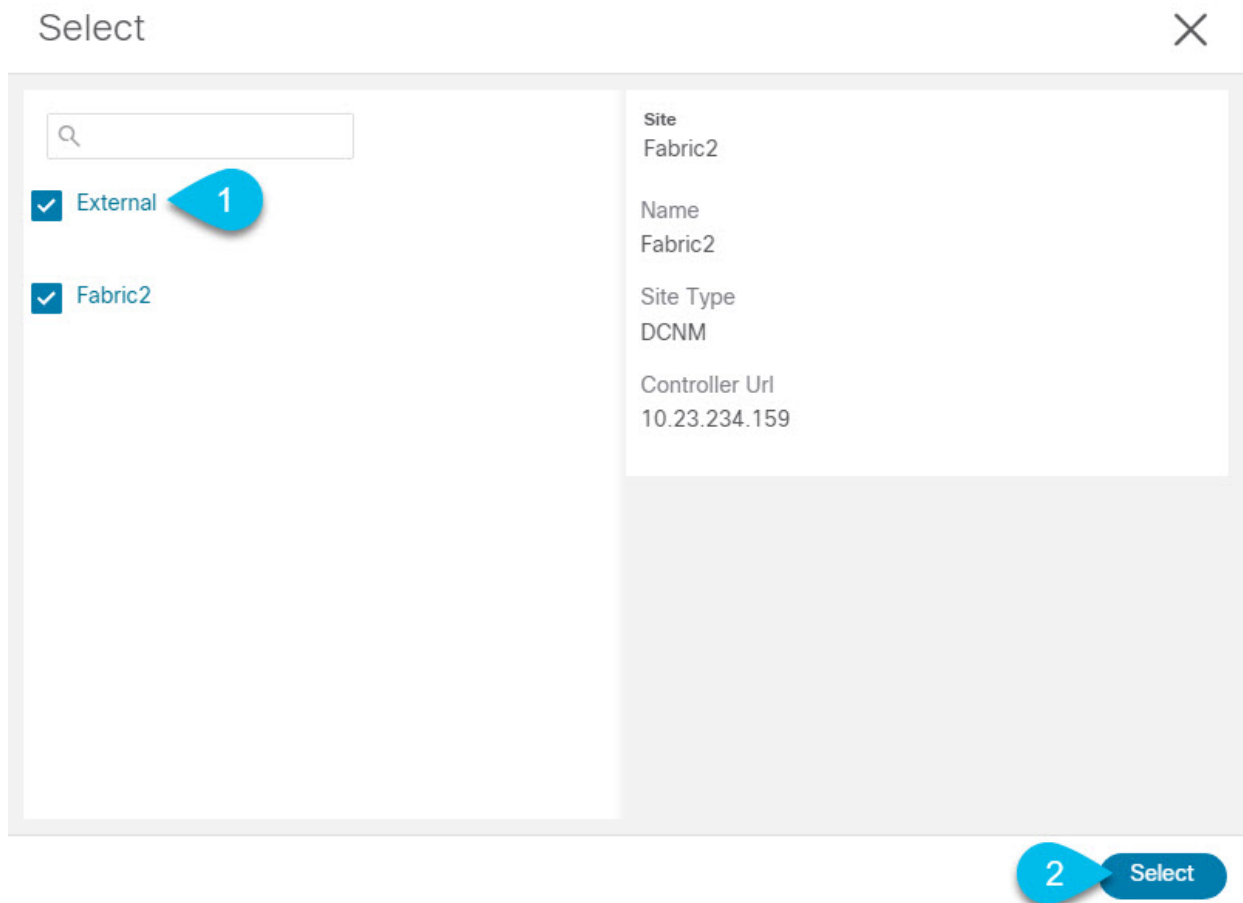


Figure 5.

- d. Click **Add Security Domains** to select one or more security domains that will have access to this fabric.
4. Repeat the previous steps for any additional NDFC fabrics.
5. From the Nexus Dashboard's **Service Catalog** page, open the Nexus Dashboard Orchestrator service.

You will be automatically logged in using the Nexus Dashboard user's credentials.

6. In the Nexus Dashboard Orchestrator GUI, manage the fabrics.

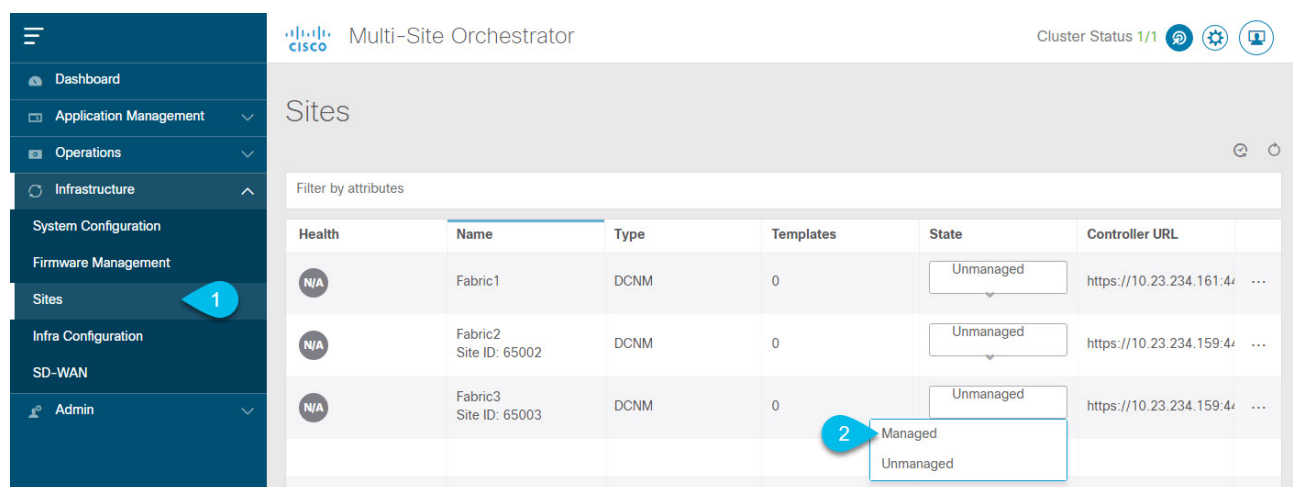


Figure 6.

- a. From the left navigation menu, select **Infrastructure > Fabrics**.

- b. In the main pane, change the **State** from **Unmanaged** to **Managed** for each fabric that you want the NDO to manage.

If the fabric you are managing is part of a Multi-Fabric Domain (MSD), it will have a **Fabric ID** already associated with it. In this case, simply changing the **State** to **Managed** will manage the fabric.

However, if the fabric is not part of an MSD, you will also be prompted to provide a **Fabric ID** for the fabric when you change its state to **Managed**.



If you want to manage both kinds of fabrics, those that are part of an existing MSD and those that are not, you must on-board the MSD fabrics first, followed by any standalone fabrics.

Removing Fabrics

Before you begin:

You must ensure that all templates associated with the fabric you want to remove are not deployed.

This section describes how to disable fabric management for one or more fabrics using the Nexus Dashboard Orchestrator GUI. The fabrics will remain present in the Nexus Dashboard.

1. Open the Nexus Dashboard Orchestrator GUI.

You can open the NDO service from the Nexus Dashboard's **Service Catalog**. You will be automatically logged in using the Nexus Dashboard user's credentials.

2. Remove the fabric from all templates.

You must remove the fabric from all templates with which it is associated before you can unmanage the fabric and remove it from your Nexus Dashboard.

- a. Navigate to **Application Management > Schemas**.
- b. Click a schema that contains one or more templates associated with the fabric.
- c. From the **View** dropdown, choose a template that's associated with the fabric that you want to remove.
- d. From the **Actions** dropdown, choose **Fabrics Association** and uncheck the fabric you want to remove.

This will remove configurations that were deployed using this template to this fabric.



For non-stretched templates, you can choose to preserve the configurations deployed by the template to the fabrics by selecting **Actions > Dissociate Fabrics** instead. This option will allow you to retain configurations deployed by NDO but no longer manage those objects from NDO.

- e. Repeat this step for all templates associated with the fabric that you want to unmanage in this and all other schemas.
3. In the Nexus Dashboard Orchestrator GUI, disable the fabrics.
 - a. From the left navigation menu, select **Fabrics**.

- b. In the main pane, change the **State** from **Managed** to **Unmanaged** for the fabric that you want to unmanage.



If the fabric is associated with one or more deployed templates, you will not be able to change its state to **Unmanaged** until you undeploy those templates, as described in the previous step.

4. Delete the fabric from Nexus Dashboard.

If you no longer want to manage this fabric or use it with any other applications, you can delete the fabric from the Nexus Dashboard as well.



Note that the fabric must not be currently in use by any of the services installed in your Nexus Dashboard cluster.

- a. In the top navigation bar, click the **Home** icon to return to the Nexus Dashboard GUI.
- b. From the left navigation menu of the Nexus Dashboard GUI, select **Fabrics**.
- c. Select one or more fabrics you want to delete.
- d. In the top right of the main pane, select **Actions > Delete Fabric**.
- e. Provide the fabric's login information and click **OK**.

The fabric will be removed from the Nexus Dashboard.

Cross Launch to Fabric Controllers

Cisco Nexus Dashboard Orchestrator currently supports several configuration options for each type of fabrics. For many extra configuration options, you may need to sign in directly into the fabric's controller.

You can cross-launch into the specific fabric controller's GUI from the NDO's **Operate > Fabrics** screen by selecting the actions (...) menu next to the fabric and clicking **Open in user interface**. Cross-launch works with out-of-band (OOB) management IP of the fabric.

If the same user is configured in Cisco Nexus Dashboard and the fabric, you will be signed in automatically into the fabric's controller using the same log in information as the Cisco Nexus Dashboard user. For consistency, we recommend configuring remote authentication with common users across Cisco Nexus Dashboard and the fabrics. :leveloffset: -1

First Published: 2024-03-11

Last Modified: 2024-07-26

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<http://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883