Cisco Nexus Dashboard Orchestrator Deployment Guide, Release 4.2(x)

First Published: 2023-08-23

Last Modified: 2023-11-11

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: http://www.cisco.com/go/softwareterms.Cisco product warranty information is available at http://www.cisco.com/go/warranty. US Federal Communications Commission Notices are found here http://www.cisco.com/c/en/us/products/us-fcc-notice.html.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

CONTENTS

CHAPTER 1

New and Changed Information

New and Changed Information

The following table provides an overview of the significant changes to the organization and features in this guide from the release the guide was first published to the current release. The table does not provide an exhaustive list of all changes made to the guide.

Table 1: Latest Updates

Release	New Feature or Update	Where Documented
4.1(2)	Updated the upgrade workflows.	• Upgrading Automatically Via Service Catalog, on page 69 • Upgrading Manually Using Configuration Restore, on page 89
4.1(1)	First release of this document.	--

CHAPTER 2

Deploying Nexus Dashboard Orchestrator

Deployment Overview

Cisco Nexus Dashboard Orchestrator (NDO) must be deployed as a service in Cisco Nexus Dashboard.

Cisco Nexus Dashboard is a central management console for multiple data center sites and a common platform for hosting Cisco data center operation services, such as Nexus Dashboard Insights and Nexus Dashboard Orchestrator. These services are available for all the data center sites and provide real time analytics, visibility, assurance for network policies and operations, as well as policy orchestration for the data center fabrics, such as Cisco ACI or Cisco NDFC.

Nexus Dashboard provides a common platform and modern technology stack for the above-mentioned micro-services-based applications, simplifying the life cycle management of the different modern applications and reducing the operational overhead to run and maintain these applications. It also provides a central integration point for external 3rd party applications with the locally hosted applications.

Each Nexus Dashboard cluster typically consists of 1 or 3 `master` nodes. For 3-node clusters, you can also provision a number of `worker` nodes to enable horizontal scaling and `standby` nodes for easy cluster recovery in case of a master node failure. For maximum number of `worker` and `standby` nodes supported in this release, see the "Verified Scalability Limits" sections of the *Cisco Nexus Dashboard Release Notes*.

For detailed information about Nexus Dashboard cluster initial deployment and configuration, see *Cisco Nexus Dashboard Deployment Guide*. For more information about using Nexus Dashboard, see the *Cisco Nexus Dashboard User Guide*.

This document describes initial installation requirements and procedures for the Nexus Dashboard Orchestrator service. Detailed configuration and use case information is available from the *Cisco Nexus Dashboard Orchestrator Configuration Guide for Cisco ACI* or *Cisco Nexus Dashboard Orchestrator Configuration Guide for Cisco NDFC* for your release and the Cisco Cloud Network Controller use case documents, depending on the type of fabrics you plan to manage.

Prerequisites and Guidelines

Nexus Dashboard

You must have Cisco Nexus Dashboard cluster deployed and its fabric connectivity configured, as described in *Cisco Nexus Dashboard Deployment Guide* before proceeding with any additional requirements and the Nexus Dashboard Orchestrator service installation described here.

Orchestrator Release	Minimum Nexus Dashboard Release
Release 4.2(1) and later	Cisco Nexus Dashboard, Release 3.0(1)

Nexus Dashboard Networks

When first configuring Nexus Dashboard, you will need to provide two IP addresses for the two Nexus Dashboard interfaces—one connected to the Data Network and the other to the Management Network. The data network is used for the nodes' clustering and Cisco fabrics traffic. The management network is used to connect to the Cisco Nexus Dashboard GUI, CLI, or API.

Note The two interfaces must be in different subnets.

Connectivity between the nodes is required on both networks with the round trip time (RTT) not exceeding 150ms for Nexus Dashboard Orchestrator. Other services running in the same Nexus Dashboard cluster may have lower RTT requirements and you must always use the lowest RTT requirement when deploying multiple services in the same Nexus Dashboard cluster. We recommend consulting the *Cisco Nexus Dashboard Deployment Guide* for more information.

When Nexus Dashboard Orchestrator service is deployed in Nexus Dashboard, it uses each of the two networks for different purposes as shown in the following table:

NDO Traffic Type	Nexus Dashboard Network
Any traffic to and from: • Cisco APIC • Cisco NDFC • Any other remote devices or controllers	Data network
Intra-cluster communication	Data network
Audit log streaming (Splunk/syslog)	Management network
Remote backup	Management network

Nexus Dashboard Cluster Sizing and Services Cohosting

Nexus Dashboard supports co-hosting of services. Depending on the type and number of services you choose to run, you may be required to deploy additional worker nodes in your cluster. For cluster sizing information and recommended number of nodes based on specific use cases, see the Cisco Nexus Dashboard Capacity Planning tool.

If you plan to host other services in addition to the Nexus Dashboard Orchestrator, ensure that you deploy and configure additional Nexus Dashboard nodes based on the cluster sizing tool recommendation, as described in the *Cisco Nexus Dashboard User Guide*, which is also available directly from the Nexus Dashboard GUI.

Note This release of Nexus Dashboard Orchestrator can be co-hosted with other services on physical or virtual (ESX) Nexus Dashboard clusters only. If you are deploying the Nexus Dashboard Orchestrator service in a virtual (KVM) or cloud Nexus Dashboard cluster, you must not install other services in the same cluster.

Network Time Protocol (NTP) and Domain Name System (DNS)

The Nexus Dashboard nodes require valid DNS and NTP servers for all deployments and upgrades.

Lack of valid DNS connectivity (such as if using an unreachable or a placeholder IP address) can prevent the system from deploying or upgrading successfully.

Note Nexus Dashboard acts as both a DNS client and resolver. It uses an internal Core DNS server which acts as DNS resolver for internal services. It also acts as a DNS client to reach external hosts within the intranet or the Internet, hence it requires an external DNS server to be configured.

Additionally, Nexus Dashboard does not support DNS servers with wildcard records.

Hardware Requirements For ACI Fabrics

Spine Switch Requirements

Multi-Site requires second generation (Cloud Scale) spine switches for intersite connectivity. All Cloud Scale spine switches supported by a given ACI release are supported by Nexus Dashboard Orchestrator.

Nexus 9000 first generation switches are not supported for Multi-Site intersite connectivity, but can still be used within a single fabric as long as that fabric is running an APIC release prior to 5.0(1).

Refer to the ACI-mode Switches Hardware Support Matrix for the complete list of supported spines for each release.

Leaf Switch Requirements

Multi-Site has no dependency on the fabrics' leaf switches and as such supports the same leaf switch models as the Cisco APIC. The full list of supported hardware is available in the ACI-mode Switches Hardware Support Matrix.

IPN Connectivity Across Sites

The following figure shows how spine switches supported with Multi-Site are connected to the intersite network.

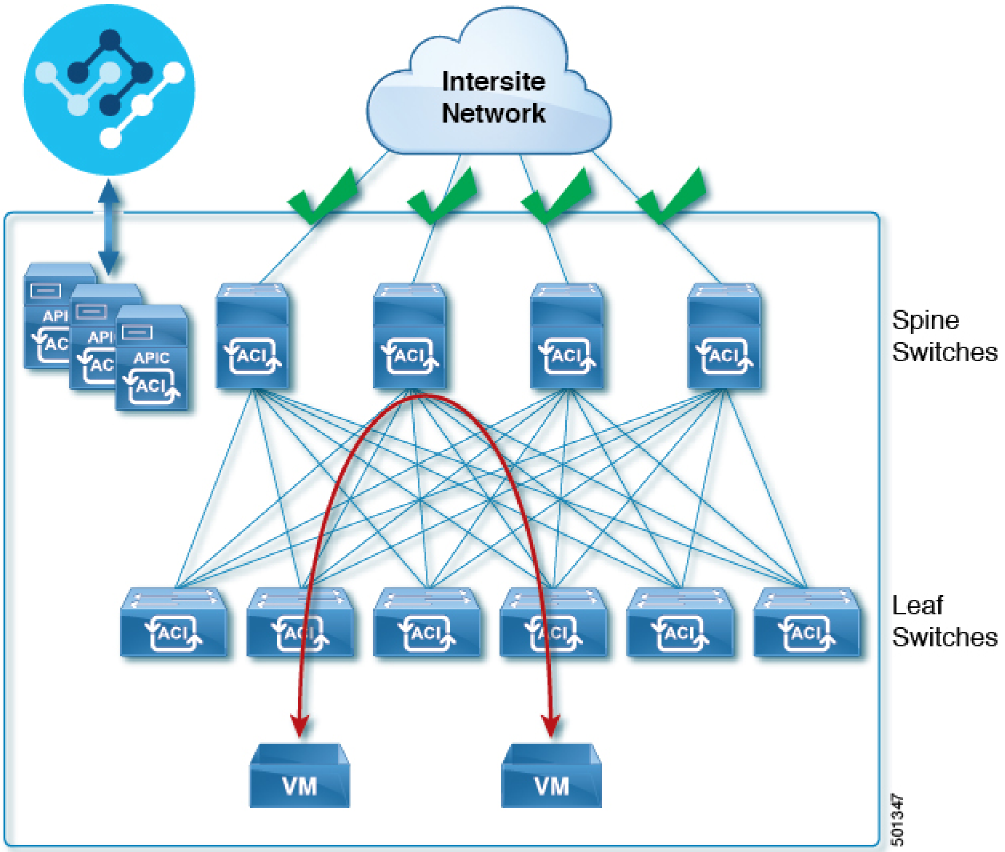

You can choose to mix spine switches supported by Multi-Site with switches that are not supported within the same Cisco APIC fabric, but only the supported switches can connect to the intersite network as shown in the following figure.

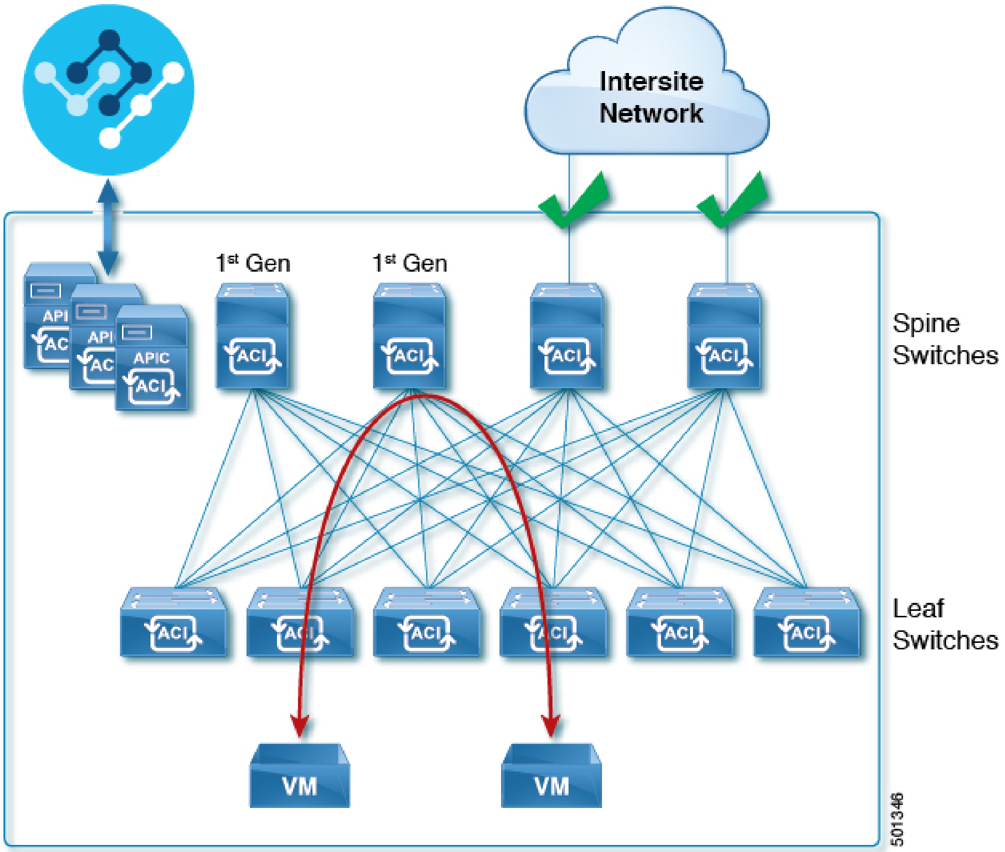

Hardware Requirements For NDFC Fabrics

Border Gateways Requirements

The following table summarizes the hardware requirements for EVPN Multi-Site Architecture:

- Cisco Nexus 9300 EX platform
- Cisco Nexus 9300 FX platform
- Cisco Nexus 9300 FX2 platform
- Cisco Nexus 9300-GX platform
- Cisco Nexus 9332C platform
- Cisco Nexus 9364C platform
- Cisco Nexus 9500 platform with X9700-EX line card
- Cisco Nexus 9500 platform with X9700-FX line card

The hardware requirements for the site-internal BGP Route Reflector (RR) and VTEP of a VXLAN BGP EVPN site remain the same as those without the EVPN Multi-Site Border Gateways (BGW). This document does not cover the hardware and software requirements for the VXLAN EVPN site-internal network.

Installing Cisco Nexus Dashboard Orchestrator Service Using App Store

This section describes how to install the Cisco Nexus Dashboard Orchestrator service in an existing Cisco Nexus Dashboard cluster.

Before you begin

- Ensure that you meet the requirements and guidelines that are described in Prerequisites and Guidelines, on page 4.
- The Cisco DC App Center must be reachable from the Cisco Nexus Dashboard through the Management Network directly or using a proxy configuration. Cisco Nexus Dashboard proxy configuration is described in the *Nexus Dashboard User Guide*.

  If you are unable to establish the connection to the DC App Center, skip this section and follow the steps that are described in Installing Nexus Dashboard Orchestrator Service Manually, on page 9.
- The App Store allows you to install the latest version of the service only.

  If you want to install a version that's different from the one in the App Store, you must follow the steps in Installing Nexus Dashboard Orchestrator Service Manually, on page 9 instead.

Step 1 Log in to your Cisco Nexus Dashboard GUI.

When deploying a service, you must install it in only one of the Cisco Nexus Dashboard nodes, the service will be replicated to the other nodes in the cluster automatically. So you can sign in to any one of your Cisco Nexus Dashboard nodes using its management IP address.

Step 2 From the drop-down in the top navigation menu, select **Admin Console**.

You must have `admin` privileges to deploy services.

Step 3 Navigate to the App Store and choose the Cisco Nexus Dashboard Orchestrator app.

a) From the left navigation menu, select **Operate** > **Services**.
b) Select the **App Store** tab.
c) In the **Nexus Dashboard Orchestrator** tile, click **Install**.

Step 4 In the License Agreement window that opens, click **Agree and Download**.

Step 5 Wait for the service to be downloaded to your Cisco Nexus Dashboard and installed.

Step 6 Enable and launch the app.

After installation is complete, click the Cisco Nexus Dashboard Orchestrator tile to see the list of available versions, choose the latest version to upgrade.

a) In the **Services** page, select the **Installed Services** tab.
b) In the **Nexus Dashboard Orchestrator** tile, click **Enable**.

   Once the service is enabled, the **Enable** button changes to **Open**.
c) In the **Nexus Dashboard Orchestrator** tile, click **Open**.

The single sign-on (SSO) feature allows you to log in to the service using the same credentials as you used for the Cisco Nexus Dashboard.

Installing Nexus Dashboard Orchestrator Service Manually

This section describes how to manually upload and install Cisco Nexus Dashboard Orchestrator service in an existing Cisco Nexus Dashboard cluster.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 4.

Step 1 Download the Cisco Nexus Dashboard Orchestrator service.

a) Browse to the Nexus Dashboard Orchestrator page on DC App Center:

https://dcappcenter.cisco.com/nexus-dashboard-orchestrator.html

b) From the **Version** drop-down, choose the version that you want to install and click **Download**.

c) Click **Agree and download** to accept the license agreement and download the image.

Step 2 Log in to your Cisco Nexus Dashboard GUI.

When deploying a service, you must install it in only one of the Cisco Nexus Dashboard nodes, the service will be replicated to the other nodes in the cluster automatically. So you can sign in to any one of your Cisco Nexus Dashboard nodes using its management IP address.

Step 3 From the drop-down in the top navigation menu, select **Admin Console**.

You must have `admin` privileges to deploy services.

Step 4 Manually upload the image.

a) From the left navigation menu, select **Operate** > **Services**.

b) Select the **Installed Services** tab.

c) In the top right of the main pane, select **Actions** > **Upload Service**.

d) Choose the location of the image.

If you downloaded the service image to your system, choose **Local**.

If you are hosting the image on a server, choose **Remote**.

e) Choose the image file.

If you chose **Local** in the previous substep, click **Select File** and locate the image that you downloaded.

If you chose **Remote**, provide the full URL to the image file, for example `http://<ip-address>:<port>/<full-path>/cisco-mso-<version>.nap`.

f) Click **Upload** to add the service to the cluster.

Step 5 Wait for the service to be downloaded to your Cisco Nexus Dashboard and installed.

Step 6 Enable and launch the app.

After installation is complete, click Cisco Nexus Dashboard Orchestrator tile to see the list of available versions, choose the latest version to upgrade.

a) In the **Services** page, select the **Installed Services** tab.
b) In the **Nexus Dashboard Orchestrator** tile, click **Enable**.

   Once the service is enabled, the **Enable** button will change to **Open**.

c) In the **Nexus Dashboard Orchestrator** tile, click **Open**.

   The single sign-on (SSO) feature allows you to sign in to the service using the same credentials as you used for the Cisco Nexus Dashboard.

PART I

Day-0 Operations for APIC and Cloud Network Controller Sites

CHAPTER 3

Preparing Cisco APIC Sites

Pod Profile and Policy Group

In each site's APIC, you must have one Pod profile with a Pod policy group. If your site does not have a Pod policy group you must create one. Typically, these settings will already exist as you will have configured them when you first deployed the fabric.

Step 1 Log in to the site's APIC GUI.

Step 2 Check that the Pod profile contains a Pod policy group.

Navigate to **Fabric** > **Fabric Policies** > **Pods** > **Profiles** > **Pod Profile default**.

Step 3 If necessary, create a Pod policy group.

a) Navigate to **Fabric** > **Fabric Policies** > **Pods** > **Policy Groups**.
b) Right-click **Policy Groups** and select **Create Pod Policy Group**.
c) Enter the appropriate information and click **Submit**.

Step 4 Assign the new Pod policy group to the default Pod profile.

a) Navigate to **Fabric** > **Fabric Policies** > **Pods** > **Profiles** > **Pod Profile default**
b) Select the default profile.
c) Choose the new pod policy group and click **Update**.

Configuring Fabric Access Policies for All APIC Sites

Before your APIC fabrics can be added to and managed by the Nexus Dashboard Orchestrator, there is a number of fabric-specific access policies that you must configure on each site.

Configuring Fabric Access Global Policies

This section describes the global fabric access policy configurations that must be created for each APIC site before it can be added to and managed by the Nexus Dashboard Orchestrator.

Step 1 Log in directly to the site's APIC GUI.

Step 2 From the main navigation menu, select **Fabric** > **Access Policies**.

You must configure a number of fabric policies before the site can be added to the Nexus Dashboard Orchestrator. From the APIC's perspective, this is something you do just like you would if you were connecting a bare-metal host, where you would configure domains, AEPs, policy groups, and interface selectors; you must configure the same options for connecting the spine switch interfaces to the inter-site network for all the sites that will be part of the same Multi-Site domain.

Step 3 Specify the VLAN pool.

The first thing you configure is the VLAN pool. We use Layer 3 sub-interfaces tagging traffic with VLAN-4 to connect the spine switches to the inter-site network.

a) In the left navigation tree, browse to **Pools** > **VLAN**.
b) Right-click the **VLAN** category and choose **Create VLAN Pool**.

In the **Create VLAN Pool** window, specify the following:

- For the **Name** field, specify the name for the VLAN pool, for example `msite`.
- For **Allocation Mode**, specify `Static Allocation`.
- And for the **Encap Blocks**, specify just the single VLAN 4. You can specify a single VLAN by entering the same number in both **Range** fields.

Step 4 Configure Attachable Access Entity Profiles (AEP).

a) In the left navigation tree, browse to **Global Policies** > **Attachable Access Entity Profiles**.
b) Right-click the **Attachable Access Entity Profiles** category and choose **Create Attachable Access Entity Profiles**.

In the **Create Attachable Access Entity Profiles** window, specify the name for the AEP, for example `msite-aep`.

c) Click **Next** and **Submit**

No additional changes, such as interfaces, are required.

Step 5 Configure domain.

The domain you configure is what you will select from the Nexus Dashboard Orchestrator when adding this site.

a) In the left navigation tree, browse to **Physical and External Domains** > **External Routed Domains**.
b) Right-click the **External Routed Domains** category and choose **Create Layer 3 Domain**.

In the **Create Layer 3 Domain** window, specify the following:

- For the **Name** field, specify the name the domain, for example `msite-l3`.
- For **Associated Attachable Entity Profile**, select the AEP you created in Step 4.
- For the **VLAN Pool**, select the VLAN pool you created in Step 3.

c) Click **Submit**.

No additional changes, such as security domains, are required.

What to do next

After you have configured the global access policies, you must still add interfaces policies as described in Configuring Fabric Access Interface Policies, on page 15.

Configuring Fabric Access Interface Policies

This section describes the fabric access interface configurations that must be done for the Nexus Dashboard Orchestrator on each APIC site.

Before you begin

You must have configured the global fabric access policies, such as VLAN Pool, AEP, and domain, in the site's APIC, as described in Configuring Fabric Access Global Policies, on page 14.

Step 1 Log in directly to the site's APIC GUI.

Step 2 From the main navigation menu, select **Fabric** > **Access Policies**.

In addition to the VLAN, AEP, and domain you have configured in previous section, you must also create the interface policies for the fabric's spine switch interfaces that connect to the Inter-Site Network (ISN).

Step 3 Configure a spine policy group.

a) In the left navigation tree, browse to **Interface Policies** > **Policy Groups** > **Spine Policy Groups**.

This is similar to how you would add a bare-metal server, except instead of a Leaf Policy Group, you are creating a Spine Policy Group.

b) Right-click the **Spine Policy Groups** category and choose **Create Spine Access Port Policy Group**.

In the **Create Spine Access Port Policy Group** window, specify the following:

- For the **Name** field, specify the name for the policy group, for example `Spine1-PolGrp`.
- For the **Link Level Policy** field, specify the link policy used between your spine switch and the ISN.
- For **CDP Policy**, choose whether you want to enable CDP.
- For the **Attached Entity Profile**, select the AEP you have configured in previous section, for example `msite-aep`.

c) Click **Submit**.

No additional changes, such as security domains, are required.

Step 4 Configure a spine profile.

a) In the left navigation tree, browse to **Interface Policies** > **Profiles** > **Spine Profiles**.

b) Right-click the **Spine Profiles** category and choose **Create Spine Interface Profile**.

In the **Create Spine Interface Profile** window, specify the following:

- For the **Name** field, specify the name for the profile, for example `Spine1-ISN`.
- For **Interface Selectors**, click the **+** sign to add the port on the spine switch that connects to the ISN. Then in the **Create Spine Access Port Selector** window, provide the following:
  - For the **Name** field, specify the name for the port selector, for example `Spine1-ISN`.
  - For the **Interface IDs**, specify the switch port that connects to the ISN, for example `5/32`.
  - For the **Interface Policy Group**, choose the policy group you created in the previous step, for example `Spine1-PolGrp`.

  Then click **OK** to save the port selector.

c) Click **Submit** to save the spine interface profile.

Step 5 Configure a spine switch selector policy.

a) In the left navigation tree, browse to **Switch Policies** > **Profiles** > **Spine Profiles**.

b) Right-click the **Spine Profiles** category and choose **Create Spine Profile**.

In the **Create Spine Profile** window, specify the following:

- For the **Name** field, specify the name for the profile, for example `Spine1`.
- For **Spine Selectors**, click the **+** to add the spine and provide the following:
  - For the **Name** field, specify the name for the selector, for example `Spine1`.
  - For the **Blocks** field, specify the spine node, for example `201`.

c) Click **Update** to save the selector.

d) Click **Next** to proceed to the next screen.

e) Select the interface profile you have created in the previous step

For example `Spine1-ISN`.

f) Click **Finish** to save the spine profile.

Configuring Sites That Contain Remote Leaf Switches

Multi-Site architecture supports APIC sites with Remote Leaf switches. The following sections describe guidelines, limitations, and configuration steps required to allow Nexus Dashboard Orchestrator to manage these sites.

Remote Leaf Guidelines and Limitations

If you want to add an APIC site with a Remote Leaf to be managed by the Nexus Dashboard Orchestrator, the following restrictions apply:

- You must upgrade your Cisco APIC to Release 4.2(4) or later.
- Only physical Remote Leaf switches are supported in this release

- Only -EX and -FX or later switches are supported as Remote Leaf switches for use with Multi-Site
- Remote Leaf is not supported with back-to-back connected sites without IPN switches
- Remote Leaf switches in one site cannot use another site's L3Out
- Stretching a bridge domain between one site and a Remote Leaf in another site is not supported

You must also perform the following tasks before the site can be added to and managed by the Nexus Dashboard Orchestrator:

- You must enable Remote Leaf direct communication and configure routable subnets directly in the site's APIC, as described in the following sections.
- You must add the routable IP addresses of Cisco APIC nodes in the DHCP-Relay configuration applied on the interfaces of the Layer 3 routers connecting to the Remote Leaf switches.

  The routable IP address of each APIC node is listed in the **Routable IP** field of the **System** > **Controllers** > **<controller-name>** screen of the APIC GUI.

Configuring Routable Subnets for Remote Leaf Switches

Before you can add a site that contains one or more Remote Leaf switches to the Nexus Dashboard Orchestrator, you must configure routable subnets for the pod with which the Remote Leaf nodes are associated.

Step 1 Log in directly to the site's APIC GUI.

Step 2 From the menu bar, select **Fabric** > **Inventory**.

Step 3 In the Navigation pane, click **Pod Fabric Setup Policy**.

Step 4 In the main pane, double-click the pod where you want to configure the subnets.

Step 5 In the **Routable Subnets** area, click the **+** sign to add a subnet.

Step 6 Enter the **IP** and **Reserve Address Count**, set the state to `Active` or `Inactive`, then click **Update** to save the subnet.

When configuring routable subnets, you must provide a netmask between `/22` and `/29`.

Step 7 Click **Submit** to save the configuration.

Enabling Direct Communication for Remote Leaf Switches

Before you can add a site that contains one or more Remote Leaf switches to the Nexus Dashboard Orchestrator, you must configure direct remote leaf communication for that site. Additional information about remote leaf direct communication feature is available in the *Cisco APIC Layer 3 Networking Configuration Guide*. This section outlines the steps and guidelines specific to the integration with Multi-Site.

Note Once you enable Remote Leaf switch direct communication, the switches will function in the new mode only

Step 1 Log in directly to the site's APIC.

Step 2 Enable direct traffic forwarding for Remote Leaf switches.

a) From the menu bar, navigate to **System** > **System Settings**.

b) From the left side bar, select **Fabric Wide Setting**.

c) Check the **Enable Remote Leaf Direct Traffic Forwarding** checkbox.

Note You cannot disable this option after you enable it.

d) Click **Submit** to save the changes.

Cisco Mini ACI Fabrics

Cisco Multi-Site supports Cisco Mini ACI fabrics as typical on-premises sites without requiring any additional configuration. This section provides a brief overview of Mini ACI fabrics, detailed info on deploying and configuring this type of fabrics is available in *Cisco Mini ACI Fabric and Virtual APICs*.

Cisco ACI, Release 4.0(1) introduced Mini ACI Fabric for small scale deployment. Mini ACI fabric works with Cisco APIC cluster consisting of one physical APIC and two virtual APICs (vAPIC) running in virtual machines. This reduces the physical footprint and cost of the APIC cluster, allowing ACI fabric to be deployed in scenarios with limited rack space or initial budget, such as a colocation facility or a single-room data center, where a full-scale ACI installations may not be practical due to physical footprint or initial cost.

The following diagram shows an example of a mini Cisco ACI fabric with a physical APIC and two virtual APICs (vAPICs):

Figure 1: Cisco Mini ACI Fabric

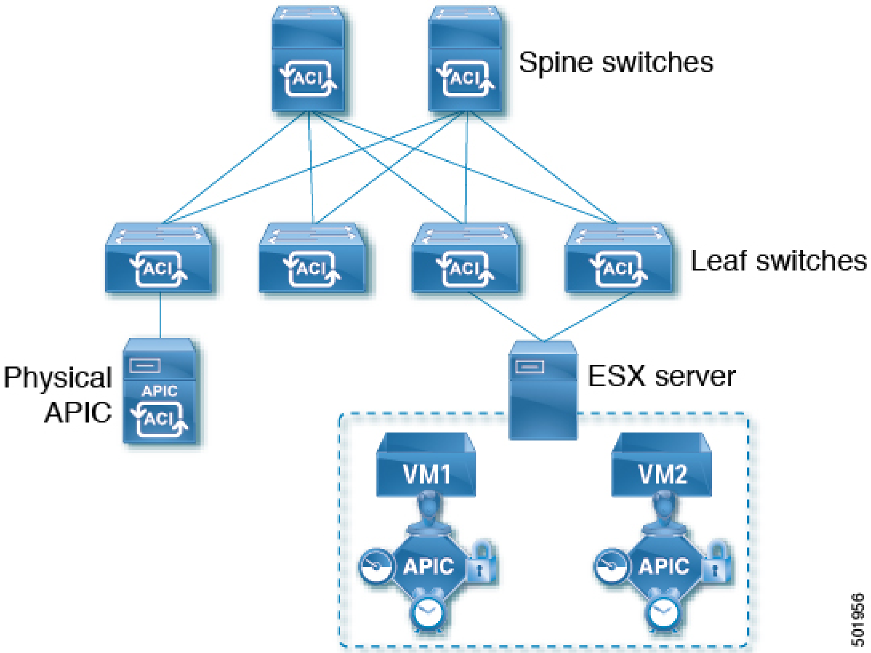

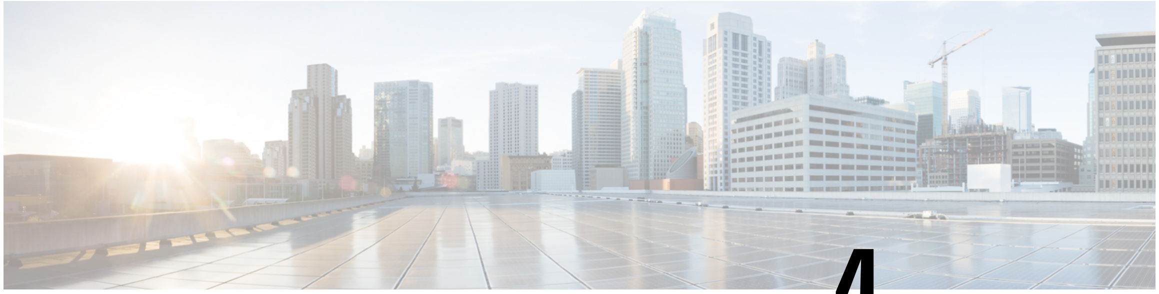

CHAPTER 4

Adding and Deleting Sites

Cisco NDO and APIC Interoperability Support

Cisco Nexus Dashboard Orchestrator (NDO) does not require a specific version of APIC to be running in all sites. The APIC clusters in each site as well as the NDO itself can be upgraded independently of each other and run in mixed operation mode as long as the fabric can be on-boarded to the Nexus Dashboard where the Nexus Dashboard Orchestrator service is installed. As such, we recommend that you always upgrade to the latest release of the Nexus Dashboard Orchestrator.

However, keep in mind that if you upgrade the NDO before upgrading the APIC clusters in one or more sites, some of the new NDO features may not yet be supported by an earlier APIC release. In that case a check is performed on each template to ensure that every configured option is supported by the target sites.

The check is performed when you save a template or deploy a template. If the template is already assigned to a site, any unsupported configuration options will not be saved; if the template is not yet assigned, you will be able to assign it to a site, but not be able to save or deploy the schema if it contains configuration unsupported by that site.

In case an unsupported configuration is detected, an error message will show, for example: `This APIC site version <site-version> is not supported by NDO. The minimum version required for this <feature> is <required-version> or above.`

The following table lists the features and the minimum required APIC release for each one:

Note While some of the following features are supported on earlier Cisco APIC releases, Release 4.2(4) is the earliest release that can be on-boarded to the Nexus Dashboard and managed by this release of Nexus Dashboard Orchestrator.

Feature	Minimum APIC Version
ACI Multi-Pod Support	Release 4.2(4)

Feature	Minimum APIC Version
Service Graphs (L4-L7 Services)	Release 4.2(4)
External EPGs	Release 4.2(4)
ACI Virtual Edge VMM Support	Release 4.2(4)
DHCP Support	Release 4.2(4)
Consistency Checker	Release 4.2(4)
vzAny	Release 4.2(4)
Host Based Routing	Release 4.2(4)
CloudSec Encryption	Release 4.2(4)
Layer 3 Multicast	Release 4.2(4)
MD5 Authentication for OSPF	Release 4.2(4)
EPG Preferred Group	Release 4.2(4)
Intersite L3Out	Release 4.2(4)
EPG QoS Priority	Release 4.2(4)
Contract QoS Priority	Release 4.2(4)
Single Sign-On (SSO)	Release 5.0(1)
Multicast Rendezvous Point (RP) Support	Release 5.0(1)
Transit Gateway (TGW) support for AWS and Azure Sites	Release 5.0(1)
SR-MPLS Support	Release 5.0(1)
Cloud LoadBalancer High Availability Port	Release 5.0(1)
Service Graphs (L4-L7 Services) with UDR	Release 5.0(2)
3rd Party Device Support in Cloud	Release 5.0(2)
Cloud Loadbalancer Target Attach Mode Feature	Release 5.1(1)
Support security and service insertion in Azure for non-ACI networks reachable through Express Route	Release 5.1(1)
CSR Private IP Support	Release 5.1(1)
Extend ACI policy model and automation for Cloud native services in Azure	Release 5.1(1)

Feature	Minimum APIC Version
Flexible segmentation through multiple VRF support within a single VNET for Azure	Release 5.1(1)
Private Link automation for Azure PaaS and third-party services	Release 5.1(1)
Openshift 4.3 IPI on Azure with ACI-CNI	Release 5.1(1)
Cloud Site Underlay Configuration	Release 5.2(1)

Adding Cisco ACI Sites

This section describes how to add a Cisco APIC or Cloud Network Controller site using the Cisco Nexus Dashboard GUI and then enable that site to be managed by Cisco Nexus Dashboard Orchestrator.

Before you begin

- If you are adding on-premises ACI site, you must have completed the site-specific configurations in each site's APIC, as described in previous sections in this chapter.
- You must ensure that one or more sites you are adding are running Release 4.2(4) or later.

Step 1 Log in to your Cisco Nexus Dashboard and open the **Admin Console**.

Step 2 From the left navigation menu, choose **Operate** and click **Sites**..

Step 3 Choose **Add Site** and provide site information.

a) For **Site Type**, select **ACI** or **Cloud Network Controller** depending on the type of ACI fabric you are adding.

b) Provide the controller information.

- You must provide the **Host Name/IP Address**, **User Name**, and **Password.** for the APIC controller currently managing your ACI fabrics.

  Note For APIC fabrics, if you use the site with Cisco Nexus Dashboard Orchestrator service only, you can provide either the in-band or out-of-band IP address of the APIC. If you use the site with Cisco Nexus Dashboard Insights as well, you must provide the in-band IP address.

- For on-premises ACI sites managed by Cisco APIC, if you plan to use this site with Day-2 Operations applications such as Cisco Nexus Insights, you must also provide the **In-Band EPG** name that is used to connect the Cisco Nexus Dashboard to the fabric you are adding. Otherwise, if you use this site with Cisco Nexus Dashboard Orchestrator only, you can leave this field blank.

- For Cloud Network Controller sites, **Enable Proxy** if your cloud site is reachable through a proxy.

  Proxy must be already configured in your Cisco Nexus Dashboard's cluster settings. If the proxy is reachable through management network, a static management network route must also be added for the proxy IP address. For more information about proxy and route configuration, see Nexus Dashboard User Guide for your release.

c) Click **Save** to finish adding the site.

Currently, the sites are available in the Cisco Nexus Dashboard, but you still must enable them for Cisco Nexus Dashboard Orchestrator management as described in the following steps.

Step 4 Repeat the previous steps for any additional ACI or Cloud Network Controller sites.

Step 5 From the Cisco Nexus Dashboard's **Services** page, open the Cisco Nexus Dashboard Orchestrator service.

You are automatically signed in using the Cisco Nexus Dashboard user's credentials.

Step 6 In the Cisco Nexus Dashboard Orchestrator GUI, manage the sites.

a) From the left navigation menu, select **Sites**.

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

When managing the sites, you must provide a unique site ID for each site.

Removing Sites

This section describes how to disable site management for one or more sites using the Cisco Nexus Dashboard Orchestrator GUI. The sites remain present in the Cisco Nexus Dashboard.

Before you begin

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

Step 1 Open the Cisco Nexus Dashboard Orchestrator GUI.

You can open the NDO service from the Cisco Nexus Dashboard's **Service Catalog**. You are automatically signed in using the Cisco Nexus Dashboard user's credentials.

Step 2 Remove the site from all templates.

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

a) Navigate to **Configure > Tenant Template** > **Applications**.

b) Click a **Schema** that contains one or more templates that are associated with the site.

c) From the **Overview** drop-down, choose a template that's associated with the site that you want to remove.

d) From the **Actions** drop-down, choose **Add/Remove Sites** and uncheck the site that you want to remove.

This removes configurations that were deployed using this template to this site.

Note For nonstretched templates, you can choose to preserve the configurations that are deployed by the template to the sites by selecting **Actions** > **Dissociate Sites** instead. This option allows you to retain configurations that are deployed by NDO but no longer manage those objects from NDO.

e) Repeat this step for all templates associated with the site that you want to unmanage in this and all other schemas.

Step 3 Remove the site's underlay configuration.

a) From the left navigation menu, select **Configure** > **Site To Site Connectivity**.

b) In the main pane, click **Configure**.

c) In the left sidebar, select the site that you want to unmanage.

d) Click **View Details** to load site settings.

Figure 2:

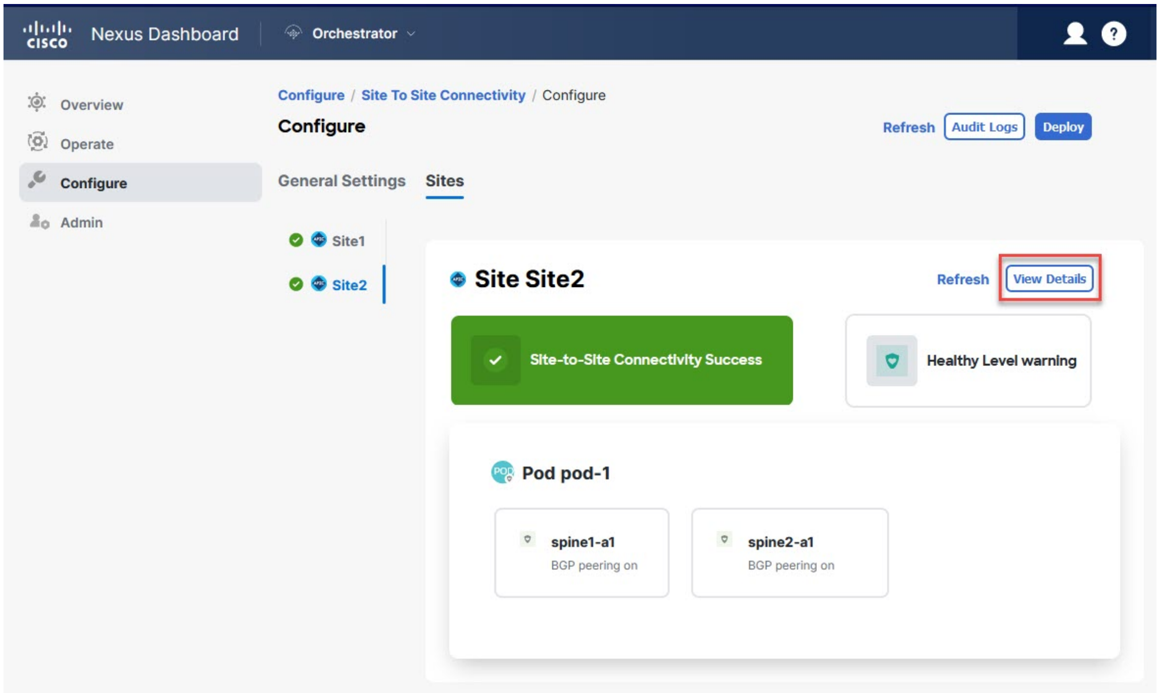

e) In right sidebar's **Inter-Site Connectivity** tab, disable the **Multi-Site** check box.

This disables EVPN peering between this site and other sites.

f) Click **Deploy** to deploy the changes to the site.

Step 4 In the Cisco Nexus Dashboard Orchestrator GUI, disable the sites.

a) From the left navigation menu, select **Sites**.

b) In the main pane, change the **State** from `Managed` to `Unmanaged` for the site that you want to unmanage.

Note If the site 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.

Step 5 Delete the site from Cisco Nexus Dashboard.

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

Note The site must not be currently in use by any of the services that are installed in your Cisco Nexus Dashboard cluster.

a) In the top navigation bar, click the **Home** icon to return to the Cisco Nexus Dashboard GUI.

b) From the left navigation menu of the Cisco Nexus Dashboard GUI, select **Operate** > **Sites**.

c) Select one or more sites that you want to delete.

d) In the top right of the main pane, select **Actions** > **Delete Site**.

e) Provide the site's sign-in information and click **OK**.

The site will be removed from the Cisco 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 site controller's GUI from the NDO's **Operate** > **Sites** screen by selecting the actions (. . .) menu next to the site 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.

CHAPTER 5

Configuring Infra General Settings

Infra Configuration Dashboard

The **Config** > **Site To Site Connectivity** page displays a summary of all sites and intersite connectivity in your Cisco Nexus Dashboard Orchestrator deployment and contains the following information:

Figure 3: Infra Configuration Overview

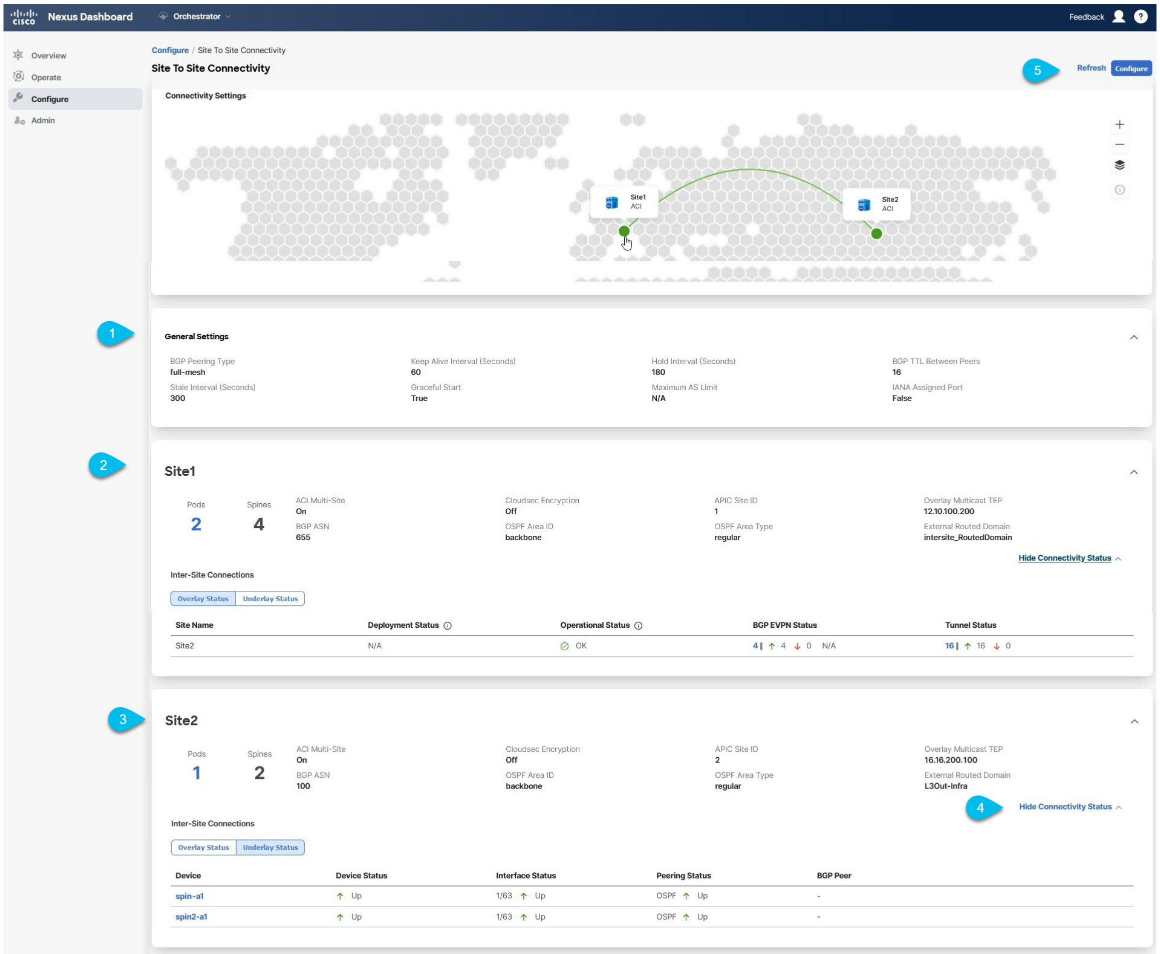

General Settings

BGP Peering Type: full-mesh
Keep Alive Interval (Seconds): 60
Hold Interval (Seconds): 180
BGP TTL Between Peers: 16
Stale Interval (Seconds): 300
Graceful Start: True
Maximum AS Limit: N/A
IANA Assigned Port: False

Site1

Pods: 2
Spines: 4
ACI Multi-Site: On
BGP ASN: 655
Cloudsec Encryption: Off
OSPF Area ID: backbone
APIC Site ID: 1
OSPF Area Type: regular
Overlay Multicast TEP: 12.10.100.200
External Routed Domain: intersite_RoutedDomain

Inter-Site Connections

Site Name	Deployment Status	Operational Status	BGP EVPN Status	Tunnel Status
Site2	N/A	✓ OK	4 | ↑ 4 ↓ 0 N/A	16 | ↑ 16 ↓ 0

Site2

Pods: 1
Spines: 2
ACI Multi-Site: On
BGP ASN: 100
Cloudsec Encryption: Off
OSPF Area ID: backbone
APIC Site ID: 2
OSPF Area Type: regular
Overlay Multicast TEP: 16.16.200.100
External Routed Domain: L3Out-Infra

Inter-Site Connections

Device	Device Status	Interface Status	Peering Status	BGP Peer
spin-a1	↑ Up	1/63 ↑ Up	OSPF ↑ Up	-
spin2-a1	↑ Up	1/63 ↑ Up	OSPF ↑ Up	-

1. The **General Settings** tile displays information about BGP peering type and its configuration.

   This is described in detail in the next section.

2. The **On-Premises** tiles display information about every on-premises site that is part of your Multi-Site domain along with their number of Pods and spine switches, OSPF settings, and overlay IPs.

   You can click the **Pods** tile that displays the number of Pods in the site to show information about the Overlay Unicast TEP addresses of each Pod.

   This is described in detail in Configuring Infra for Cisco APIC Sites, on page 33.

3. The **Cloud** tiles display information about every cloud site that is part of your Multi-Site domain along with their number of regions and basic site information.

   This is described in detail in Configuring Infra for Cisco Cloud Network Controller Sites, on page 41.

4. You can click **Show Connectivity Status** to display intersite connectivity details for a specific site.

5. You can use the **Configure** button to navigate to the intersite connectivity configuration, which is described in detail in the following sections.

The following sections describe the steps necessary to configure the general fabric Infra settings. Fabric-specific requirements and procedures are described in the following chapters based on the specific type of fabric that you are managing.

Before you proceed with Infra configuration, you must have configured and added the sites as described in previous sections.

In addition, any infrastructure changes such as adding and removing spine switches or spine node ID changes require a Cisco Nexus Dashboard Orchestrator fabric connectivity information refresh described in the Refreshing Site Connectivity Information, on page 33 as part of the general Infra configuration procedures.

Partial Mesh Intersite Connectivity

In addition to full mesh connectivity where you configure intersite connectivity from every site managed by your Nexus Dashboard Orchestrator to every other site, this release also supports partial mesh configuration. In partial mesh configuration, you can manage sites in standalone mode with no intersite connectivity to any other site or limit the intersite configuration to only a subset of other sites in your Multi-Site domain.

Prior to Nexus Dashboard Orchestrator, Release 3.6(1), you could stretch templates between sites and refer to policies from other templates, which were deployed to other sites, even if the intersite connectivity between those sites was not configured, resulting in intended traffic flow between the sites to not work.

Beginning with release 3.6(1), the Orchestrator will allow you to stretch template and remote reference policies from other templates (deployed on other sites) between two or more sites only if the intersite connectivity between those sites is properly configured and deployed.

When configuring site infra for Cisco APIC and Cisco Cloud Network Controller sites as described in the following sections, for each site you can explicitly choose to which other sites infra connectivity will be established and provide that configuration information only.

Partial Mesh Connectivity Guidelines

When configuring partial mesh connectivity, consider the following guidelines:

- Partial mesh connectivity is supported between two cloud sites or a cloud and on-premises site.

  Full mesh connectivity is automatically established between all on-premises sites.

- Partial mesh connectivity is supported using BGP-EVPN or BGP-IPv4 protocols.

  Note however that stretching a template is allowed only for sites that are connected using BGP-EVPN protocol. If you are using BGP-IPv4 to connect two or more sites, any template assigned to any of those sites can be deployed to one site only.

Configuring Infra: General Settings

This section describes how to configure general Infra settings for all the sites.

Note Some of the following settings apply to all sites, while others are required for specific type of sites (for example, Cloud Network Controller sites). Ensure that you complete all the required configurations in infra general settings before proceeding to the site-local settings specific to each site.

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 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 establishes peer connectivity with remote sites' border gateway switches.

  In `full-mesh` configuration, Cisco Nexus Dashboard Orchestrator uses the spine switches for ACI-managed fabrics and border gateways for NDFC-managed fabrics.

- `route-reflector`—The route-reflector option allows you to specify one or more control-plane nodes to which each site establishes MP-BGP EVPN sessions. The use of route-reflector nodes avoids creating MP-BGP EVPN full mesh adjacencies between all the sites that are managed by NDO.

  For ACI fabrics, the `route-reflector` option is effective only for fabrics that are part of the same BGP ASN.

c) In the **Keepalive Interval (Seconds)** field, enter the keepalive interval seconds.

We recommend keeping the default value.

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

We recommend keeping the default value.

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

We recommend keeping the default value.

f) Choose whether you want to turn on the **Graceful Helper** option.

g) Provide the **Maximum AS Limit**.

We recommend keeping the default value.

h) Provide the **BGP TTL Between Peers**.

We recommend keeping the default value.

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

If you do not have any Cloud Network Controller sites, this field will not be present in the UI.

This is OSPF area ID used by cloud sites for on-premises IPN peering.

j) (Optional) Enable **IANA Assigned Port** for CloudSec encryption.

By default, CloudSec uses a proprietary UDP port. This option allows you to configure CloudSec to use the official IANA-reserved port 8017 for CloudSec encryption between sites.

Note The IANA-reserved port is supported for Cisco APIC sites running release 5.2(4) or later.

To change this setting, CloudSec must be disabled on all sites. If you want to enable IANA reserved port, but already have CloudSec encryption that is enabled for one or more of your sites, disable CloudSec for all sites, enable **IANA Reserve UDP Port** option, then re-enable CloudSec for the required sites.

For detailed information and steps for configuring CloudSec, see the "CloudSec Encryption" chapter of the *Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics*.

Step 6 Provide the **IPN Devices** information.

If you do not plan to configure intersite connectivity between on-premises and cloud sites, you can skip this step.

When you configure intersite underlay connectivity between on-premises and cloud sites 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 site configuration screen, which is described in more detail in Configuring Infra: On-Premises Site Settings, on page 33.

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 **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 that you want to add.

Step 7 Provide the **External Devices** information.

If you do not have any Cloud Network Controller sites, this tab will not be present in the UI.

If you do not have any Cloud Network Controller sites in your Multi-Site domain or you do not plan to configure connectivity between cloud sites and branch routers or other external devices, you can skip this step.

The following steps describe how to provide information about any branch routers or external devices to which you want to configure connectivity from your cloud sites.

a) Select the **External Devices** tab.

   This tab will only be available if you have at least one cloud site in your Multi-Site domain.

b) Click **Add External Device**.

   The **Add External Device** dialogue opens.

c) Provide the **Name**, **IP Address**, and **BGP Autonomous System Number** for the device.

   The IP address that you provide will be used as the tunnel peer address from the Cloud Network Controller's CSRs, not the device's management IP address. The connectivity will be established over public Internet using IPsec.

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.

After you have added all the external devices, ensure to complete the next step to provide the IPsec tunnel subnet pools from with the internal IP addresses will be allocated for these tunnels.

Step 8 Provide the **IPsec Tunnel Subnet Pools** information.

If you do not have any Cloud Network Controller sites, this tab will not be present in the UI.

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 Cisco Nexus Dashboard Orchestrator. The Orchestrator, creates smaller subnets from these pools and allocates them to sites to be used for intersite 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 that are used for on-premises connectivity, which you previously configured in the Cloud Network Controller for intersite 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 **Site-Specific Subnet Pools**:

a) Select the **IsSec 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 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.

What to do next

After you have configured general infra settings, you must still provide additional information for site-specific configurations based on the type of sites (ACI, Cloud Network Controller, or NDFC) you are managing. Follow the instructions described in the following sections to provide site-specific infra configurations.

CHAPTER 6

Configuring Infra for Cisco APIC Sites

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 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.

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 4:

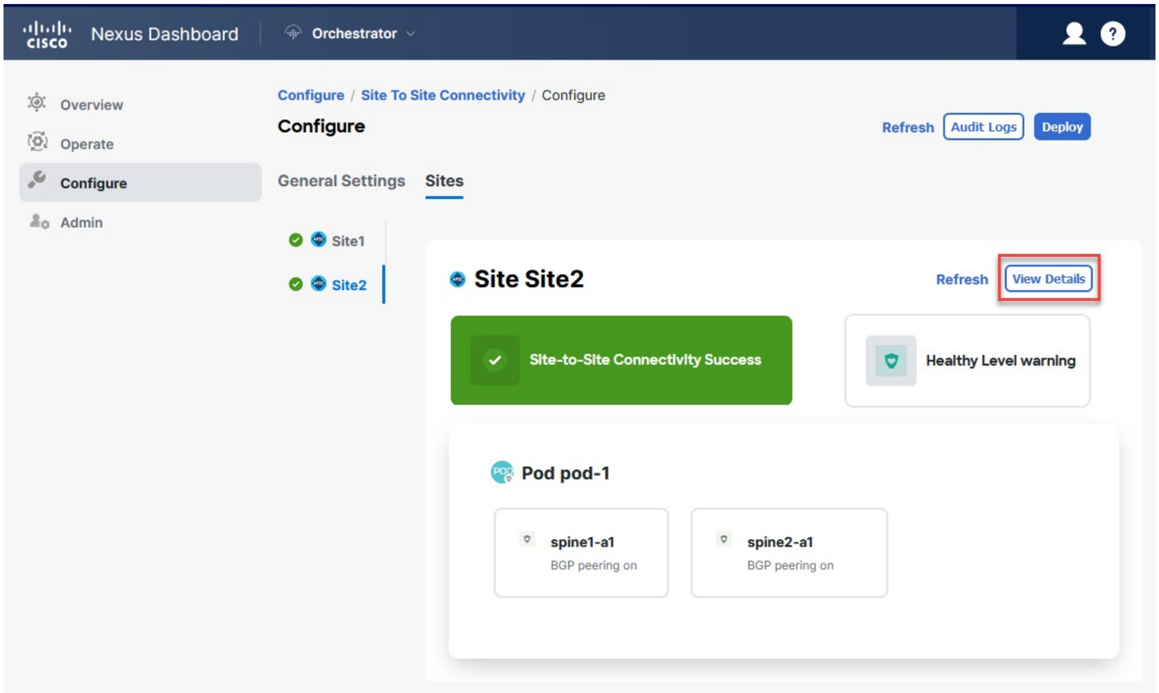

Step 6 Provide the **Inter-Site Connectivity** information.

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 37.

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 37.

   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 37.

   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, on page 27

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, on page 41.

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, on page 47

Configuring Infra: Pod Settings

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

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.

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 33.

  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.

CHAPTER 7

Configuring Infra for Cisco Cloud Network Controller Sites

Refreshing Cloud Site Connectivity Information

Any infrastructure changes, such as CSR and Region addition or removal, 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 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 discover any new or changed CSRs and regions.

Step 6 Finally, click **Yes** to confirm and load the connectivity information.

This discovers any new or removed CSRs and regions.

Step 7 Click **Deploy** to propagate the cloud site changes to other sites that have connectivity to it.

After you Refresh a cloud site's connectivity and CSRs or regions are added or removed, you must deploy infra configuration so other sites that have underlay connectivity to that cloud site get updated configuration.

Configuring Infra: Cloud Site Settings

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

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 cloud site.

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

Figure 5:

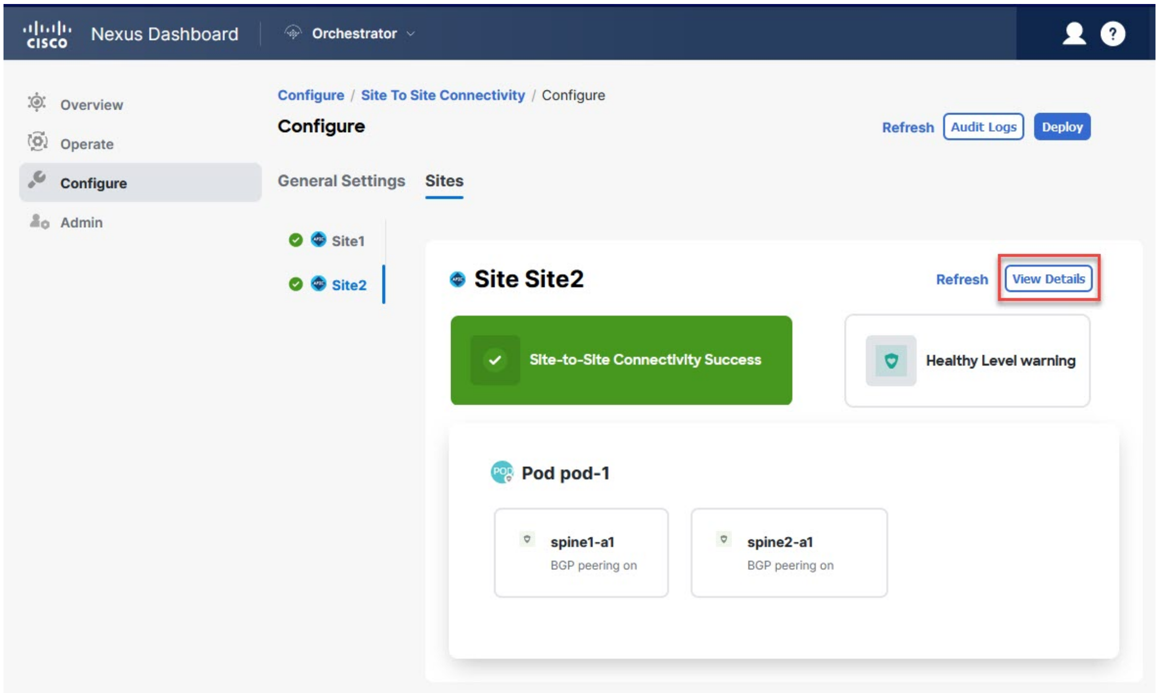

Step 6 Provide the general **Inter-Site Connectivity** information.

a) In the right ***<Site>*** **Settings** pane, select the **Inter-Site Connectivity** tab.

b) Enable the **Multi-Site** knob.

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

The overlay configuration will not be pushed to sites which do not have the underlay intersite connectivity that is established as described in the next step.

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

Step 7 Provide site-specific **Inter-Site Connectivity** information.

a) In the right properties sidebar for the cloud site, click **Add Site**.

The **Add Site** window opens.

b) Under **Connected to Site**, click **Select a Site** and select the site (for example, `Site2`) to which you want to establish connectivity from the site you are configuring (for example, `Site1`) .

When you select the remote site, the **Add Site** window updates to reflect both directions of connectivity: **Site1** > **Site2** and **Site2** > **Site1**.

c) In the **Site1** > **Site2** area, from the **Connection Type** drop-down, choose the type of connection between the sites.

The following options are available:

- `Public Internet`—Connectivity between the two sites is established through the Internet.

  This type is supported between any two cloud sites or between a cloud site and an on-premises site.

- `Private Connection`—Connectivity is established using a private connection between the two sites.

  This type is supported between a cloud site and an on-premises site.

- `Cloud Backbone`—Connectivity is established using cloud backbone.

  This type is supported between two cloud sites of the same type, such as Azure-to-Azure or AWS-to-AWS.

If you have multiple types of sites (on-premises, AWS, and Azure), different pairs of site can use different connection type.

d) Choose the **Protocol** that you want to use for connectivity between these two sites.

If using **BGP-EVPN** connectivity, you can optionally enable **IPSec** and choose which version of the Internet Key Exchange (IKE) protocol to use: IKEv1 (`Version 1`) or IKEv2 (`Version 1`) depending on your configuration.

- For `Public Internet` connectivity, IPsec is always enabled.
- For `Cloud Backbone` connectivity, IPsec is always disabled.
- For `Private Connection`, you can choose to enable or disable IPsec.

If using **BGP-IPv4** connectivity instead, you must provide an external VRF which will be used for route leaking configuration from the cloud site you are configuring.

After **Site1** > **Site2** connectivity information is provided, the **Site2** > **Site1** area will reflect the connectivity information in the opposite direction.

e) Click **Save** to save the intersite connectivity configuration.

When you save connectivity information from `Site1` to `Site2`, the reverse connectivity is automatically created from `Site2` to `Site1`, which you can see by selecting the other site and checking the **Inter-site Connectivity** information in the right sidebar.

f) Repeat this step to add intersite connectivity for other sites.

When you establish underlay connectivity from `Site1` to `Site2`, the reverse connectivity is done automatically for you.

However, if you also want to establish intersite connectivity from `Site1` to `Site3`, you must repeat this step for that site as well.

Step 8 Provide **External Connectivity** information.

If you do not plan to configure connectivity to external sites or devices that are not managed by NDO, you can skip this step.

Detailed description of an external connectivity use case is available in the *Configuring External Connectivity from Cloud CSRs Using Nexus Dashboard Orchestrator* document.

a) In the right ***<Site>* Settings** pane, select the **External Connectivity** tab.
b) Click **Add External Connection**.

The **Add External Connectivity** dialog opens.

c) From the **VRF** drop-down, select the VRF you want to use for external connectivity.

This is the VRF which will be used to leak the cloud routes. The **Regions** section displays the cloud regions that contain the CSRs to which this configuration be applied.

d) From the **Name** drop-down in the **External Devices** section, select the external device.

This is the external device that you added in the **General Settings** > **External Devices** list during general infra configuration and must already be defined as described in Configuring Infra: General Settings, on page 27.

e) From the **Tunnel IKE Version** drop-down, pick the IKE version that will be used to establish the IPsec tunnel between the cloud site's CSRs and the external device.

f) (Optional) From the **Tunnel Subnet Pool** drop-down, choose one of the named subnet pools.

Named subnet pools are used to allocate IP addresses for IPsec tunnels between cloud site CSRs and external devices. If you do not provide any **named** subnet pools here, the **external** subnet pool will be used for IP allocation.

Providing a dedicated subnet pool for external device connectivity is useful for cases where a specific subnet is already being used to allocate IP addresses to the external router and you want to continue to use those subnets for IPsec tunnels for NDO and cloud sites.

If you want to provide a specific subnet pool for this connectivity, it must already be created as described in Configuring Infra: General Settings, on page 27.

g) (Optional) In the **Pre-Shared Key** field, provide the custom keys that you want to use to establish the tunnel.

h) If necessary, repeat the previous substeps for any additional external devices you want to add for the same external connection (same VRF).

i) If necessary, repeat this step for any additional external connections (different VRFs).

There's a one-to-one relationship for tunnel endpoints between CSRs and external devices, so while you can create extra external connectivity using different VRFs, you cannot create extra connectivity to the same external devices.

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, on page 47

Recovering from Cloud Network Controller Site Downtime

When Cloud Network Controller (formerly Cloud APIC) instance/VM goes down for any reason while still being managed by NDO, you may be unable to undeploy or delete any existing templates associated with that cloud site. In this case, attempting to forcefully unmanage the site in NDO can cause stale configuration and deployment errors even if the site recovers.

To recover from this:

Step 1 Bring up the new Cloud Network Controller sites and reregister the cloud sites.

a) Log in to NDO.

b) Open the admin console.

c) Navigate to the **Operate** > **Sites** page.
d) From the action **(…)** menu next to the site you redeployed, choose **Edit Site**.
e) Check the "Reregister site" check box.
f) Provide the new site details.

   You must provide the new public IP address of site and sign-in credentials.

g) Click **Save** to reregister the site.

When the connectivity status of the site shows `UP`, the site IPs in NDO are also updated and the new sites are in 'managed' state.

Step 2 Undeploy the previously deployed templates for each schema.

a) Log in to NDO.
b) Navigate to **Configure** and select **Tenant Template** > **Applications**.
c) Click a schema with the deployed templates.
d) From the **Actions** menu next to the **Template Properties**, choose **Undeploy Template** and wait until the template is successfully undeployed.

Step 3 Refresh the site's infra configuration to ensure that the new Cisco Catalyst 8000V switches are added in NDO.

a) Navigate to **Configure** and select **SiteTo Site Connectivity**.
b) Click **Configure** at the top right of the screen.
c) Select the cloud site under the **Sites** panel and click **Refresh.**
d) Click **Deploy** on the top right of the screen and wait until all sites are successfully deployed.

Step 4 Redeploy all templates associated with this Cloud Network Controller site.

a) Navigate to **Configure** > **Tenant Templates** under the **Applications** tab.
b) Click a schema with the templates undeployed earlier.
c) Click **Deploy to Sites and** wait until the template is deployed.

CHAPTER 8

Deploying Infra Configuration for ACI Sites

Deploying Infra Configuration

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

Step 1 In the top right of the main pane, click **Deploy** and choose the appropriate option to deploy the configuration.

If you have configured 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 additional options may be available:

- **Deploy & Download IPN Device Config files:** Pushes the configuration to both the on-premises APIC site and the Cloud Network Controller 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 from the IPN devices to Cisco Cloud Services Router (CSR). A followup screen appears that allows you to select all or some of the configuration files to download.

- **Deploy & Download External Device Config files:** Pushes the configuration to both the Cloud Network Controller sites and enables the end-to-end interconnect between the cloud sites and external devices.

  In addition, this option downloads a zip file that contains configuration information that you will use to enable connectivity from external devices to the Cisco Cloud Services Router (CSR) deployed in your cloud sites. 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 from the IPN devices to Cisco Cloud Services Router (CSR) without deploying the configuration.

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

Step 2 In the confirmation window, click **Yes**.

The `Deployment started, refer to left menu for individual site deployment status` message will indicate that Infra configuration deployment began and you can verify each site's progress by the icon displayed next to the site's name in the left pane.

What to do next

The Infra overlay and underlay configuration settings are now deployed to all sites' controllers and cloud CSRs. The last remaining step is to configure your IPN devices with the tunnels for cloud CSRs as descrbied in Refreshing Site Connectivity Information, on page 33.

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 Network Controller sites.

By default, the Cisco Cloud Network Controller 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. If you have multiple on-premises IPsec devices, you will need to configure the same tunnels to the CSRs on each of the on-premises devices.

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 Nexus Dashboard Orchestrator as part of the procedures provided in Deploying Infra Configuration, on page 47.

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 Nexus Dashboard Orchestrator, but the following fields describe the important values for your specific deployment:

- *<first-csr-tunnel-ID>*—unique tunnel ID that you assign to this tunnel.
- *<first-csr-ip-address>*—public IP address of the third network interface of the first CSR.

  The destination of the tunnel depends on the type of underlay connectivity:

  - The destination of the tunnel is the public IP of the cloud router interface if the underlay is via public internet
  - The destination of the tunnel is the private IP of the cloud router interface if the underlay is via private connectivity, such as DX on AWS or ER on Azure

- *<first-csr-preshared-key>*—preshared key of the first CSR.

- *<onprem-device-interface>*—interface that is used for connecting to the Cisco Cloud Services Router 1000v deployed in Amazon Web Services.
- *<onprem-device-ip-address>*—IP address for the *<interface>* 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>*—peer tunnel IP address for the on-premises IPsec device to the first cloud CSR.
- *<process-id>* —OSPF process ID.
- *<area-id>*—OSPF area ID.

The following example shows intersite connectivity configuration using the IKEv2 protocol supported starting with Nexus Dashboard Orchestrator, Release 3.3(1) and Cloud Network Controller, Release 5.2(1). If you are using IKEv1, the IPN configuration file you downloaded form NDO may look slightly differently, but the principle remains the same.

```
crypto ikev2 proposal ikev2-proposal-default
    encryption aes-cbc-256 aes-cbc-192 aes-cbc-128
    integrity sha512 sha384 sha256 sha1
    group 24 21 20 19 16 15 14 2
exit

crypto ikev2 policy ikev2-policy-default
    proposal ikev2-proposal-default
exit

crypto ikev2 keyring key-ikev2-infra:overlay-1-<first-csr-tunnel-id>
    peer peer-ikev2-keyring
        address <first-csr-ip-address>
        pre-shared-key <first-csr-preshared-key>
    exit
exit

crypto ikev2 profile ikev2-infra:overlay-1-<first-csr-tunnel-id>
    match address local interface <onprem-device-interface>
    match identity remote address <first-csr-ip-address> 255.255.255.255
    identity local address <onprem-device-ip-address>
    authentication remote pre-share
    authentication local pre-share
    keyring local key-ikev2-infra:overlay-1-<first-csr-tunnel-id>
    lifetime 3600
    dpd 10 5 on-demand
exit

crypto ipsec transform-set infra:overlay-1-<first-csr-tunnel-id> esp-gcm 256
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-<first-csr-tunnel-id>
    set pfs group14
    set ikev2-profile ikev2-infra:overlay-1-<first-csr-tunnel-id>
    set transform-set infra:overlay-1-<first-csr-tunnel-id>
exit

interface tunnel 2001
    ip address <peer-tunnel-for-onprem-IPsec-to-first-CSR> 255.255.255.252
    ip virtual-reassembly
    tunnel source <onprem-device-interface>
    tunnel destination <first-csr-ip-address>
    tunnel mode ipsec ipv4
    tunnel protection ipsec profile infra:overlay-1-<first-csr-tunnel-id>
    ip mtu 1400
    ip tcp adjust-mss 1400
    ip ospf <process-id> area <area-id>
    no shut
exit
```

Example:

```
crypto ikev2 proposal ikev2-proposal-default
    encryption aes-cbc-256 aes-cbc-192 aes-cbc-128
    integrity sha512 sha384 sha256 sha1
    group 24 21 20 19 16 15 14 2
exit

crypto ikev2 policy ikev2-policy-default
    proposal ikev2-proposal-default
exit

crypto ikev2 keyring key-ikev2-infra:overlay-1-2001
    peer peer-ikev2-keyring
        address 52.12.232.0
        pre-shared-key 1449047253219022866513892194096727146110
    exit
exit

crypto ikev2 profile ikev2-infra:overlay-1-2001
    ! Please change GigabitEthernet1 to the appropriate interface
    match address local interface GigabitEthernet1
    match identity remote address 52.12.232.0 255.255.255.255
    identity local address 128.107.72.62
    authentication remote pre-share
    authentication local pre-share
    keyring local key-ikev2-infra:overlay-1-2001
    lifetime 3600
    dpd 10 5 on-demand
exit

crypto ipsec transform-set infra:overlay-1-2001 esp-gcm 256
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-2001
    set pfs group14
    set ikev2-profile ikev2-infra:overlay-1-2001
    set transform-set infra:overlay-1-2001
exit

! These tunnel interfaces establish point-to-point connectivity between the on-prem device and the
cloud Routers
! The destination of the tunnel depends on the type of underlay connectivity:
! 1) The destination of the tunnel is the public IP of the cloud Router interface if the underlay is
 via internet
! 2) The destination of the tunnel is the private IP of the cloud Router interface if the underlay
is via private
     connectivity like DX on AWS or ER on Azure

interface tunnel 2001
    ip address 5.5.1.26 255.255.255.252
    ip virtual-reassembly
    ! Please change GigabitEthernet1 to the appropriate interface
    tunnel source GigabitEthernet1
    tunnel destination 52.12.232.0
    tunnel mode ipsec ipv4
    tunnel protection ipsec profile infra:overlay-1-2001
    ip mtu 1400
    ip tcp adjust-mss 1400
    ! Please update process ID according with your configuration
    ip ospf 1 area 0.0.0.1
    no shut
exit
```

Step 4 Repeat the previous step for the 2nd and any additional CSRs that you need to configure.

Step 5 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
```

PART II

Day-0 Operations for NDFC Fabrics

CHAPTER 9

Adding and Deleting Sites

Adding Cisco NDFC Sites

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

Before you begin

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

Step 1 Log in to your Nexus Dashboard and open the **Admin Console**.

Step 2 From the left navigation menu, choose **Sites** and click **Add Site**..

Step 3 Provide site information.

a) For **Site 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 Sites** 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**:

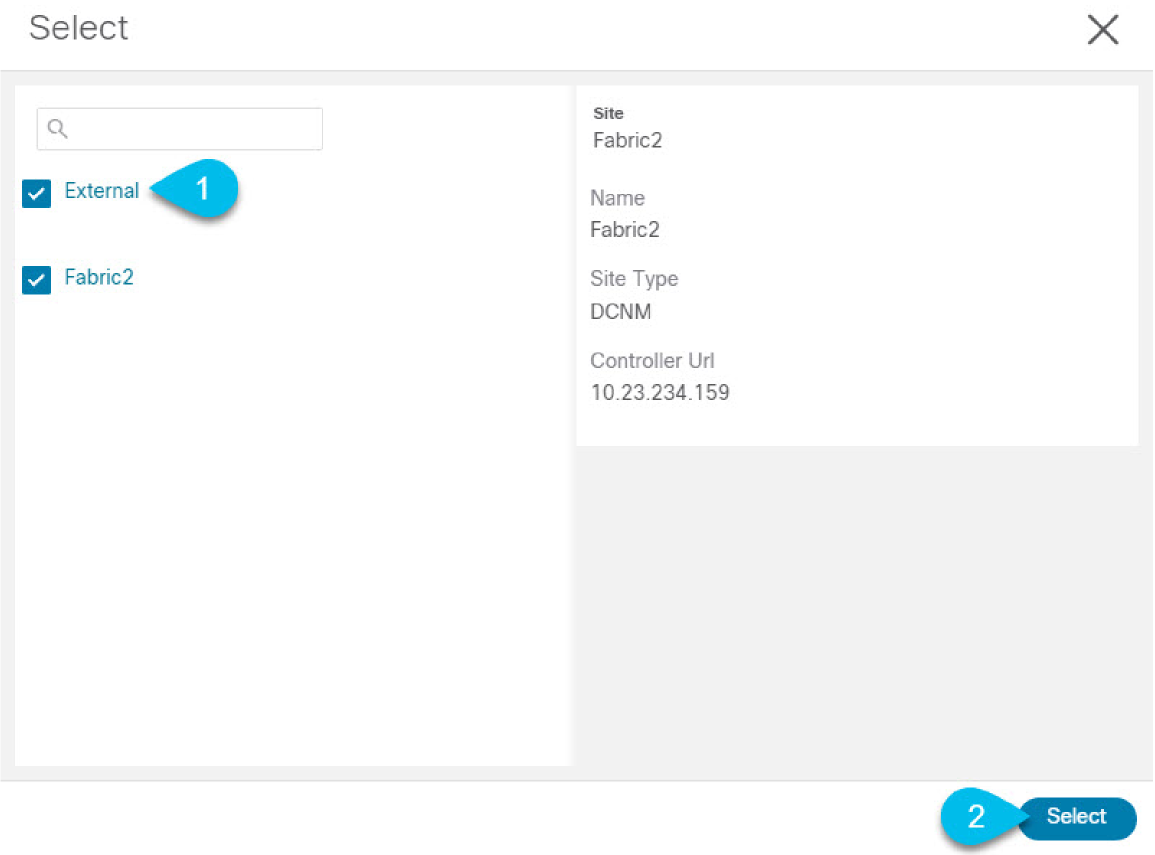

d) Click **Add Security Domains** to select one or more security domains that will have access to this site.

Step 4 Repeat the previous steps for any additional NDFC sites.

Step 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.

Step 6 In the Nexus Dashboard Orchestrator GUI, manage the sites.

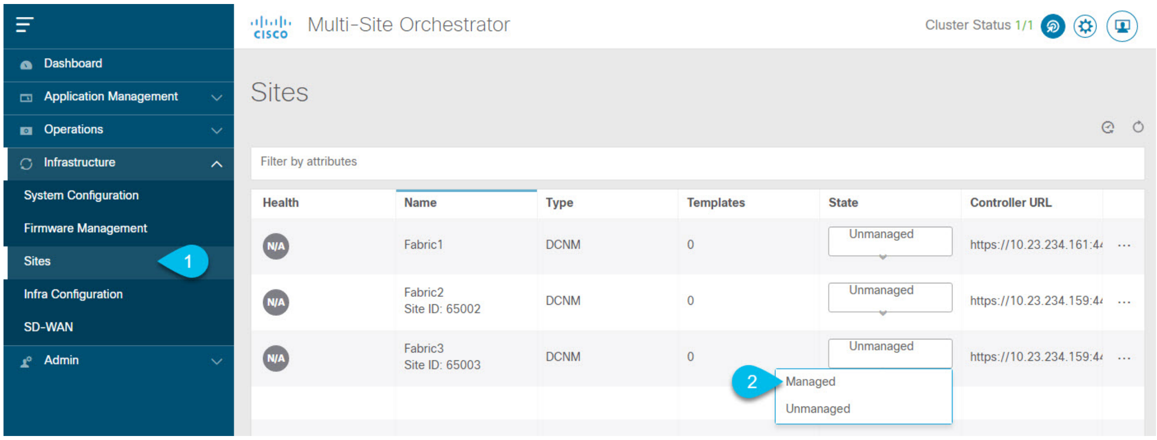

Health	Name	Type	Templates	State	Controller URL
N/A	Fabric1	DCNM	0	Unmanaged	https://10.23.234.161:4
N/A	Fabric2 Site ID: 65002	DCNM	0	Unmanaged	https://10.23.234.159:4
N/A	Fabric3 Site ID: 65003	DCNM	0	Unmanaged	https://10.23.234.159:4

a) From the left navigation menu, select **Infrastructure** > **Sites**.

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-Site Domain (MSD), it will have a **Site 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 site when you change its state to `Managed`.

Note 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 Sites

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

Before you begin

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

Step 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.

Step 2 Remove the site from all templates.

You must remove the site from all templates with which it is associated before you can unmanaged the site 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 site.
c) From the **View** dropdown, choose a template that's associated with the site that you want to remove.
d) From the **Actions** dropdown, choose **Sites Association** and uncheck the site you want to remove.

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

Note For non-stretched templates, you can choose to preserve the configurations deployed by the template to the sites by selecting **Actions** > **Dissociate Sites** 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 site that you want to unmanage in this and all other schemas.

Step 3 In the Nexus Dashboard Orchestrator GUI, disable the sites.

a) From the left navigation menu, select **Sites**.
b) In the main pane, change the **State** from `Managed` to `Unmanaged` for the site that you want to unmanage.

Note If the site 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.

Step 4 Delete the site from Nexus Dashboard.

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

Note Note that the site 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 **Sites**.
c) Select one or more sites you want to delete.
d) In the top right of the main pane, select **Actions** > **Delete Site**.
e) Provide the site's login information and click **OK**.

The site 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 site controller's GUI from the NDO's **Operate** > **Sites** screen by selecting the actions (. . .) menu next to the site 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.

CHAPTER 10

Configuring Infra for Cisco NDFC Sites

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 62 as part of the general Infra configuration procedures.

Configuring Infra: General Settings

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

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

Step 2 In the left navigation menu, choose **Configure** > **Site to Site Connectivity**.

Step 3 In the main pane, choose the **Configure** button.

Step 4 Choose the left tab, which is for **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 establishes 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 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 sites 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 **Site** drop-down, select the site that you want to connect to the route server.
- The **ASN** field will be autopopulated with the site'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 site 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.

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

If your intersite 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 intersite underlay connectivity between on-premises and cloud sites 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 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 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 **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 that you want to add.

Step 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 site CSRs and other sites (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 sites to be used for intersite 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 the 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 that are used for on-premises connectivity, which you previously configured in the Cloud Network Controller for intersite 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 **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 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.

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 autopopulate the **Multi-Site TEP** field for each fabric, which is described in Configuring Infra: NDFC Site-Specific Settings, on page 62.

For sites that were previously part of a Multi-Site Domain **(MSD)** in NDFC, this field will be prepopulated 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 Cisco 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 **Configure** > **Site to 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 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.

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

This discovers any new or removed border gateways and all site-related fabric connectivity will be reimported 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 Cisco Nexus Dashboard and open the Cisco Nexus Dashboard Orchestrator service.

Step 2 In the left navigation menu, choose **Configure** > **Site to Site Connectivity**.

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

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

Step 5 In the right *<Site>* **Settings** sidebar, specify the **Multi-Site VIP**.

This address is used for the intersite L2 BUM and L3 multicast traffic. The IP address is deployed in 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 (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 sites that are part of the MSD.

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

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

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 sites to the Internet. For more information on see, Configuring Shared Border in Cisco Nexus Dashboard Fabric Controller.

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 you are configuring a site which is already part of the NDFC Multi-Site 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, sites and fabrics.

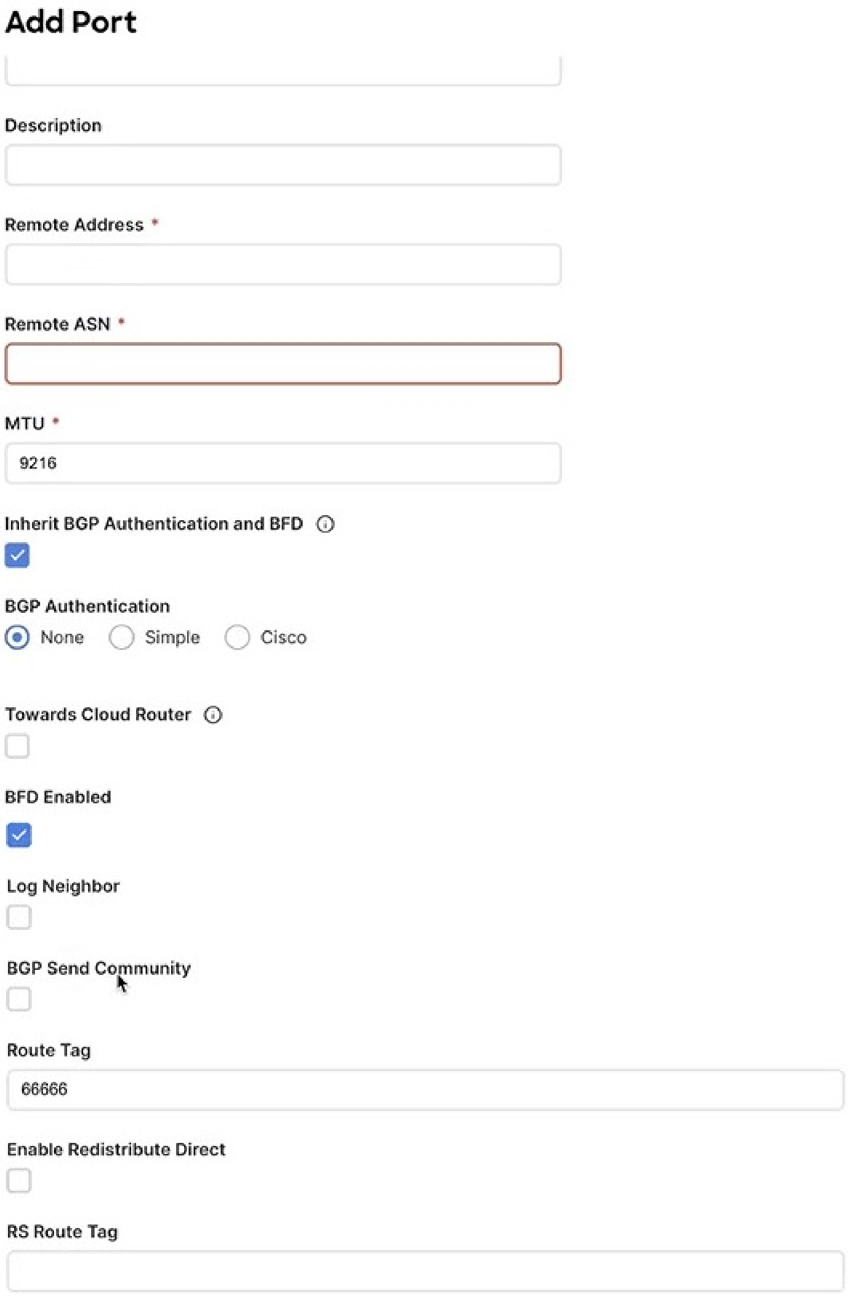

Add Port

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

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** drop-down, choose 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 **Autonomous System Number**.

- 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`.

- 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.

- Check the **BFD Enabled**, **Log Neighbor**, and **BGP Send Community** radio buttons to inherit all these features to all the multisite underlay interfaces.

- **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

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 the 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.

If configuration conflicts:

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

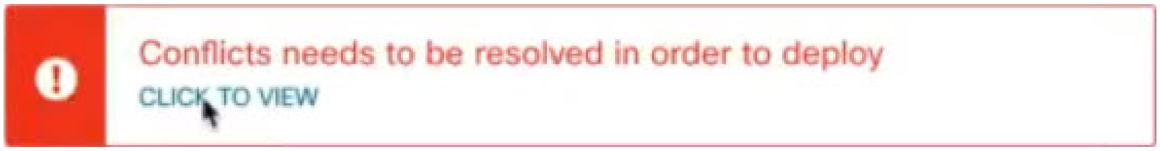

Conflicts needs to be resolved in order to deploy

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` sites.

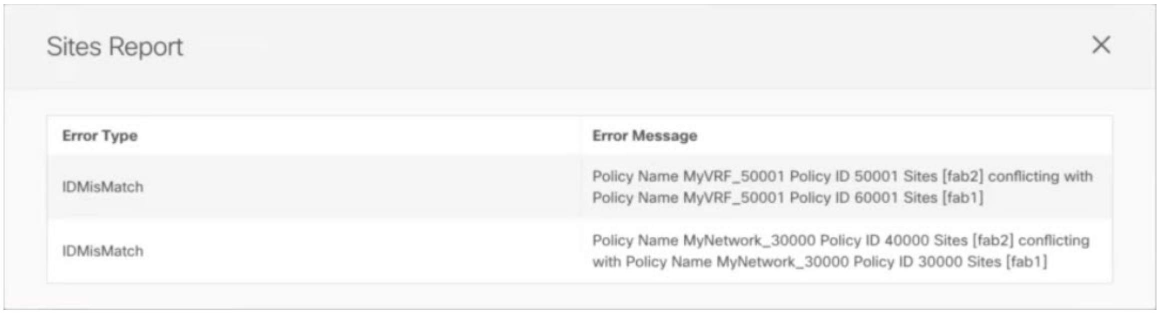

Sites Report

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]

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, on page 57.

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

e) Resolve the existing configuration conflicts.

f) Manage the site again, as described in Adding Cisco NDFC Sites, on page 55.

Since the site 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.

Step 2 Deploy configuration.

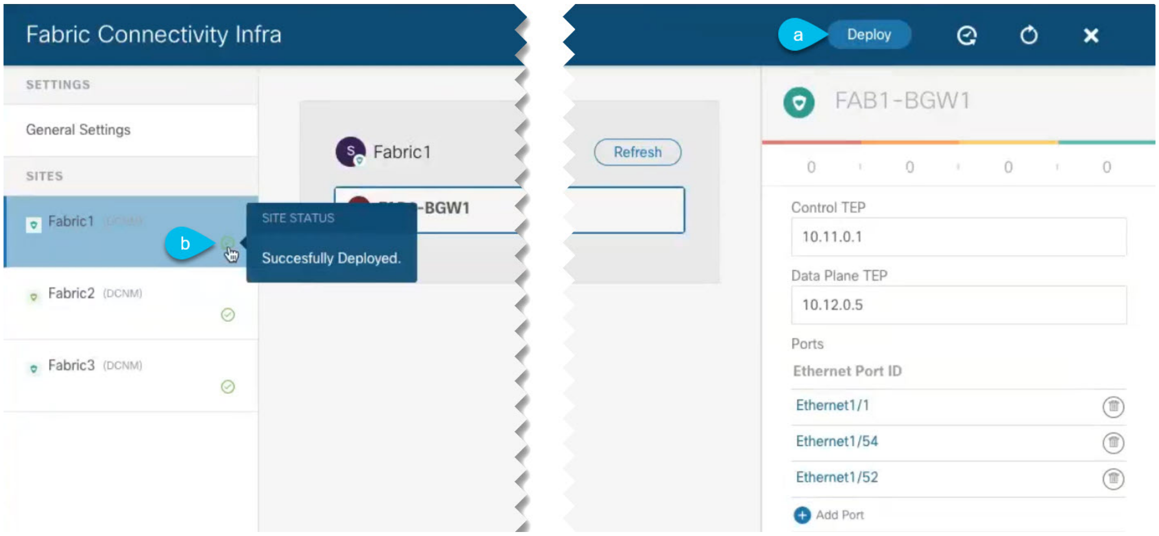

SITE STATUS

Succesfully Deployed.

Control TEP: 10.11.0.1

Data Plane TEP: 10.12.0.5

Ports

Ethernet Port ID
Ethernet1/1
Ethernet1/54
Ethernet1/52

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 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 site in the **Fabric Connectivity Infra** screen:

PART III

Upgrading Nexus Dashboard Orchestrator

CHAPTER 11

Upgrading Automatically Via Service Catalog

Overview

There are two approaches when it comes to upgrading your Nexus Dashboard Orchestrator:

- Upgrading in-place by upgrading each component (such as the Nexus Dashboard platform and the Orchestrator service) in sequence.

  This approach is described in this chapter and is recommended in the following cases:

  - If you are using a physical Nexus Dashboard cluster.
  - If you are running a recent release of Nexus Dashboard (2.2.2 or later) and Nexus Dashboard Orchestrator (3.7.1 or later).

    While you can use this approach to upgrade any Orchestrator release 3.3(1) or later, it may require upgrading the underlying Nexus Dashboard platform before you can upgrade the Orchestrator service. In those cases, an upgrade via configuration restore described below may be faster and simpler.

- Deploy a brand new Nexus Dashboard cluster, installing a new NDO service instance in it and transferring existing Orchestrator configuration via the configuration restore workflow

  This approach is described in Upgrading Manually Using Configuration Restore, on page 89 and is recommended in the following cases:

  - If you are running any release of Nexus Dashboard Orchestrator or Multi-Site Orchestrator prior to release 3.3(1).

    In this case you must upgrade using configuration restore because in-place upgrade is not supported.

- If you are using a virtual Nexus Dashboard cluster and running an older release of Nexus Dashboard Orchestrator.

  Upgrading from an old Nexus Dashboard Orchestrator release requires upgrading the underlying Nexus Dashboard platform as well, in which case deploying a new cluster and restoring configuration may shorten the required maintenance window.

  This also allows you to simply disconnect the existing cluster and keep the existing VMs until the upgrade is complete in case you want to revert to the previous version or the upgrade does not succeed.

Changes in Release 4.0(1) and Later

Beginning with Release 4.0(1), Nexus Dashboard Orchestrator will validate and enforce a number of best practices when it comes to template design and deployment:

- All policy objects must be **deployed** in order according to their dependencies.

  For example, when creating a bridge domain (BD), you must associate it with a VRF. In this case, the BD has a VRF dependency so the VRF must be deployed to the fabric before or together with the BD. If these two objects are defined in the same template, then the Orchestrator will ensure that during deployment, the VRF is created first and associate it with the bridge domain.

  However, if you define these two objects in separate templates and attempt to deploy the template with the BD first, the Orchestrator will return a validation error as the associated VRF is not yet deployed. In this case you must deploy the VRF template first, followed by the BD template.

- All policy objects must be **undeployed** in order according to their dependencies, or in other words in the opposite order in which they were deployed.

  As a corollary to the point above, when you undeploy templates, you must not undeploy objects on which other objects depend. For example, you cannot undeploy a VRF before undeploying the BD with which the VRF is associated.

- No cyclical dependencies are allowed across multiple templates.

  Consider a case of a VRF (`vrf1`) associated with a bridge domain (`bd1`), which is in turn associated with an EPG (`epg1`). If you create `vrf1` in `template1` and deploy that template, then create `bd1` in `template2` and deploy that template, there will be no validation errors since the objects are deployed in correct order. However, if you then attempt to create `epg1` in `template1`, it would create a circular dependency between the two template, so the Orchestrator will not allow you to save `template1` addition of the EPG.

Due to these additional rules and requirements, an upgrade to release 4.0(1) or later from an earlier release requires an analysis of all existing templates and conversion of any template that does not satisfy the new requirements. This is done automatically during the upgrade process described in the following sections and you will receive a detailed report of all the changes that had to be applied to your existing templates to make them compliant with the new best practices.

Note You must ensure that you complete all the requirements described in the following "Prerequisites and Guidelines" section before you back up your existing configuration for the upgrade. Failure to do so may result in template conversion to fail for one or more templates and require you to manually resolve the issues or restart the migration process.

Upgrade Workflow

The following list provides a high level overview of the upgrade process and the order of tasks you will need to perform.

1. Review the upgrade guidelines and complete all prerequisites.

2. If upgrading from a release prior to release 4.0(1), validate existing configuration using a Cisco-provided validation script.

3. If necessary, disable the existing Nexus Dashboard Orchestrator service and upgrade the Nexus Dashboard cluster.

   This is mandatory when upgrading to release 4.2(1) as you need to also upgrade the Nexus Dashboard platform software, which requires all services to be disabled during the upgrade.

   However, if your Nexus Dashboard cluster is virtual, you can choose to deploy a brand new cluster and install Nexus Dashboard release 3.0(1) or later along with the Orchestrator service release 4.2(1) in it. After the new cluster is up and running, you can disconnect the old cluster's VMs and complete the migration process on the new cluster, which allows you to preserve your existing cluster and easily bring it back in service in case of any issue with the migration procedure. This effectively turns the upgrade into a manual upgrade using the backup restore approach, and in this case we recommend following the instructions described in Upgrading Manually Using Configuration Restore, on page 89 instead

4. Re-enable your existing Orchestrator service, then upload and activate Nexus Dashboard Orchestrator release 4.2(1).

   Note Enabling the older version of NDO in the newer Nexus Dashboard is required for the upgrade purposes only and is not supported for production functionality. You must upgrade to NDO release 4.2(1) or later as described in this document.

5. Finalize the upgrade and resolve any configuration drifts.

Prerequisites and Guidelines

Before you upgrade your Cisco Nexus Dashboard Orchestrator cluster:

- Ensure that you are running Nexus Dashboard Orchestrator release 3.3(1) or later.

  Note If you are running a release prior to 3.3(1), you must skip this chapter and follow the instructions described in Upgrading Manually Using Configuration Restore, on page 89 instead.

- Note that downgrading from this release is not supported.

  If you ever want to downgrade, you can deploy a brand-new cluster using the earlier version and then restore configuration from the earlier release. Note that you cannot restore a backup created on a newer version in an older version, in other words restoring a backup from release 4.2(1) in release 3.7(1) is not supported.

- Ensure that your current Nexus Dashboard cluster is healthy.

  You can check the Nexus Dashboard cluster health in one of two ways:

  - By logging into your Nexus Dashboard GUI and verifying system status in the **System Overview** page.
  - By logging into any one of the nodes directly as `rescue-user` and running the following command:

    ```
    # acs health
    All components are healthy
    ```

- Ensure that your current Cisco Nexus Dashboard Orchestrator is healthy.

  Depending on your existing Orchestrator version, you can check the service status in one of two ways:

  - For recent NDO releases, check the Nexus Dashboard's **Overview** page that shows the platform and services' health.
  - For older Orchestrator releases, you can check the status of you Nexus Dashboard Orchestrator service by navigating to **Settings** > **System Status**:

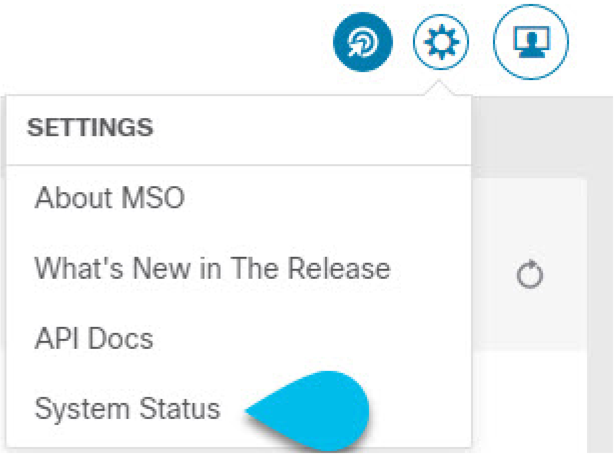

Then ensure that the status of all nodes and services is healthy:

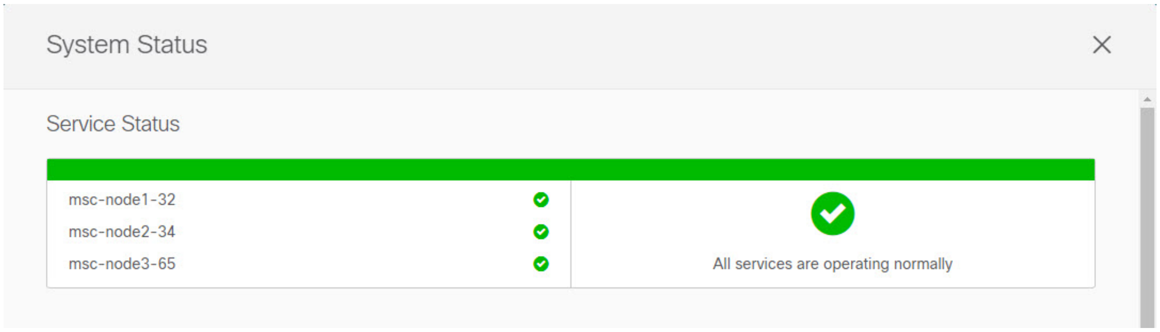

- Ensure that there are no configuration drifts before you back up your existing configuration.

  This applies to all template types available in your existing release, such as application, tenant policies, fabric policies, and fabric resource policies templates.

If your existing Nexus Dashboard Orchestrator is release 3.7(1) or later, you can use the drift reconciliation workflow for application templates, as described in the "Configuration Drifts" section of the *Nexus Dashboard Orchestrator Configuration Guide*.

- Back up and download your existing Orchestrator configurations.

  Configuration backups are described in the "Backup and Restore" chapter of the *Nexus Dashboard Orchestrator Configuration Guide* for your release.

- Back up and download your existing fabrics' configurations.

  We recommend running configuration drift reconciliation after you upgrade your Nexus Dashboard Orchestrator, which may require you to redeploy configurations to your fabrics. As such, we recommend creating configuration backups of all fabrics managed by your Nexus Dashboard Orchestrator.

  For more information on creating Cisco APIC configuration backups, see the "Management" chapter of the *Cisco APIC Basic Configuration Guide* for your release.

  For more information on creating Cisco Cloud Network Controller configuration backups, see the "Configuring Cisco Cloud Network Controller Components" chapter of the *Cisco Cloud Network Controller User Guide* for your release.

  For more information on creating Cisco Nexus Dashboard Fabric Controller configuration backups, see the "Backup and Restore" chapter of the *Cisco NDFC Fabric Controller Configuration Guide* for your release.

- Note that if you have template versioning enabled (supported since release 3.4(1)), only the latest versions of the templates are preserved during the upgrade.

  All other existing versions of templates, including older versions that are tagged `Golden`, will not be transferred.

- Ensure that all templates are in a supported state before creating the configuration backup of the existing cluster:

  - Templates that are **undeployed** or were **never deployed** after they were created require no additional attention and will be migrated during the upgrade.

  - All **deployed** templates must have no pending configuration changes.

    If you have one or more templates that have been modified since they were last deployed, you must either deploy the latest version of the template or undo the changes to the template since it was deployed by reverting to the last-deployed version and re-deploying it.

- When upgrading the Orchestrator service, you can do so in one of two ways:

  - Using the Nexus Dashboard's App Store, as described in Upgrading Orchestrator Service Using Cisco App Store, on page 82.

    In this case, the Cisco DC App Center must be reachable from the Nexus Dashboard via the Management Network directly or using a proxy configuration. Nexus Dashboard proxy configuration is described in the *Nexus Dashboard User Guide*.

Note The App Store allows you to upgrade to the latest available version of the service only. If you want to upgrade to a different release, you must use the manual upgrade process as described below.

  - By manually uploading the new app image, as described in Upgrading Orchestrator Service Manually, on page 83.

    You can use this approach if you are unable to establish the connection to the DC App Center or if you want to upgrade to a version of the service that is not the latest available release.

- SR-MPLS and SDA integration configurations are not transferred during the upgrade.

  If you have either of these integrations in your deployment, it will not affect the migration, but you will receive a notification and will need to reconfigure them after you complete the upgrade.

- If you plan to add and manage new Cloud Network Controller sites after you upgrade your Nexus Dashboard Orchestrator to this release, ensure that they are running Cloud Network Controller release 5.2(1) or later.

  On-boarding and managing Cloud Network Controller sites running earlier releases is not supported.

Validate Existing Configuration

Note If you are upgrading from release 4.0(1) or later, you can skip this section and proceed to Upgrade Your Nexus Dashboard Cluster, on page 78.

As mentioned in the Overview, on page 69, release 4.0(1) introduced a number of template validations and enforces a set of best practices when it comes to template design and deployment. The upgrade process automatically verifies the existing templates and updates them as necessary. However, some template issues cannot be addressed automatically by the upgrade and you must resolve them before upgrading from a release prior to release 4.0(1).

This section describes how to validate your existing configuration before proceeding with the upgrade.

Before you begin

You must have the following completed:

- Familiarized yourself with the migration workflow described in the Overview, on page 69
- Reviewed and completed the prerequisites described in Prerequisites and Guidelines, on page 71.

Step 1 Download and verify the configuration validation script.

You will use this script to validate your existing configuration before creating a backup and upgrading the Orchestrator service to this release.

a) Ensure that you have Python installed on your local machine.

The script requires Python 3 to run. You can check if Python is installed on your machine using the following command:

```
$ python3 --version
Python 3.9.6
```

b) Download and extract the validation script tarball.

Navigate to https://software.cisco.com/download/home/285968390/type/286317465, select the target NDO version to which you want to upgrade, download the upgrade validation script (`Final_ndo<version>-UpgradeValidationScript.tgz`), then extract it, for example:

```
$ tar -xzf Final_ndo<version>-UpgradeValidationScript.tgz
$ ls
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
Final_ndo4.1.2h-UpgradeValidationScript.tgz
UpgradeValidationScript.tgz
UpgradeValidationScript.tgz.signature
cisco_x509_verify_release.py3
```

c) Verify the validation script tarball signature.

You can use the following command to verify the Cisco signature on the configuration validation script.

```
$ ./cisco_x509_verify_release.py3 -e ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM \
  -i UpgradeValidationScript.tgz -s UpgradeValidationScript.tgz.signature -v dgst -sha512
Retrieving CA certificate from https://www.cisco.com/security/pki/certs/crcam2.cer ...
Successfully retrieved and verified crcam2.cer.
Retrieving SubCA certificate from https://www.cisco.com/security/pki/certs/innerspace.cer ...
Successfully retrieved and verified innerspace.cer.
Successfully verified root, subca and end-entity certificate chain.
Successfully fetched a public key from ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM.
Successfully verified the signature of UpgradeValidationScript.tgz using
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
```

Note If signature verification fails, you will receive the following error:

```
$ ./cisco_x509_verify_release.py3 -e ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM\
-i UpgradeValidationScript.tgz -s UpgradeValidationScript.tgz.signature.fail -v dgst
-sha512
Retrieving CA certificate from https://www.cisco.com/security/pki/certs/crcam2.cer ...
Successfully retrieved and verified crcam2.cer.
Retrieving SubCA certificate from https://www.cisco.com/security/pki/certs/innerspace.cer
 ...
Successfully retrieved and verified innerspace.cer.
Successfully verified root, subca and end-entity certificate chain.
Successfully fetched a public key from ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM.
Error log: Failed to verify dgst signature of UpgradeValidationScript.tgz.
Error log: Verification Failure
```

In this case, we recommend you re-download the `<ndo-version>-UpgradeValidationScript.tgz` tarball from the Cisco Software Download portal.

d) Once the validation script signature is verified, extract the script itself.

```
% tar -xzf UpgradeValidationScript.tgz
$ ls
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
Final_ndo4.1.2h-UpgradeValidationScript.tgz
README.md
UpgradeValidationScript.tgz
UpgradeValidationScript.tgz.signature
cisco_x509_verify_release.py3
ndo
ndoCmd.py
ndoCopy.py
requirements.txt
```

Step 2 Validate your existing configuration before you create the backup.

You can verify that your configuration backup will be compatible with upgrade to this release by running the validation script you have downloaded in the previous step. If for any reason you are unable to run the script, we recommend contacting Cisco support to have them validate your configuration backup before proceeding with the upgrade.

a) Log in to your Nexus Dashboard and open the Nexus Dashboard Orchestrator service.
b) Download the tech support logs from your existing Orchestrator.

   While for the migration you will create and download the configuration backup using the standard procedure, the validation is done on the tech support information. Note that it is normal for the tech support archive to be significantly larger than your typical configuration backup.

   You can generate the tech support logs by navigating to **Admin** > **Tech Support** page in the Orchestrator UI. Then click the **Download** icon in the **System Logs** tile. This downloads the `msc_report_<date>.zip` archive to your machine.

c) Extract the tech support archive you downloaded.

   The tech support archive comes in a standard `.zip` format, so you can use any tool of your choice to extract the contents, for example:

```
$ unzip msc_report_<date>.zip
```

   After you extract the archive, copy the `msc-db-json-<date>_temp.tar.gz` file inside into the directory where you extracted the validation script.

d) Run the validation script.

   The script requires a number of dependencies, which are all defined in the `requirements.txt` file that comes with the script, so we recommend creating a Python virtual environment before installing the dependencies and running the script:

```
$ python -m venv ndo-upgrade
$ source ndo-upgrade/bin/activate
$ pip install -r requirements.txt
```

   After the virtual environment is set up and the required modules are installed, run the script using the tech support file you downloaded and extracted in a previous step, for example:

   - `-f` allows you to provide the file on which to run the validation.
   - `-N` specifies that no configuration will be deployed to any live system.
   - `-C` generates the JSON-formatted output at the end of the script.

```
(ndo-upgrade)ndoCmd $ ./ndoCopy.py -f
msc_report_20220617_181529/msc-db-json-20220617181553_temp.tar.gz -N -C
11:49:56 Loading collection site2...4
11:49:56 Loading collection tenant...12
[...]
11:49:56 Checking template versions
11:49:56 Checking policy deployment dependencies
11:49:56 Fixing template policy flow loops
11:49:56 Fixing template dependency loops
11:49:56 Fixing policies for upgrade
11:49:56 Determine template ordering
11:49:56 Analysis completed
{
  "summaryStats": {
    "appTemplatePoliciesConverted": 139,
    "appTemplateSiteAssocMods": 7,
    "appTemplatePolicyEvictions": 2,
    "appTemplateSchemasConverted": 11,
    "appTemplatesConverted": 38,
    "appTemplatesCreated": 1,
    "tenantMods": 1
  },
[...]
}
```

After the output is generated:

- If there are no `errors` or `warnings` blocks at the end of the generated JSON, then your configuration is compliant with the migration requirements and you can proceed to the "Back up existing deployment configuration" step.
- If there is only a number of warnings but no errors, it means the migration will complete successfully, but there's a number of things that you may want to resolve before or after the upgrade. We recommend reviewing any warnings before continuing with the next step.

```
  "warnings": [
    "dropped DHCP Relay policy dhcp-tn-epgOnRL-policy: invalid provider ip address:
141.1.141.2/24",
    "dropped Route Map policy sameContract: fromPrefixLen and toPrefixLen must be larger than
 prefix",
    "dropped Mulicast Route Map policy mCastRt.map: invalid RP ip: 12.13.14.15/23",
    "dropped DHCP Option policy dhcpBdMso-option: duplicate option id: 1; duplicate option
id: 1",
    "removed dhcpLabels.0 from bd[tn-epgOnRL::Template 1::bdDhcpClient] for
unresolved policy ref key[dhcpRelayPolicies::tn-epgOnRL::dhcp-tn-epgOnRL-policy]",
    "removed dhcpLabels.0 from bd[dhcp-msite-mso::bd::bd-client-l3out] for
unresolved policy ref
key[dhcpRelayPolicies::dhcp-msite-mso::dhcp-msite-mso-relay-policy-epg-cleint-l3out]"
  ],
```

- If there is 1 or more errors listed in the JSON, the migration would fail if you continue with the current configuration.

Note You must resolve any existing errors before creating the backup and proceeding with the upgrade. We recommend re-running the validation script after you resolve any existing errors to ensure that the backup will be ready for the migration.

For example, the following sample shows 2 possible errors that can come up during validation:

```
"errors": [
  "template appTemplate[<template>] version 6 is in state EDIT_CONFIG",
  "deployed policy bd[<bd>] requires vrf[<vrf>] which is not deployed",
]
```

- As mentioned in the Prerequisites and Guidelines, on page 71 section, any deployed templates must not have undeployed changes. You must either deploy the latest version of that template or revert to the deployed version (so it is the latest version) and re-deploy the template.
- Objects must be deployed in order of their dependencies. In other words, you must not have a deployed bridge domain if the required VRF is not deployed.

e) Resolve any shown errors and repeat this step to re-validate the configuration.

Upgrade Your Nexus Dashboard Cluster

This section described how to upgrade the Nexus Dashboard cluster to release 3.0.1 which is required for this release of Nexus Dashboard Orchestrator.

Before you begin

You must have the following completed:

- Backed up and downloaded the existing Nexus Dashboard Orchestrator configuration.

Before upgrading the Nexus Dashboard cluster:

- Check your existing Nexus Dashboard release.

  At a minimum, you must upgrade to Nexus Dashboard release 3.0.1.

  Note Nexus Dashboard supports direct upgrades between specific sets of releases only:

  - If you are on release 2.2.1 or later, you can upgrade directly to release 3.0.1.
  - If you are on a release prior to release 2.2.1, you must first upgrade to release 2.2.1 and then to release 3.0.1.

- Ensure that you have read the *Release Notes* for every upgrade hop's target release for any changes in behavior, guidelines, and issues that may affect your upgrade.

  The upgrade process is the same for all Nexus Dashboard form factors. Regardless of whether you deployed your cluster using physical servers, VMware ESX, Linux KVM, Azure, or AWS, you will use the target release's ISO image to upgrade.

- Ensure that your current Nexus Dashboard cluster is healthy.

  You can check the system status on the **System Overview** page of the Nexus Dashboard GUI or by logging in to one of the nodes as `rescue-user` and ensuring that the `acs health` command returns `All components are healthy`.

- We recommend that you create a backup of the existing Nexus Dashboard cluster configuration prior to the upgrade.

- Ensure that no configuration changes are made to the cluster, such as adding worker or standby nodes, while the upgrade is in progress.

- If you are upgrading Nexus Dashboard from release 2.1(1) or earlier, you may need to clear your browser cache after the upgrade is completed for the UI to show properly.

Step 1 Download the Nexus Dashboard image.

If you have to upgrade across multiple hops as mentioned in the **Before you begin** section above, download the images for every target hop.

a) Browse to the Software Download page.

https://software.cisco.com/download/home/286327743/type/286328258

b) Choose the Nexus Dashboard version you want to download.

c) Download the Cisco Nexus Dashboard image (`nd-dk9.<version>.iso`).

Note You must download the `.iso` image for all upgrades, even if you used the VMware ESX `.ova` image or a cloud provider's marketplace for initial cluster deployment.

d) (Optional) Host the image on a web server in your environment.

When you upload the image to your Nexus Dashboard cluster, you will have an option to provide a direct URL to the image.

Step 2 Log in to your current Nexus Dashboard GUI as an Administrator user.

Step 3 Disable all installed services in the cluster.

When you upgrade Nexus Dashboard, all currently installed services must be disabled.

Step 4 Upload the new Nexus Dashboard image to the cluster.

Note Note that the UI may differ slightly across different Nexus Dashboard releases, but the navigation remains the same.

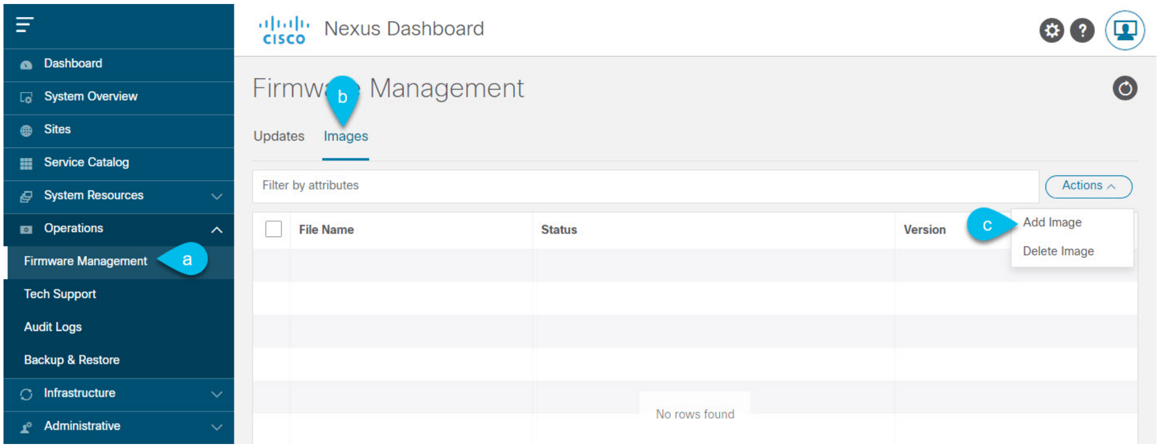

a) Navigate to **Operations** > **Firmware Management**.

b) Select the **Images** tab.

c) From the **Actions** menu, select **Add Image**.

Step 5 Select the new image.

Note If you have to upgrade over multiple hops, you must upload only one image at a time for the immediate release to which you are upgrading. Then after the upgrade to that hop is complete, you can repeat the process with the image for the next hop.

a) In the **Add Firmware Image** window, select **Local**.

Alternatively, if you hosted the image on a web server, choose **Remote** instead.

b) Click **Select file** and select the ISO image you downloaded in the first step.

If you chose to upload a remote image, provide the file path for the image on the remote server.

c) Click **Upload** to add the image.

The image will be uploaded to the Nexus Dashboard cluster, unpacked, processed, and made available for the upgrade. The whole process may take several minutes and you will be able to see the status of the process in the **Images** tab.

Step 6 Wait for the image status to change to `Downloaded`.

You can check the status of the image download progress in the **Images**.

Step 7 Set up the update.

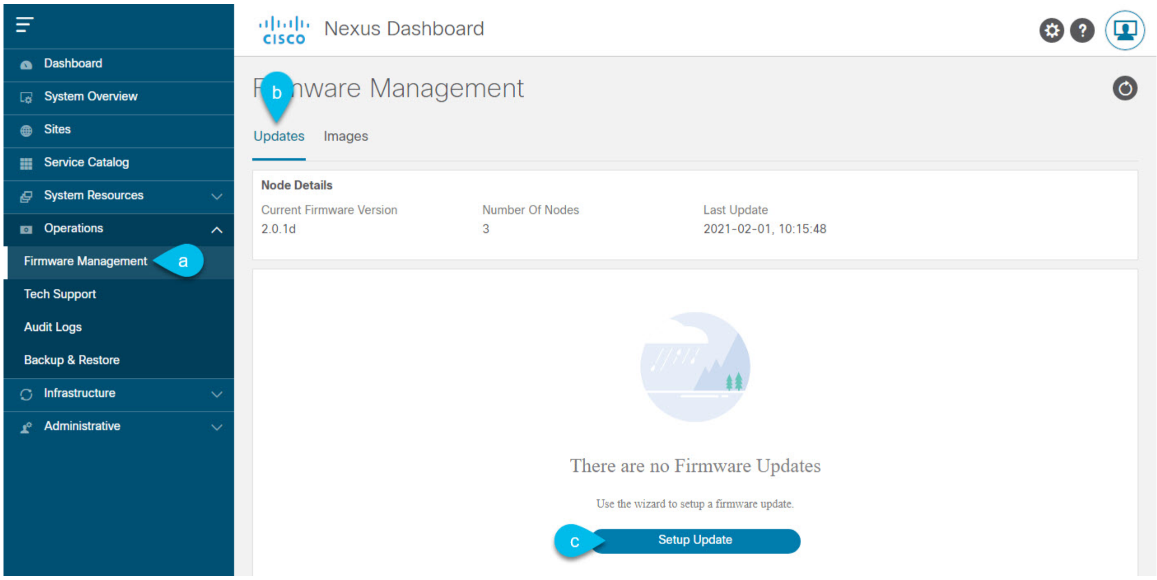

Node Details

Current Firmware Version	Number Of Nodes	Last Update
2.0.1d	3	2021-02-01, 10:15:48

There are no Firmware Updates

Use the wizard to setup a firmware update.

a) Navigate to **Operations** > **Firmware Management**.
b) Select the **Updates** tab.
c) Click **Setup Update**.

The **Firmware Update** screen opens.

Step 8 Choose the upgrade image.

a) In the **Firmware Update** > **Version selection** screen, select the firmware version you uploaded and click **Next**.
b) In the **Firmware Update** > **Confirmation** screen, verify the details and click **Begin Install**.

The installation progress window is displayed. You can navigate away from this screen while the update is in progress. To check on the update status at a later time, navigate to the **Firmware Management** screen and click **View Details** in the **Last Update Status** tile.

This will set up the required Kubernetes images and services but will not switch the cluster to the new version. The cluster will continue to run the existing version until you activate the new image in the next step. The entire process may take up to 20 minutes.

Step 9 Activate the new image.

a) Navigate back to the **Operations** > **Firmware Management** screen
b) In the **Last Update Status** tile, click **View Details**.
c) Click **Activate**.
d) In the **Activation Confirmation** window, click **Continue**.

It may take up to 20 additional minutes for all the cluster services to start and the GUI to become available. The page will automatically reload when the process is completed.

Step 10 If you upgraded a virtual cluster deployed in VMware ESX, convert the nodes to the new profile.

Note If your cluster is not deployed in VMware ESX or is already on release 2.1.1 or later, skip this step.

Starting with Release 2.1.1, Nexus Dashboard supports two different node profiles for virtual nodes deployed in VMware ESX. After the upgrade, you must convert all the nodes of the existing cluster to one of the new profiles:

- **Data node**—node profile designed for data-intensive applications, such as Nexus Dashboard Insights
- **App node**—node profile designed for non-data-intensive applications, such as Nexus Dashboard Orchestrator

The profile you choose depends on your use case scenario:

- If you plan to run only the Nexus Dashboard Orchestrator service, convert all nodes to the `App` node profile.
- If you plan to run Nexus Dashboard Insights or co-host applications, you must convert the nodes to the `Data` profile.

You convert the nodes to the new profile by deploying brand new nodes using that profile and replacing existing nodes with them one at a time.

a) Bring down one of the nodes.

You must replace one node at a time.

b) Deploy a new node in VMware ESX using the `App` or `Data` profile OVA.

When deploying the new node, you must use the same exact network configuration parameters as the node you are replacing.

c) Log in to the existing Nexus Dashboard GUI.

You can use the management IP address of one of the remaining healthy master nodes.

d) From the left navigation pane, select **System Resources** > **Nodes**.

The node you are replacing will be listed as `Inactive`.

e) Click the **(...)** menu next to the inactive master node you want to replace and select **Replace**.

The **Replace** window will open.

f) Provide the **Management IP Address** and **Password** for the node, then click **Verify**.

The cluster will connect to the new node's management IP address to verify connectivity.

g) Click **Replace**.

It may take up to 20 minutes for the node to be configured and join the cluster.

h) Wait for the cluster to become healthy, then repeat this step for the other two nodes.

Step 11 After the upgrade is successful, delete the upgrade image you had uploaded.

You must remove the older upgrade images before uploading a new image for the next upgrade.

Step 12 If you are upgrading across multiple hops, repeat these steps for each upgrade one at a time.

Upgrading Orchestrator Service Using Cisco App Store

This section describes how to upgrade Cisco Nexus Dashboard Orchestrator.

Before you begin

- Ensure that you have completed the prerequisites described in Prerequisites and Guidelines, on page 71.
- Ensure that Cisco DC App Center is reachable from the Nexus Dashboard via the Management Network directly or using a proxy configuration.

  Nexus Dashboard proxy configuration is described in the *Nexus Dashboard User Guide*

Step 1 Log in to your Nexus Dashboard.

Step 2 From the left navigation menu, select **Services**.

Step 3 Re-enable the existing version of Nexus Dashboard Orchestrator.

If you had to disable the NDO service in order to upgrade the Nexus Dashboard platform as described in the previous section, you must re-enable it before upgrading to this release.

Note Enabling the older version of NDO in the newer Nexus Dashboard is required for the upgrade purposes only and is not supported for production functionality. You must upgrade to NDO release 4.2(1) or later as described in the next steps before using the service after it's re-enabled.

Step 4 Upgrade the application using the App Store.

a) In the **Services** screen, select the **App Store** tab.
b) In the **Nexus Dashboard Orchestrator** tile, click **Upgrade**.
c) In the License Agreement window that opens, click **Agree and Download**.

Step 5 Wait for the new image to initialize.

It may take up to 20 minutes for the new application image to become available.

Step 6 Activate the new image.

a) In the **Services** page, select the **Installed Services** tab.
b) In the top right corner of the Nexus Dashboard Orchestrator tile, click the menu (. . .) and choose **Available Versions**.
c) In the available versions window, click **Activate** next to the new image.

Note Do not **Disable** the currently running image before activating the new image. The image activation process will recognize the currently running image and perform the upgrade workflows necessary for the currently running version.

It may take up to 20 additional minutes for all the application services to start and the GUI to become available. The page will automatically reload when the process is completed.

Step 7 Delete the old service image.

Because service downgrades are not supported, we recommend that you delete the older images from your Nexus Dashboard after you upgraded.

a) In the **Services** page, select the **Installed Services** tab.

b) In the top right corner of the Nexus Dashboard Orchestrator tile, click the menu (. . .) and choose **Available Versions**.

c) In the available versions window, click the delete icon next to the image you want to delete.

What to do next

After you have upgraded the NDO service, you must finalize the upgrade from the NDO UI as described in Finalize Database Upgrade, on page 84.

Upgrading Orchestrator Service Manually

This section describes how to upgrade Cisco Nexus Dashboard Orchestrator.

Before you begin

- Ensure that you have completed the prerequisites described in Prerequisites and Guidelines, on page 71.

Step 1 Download the target NDO release image.

a) Browse to the Nexus Dashboard Orchestrator page on DC App Center:

https://dcappcenter.cisco.com/nexus-dashboard-orchestrator.html

b) From the **Version** dropdown, choose the version you want to install and click **Download**.

c) Click **Agree and download** to accept the license agreement and download the image.

Step 2 Log in to your Nexus Dashboard.

Step 3 Re-enable the existing version of Nexus Dashboard Orchestrator.

If you had to disable the NDO service in order to upgrade the Nexus Dashboard platform as described in the previous section, you must re-enable it before upgrading to this release.

Note Enabling the older version of NDO in the newer Nexus Dashboard is required for the upgrade purposes only and is not supported for production functionality. You must upgrade to NDO release 4.2(1) or later as described in the next steps before using the service after it's re-enabled.

Step 4 Upload the new NDO image to your Nexus Dashboard.

a) From the left navigation menu, select **Services**.

b) In the Nexus Dashboard's **Services** screen, select the **Installed Services** tab.

c) From the **Actions** menu in the top right of main pane, select **Upload App**.

d) In the **Upload App** window, choose the location of the image

If you downloaded the application image to your system, choose **Local**.

If you are hosting the image on a server, choose **Remote**.

e) Choose the file.

If you chose **Local** in the previous substep, click **Select File** and select the app image you downloaded.

If you chose **Remote**, provide the full URL to the image file, for example `http://<ip-address>:<port>/<full-path>/cisco-mso-<version>.nap`.

f) Click **Upload** to add the app to the cluster.

A new tile will appear with the upload progress bar. Once the image upload is completed, the Nexus Dashboard will recognize the new image as an existing application and add it as a new version.

Step 5 Wait for the new image to initialize.

It may take up to 20 minutes for the new application image to become available.

Step 6 Activate the new image.

a) In the **OperateServices** page, select the **Installed Services** tab.
b) In the top right corner of the Nexus Dashboard Orchestrator tile, click the menu (. . .) and choose **Available Versions**.
c) In the available versions window, click **Activate** next to the new image.

Note Do not **Disable** the currently running image before activating the new image. The image activation process will recognize the currently running image and perform the upgrade workflows necessary for the currently running version.

It may take up to 20 additional minutes for all the application services to start and the GUI to become available. The page will automatically reload when the process is completed.

Step 7 Delete the old service image.

Because service downgrades are not supported, we recommend that you delete the older images from your Nexus Dashboard after you upgraded.

a) In the **Services** page, select the **Installed Services** tab.
b) In the top right corner of the Nexus Dashboard Orchestrator tile, click the menu (. . .) and choose **Available Versions**.
c) In the available versions window, click the delete icon next to the image you want to delete.

What to do next

After you have upgraded the NDO service, you must finalize the upgrade from the NDO UI as described in Finalize Database Upgrade, on page 84.

Finalize Database Upgrade

This section describes how to deploy and configure the new Nexus Dashboard cluster and the NDO service, which you will use to restore your previous configuration.

Before you begin

You must have the following completed:

- Backed up and downloaded the existing Nexus Dashboard Orchestrator configuration.
- Installed the target Orchestrator release as described in Deploy Nexus Dashboard and Install Nexus Dashboard Orchestrator, on page 97.

Step 1 Open the Orchestrator service UI.

Simply click **Open** on the service tile in the Nexus Dashboard's **Services** page.

When you open NDO 4.1(2) or later for the first time, it automatically starts **Post Upgrade Validation** process.

Step 2 Finalize database upgrade.

a) Ensure that there are no failures listed in the report, then click **Restore and Continue** to proceed.

Before the configuration database is updated for this release, the upgrade process performs a number of validations. The validation provides a summary of template and policy changes that are performed during this final upgrade stage in the next step and includes the following:

- Implicit template stretching – if one or more objects are implicitly stretched, the upgrade process will create new explicitly-stretched templates and move the objects into those templates.

  For example, if you have a template (`t1`) that contains `vrf1` and is associated to `site1` and another template (`t2`) that contains a BD that references `vrf1` but is associated to two sites (`site1` and `site2`), then `vrf1` will be implicitly stretched between the two sites.

  This is no longer allowed starting with release 4.0(1) and the VRF must be explicitly stretched to both sites. In such cases during the upgrade, the VRF will be either moved to a different template which will be explicitly stretched between both sites or the original template will be associated with both sites, depending on whether the other policies in that template require stretching as well.

  Any templates that are created in this case will be named `UpgradeTemplate%d`, where `%d` is an incrementing number starting with 1 to ensure that all newly added templates are unique.

- Global policy migration – all global tenant policies (such as DHCP relay or route maps) and fabric policies (such as QoS) will be moved into the new tenant and fabric policy templates that have been added in release 4.0(1).

This is the stage of the upgrade where the existing schemas and templates are recreated in your NDO configuration database according to the current best practices. These schemas and templates are then posted to the local NDO database as if it were a greenfield schema/template creation. Then the newly saved templates are deployed in the correct order that conforms to the current deployment requirements and best practices. The template deployment in this step uses a “local deploy” option to calculate the deployment plan and update the database, but does not send any configuration payload to the sites' controllers.

The upgrade process also checks for any configuration drifts between the local NDO database (configuration that is correct from NDO's point of view) and what is actually deployed in the fabrics. If this release of NDO supports additional objects or properties compared to the release from which you are upgrading, the upgrade will automatically reconcile those drifts by importing the existing configuration from the site's controller.

b) Review the report from the previous substep and click **Ok** to finish.

The final stage of the database upgrade presents a full report of the performed actions for you to review. If you close the report but want to review it again, simply click the **View Restore Report** in the **Backups** page.

Step 3 Verify that backup was restored successfully and all objects and configurations are present.

a) In the **Sites** page, verify that all sites are listed as `Managed`.

b) In the **Tenants** and **Schemas** pages, confirm that all tenants and schemas from your previous Nexus Dashboard Orchestrator cluster are present.

c) Navigate to **Infrastructure** > **Site Connectivity** and confirm that intersite connectivity is intact.

In the main pane, click **Show Connectivity Status** next to each site and verify that the existing tunnels are up and connectivity was not interrupted.

d) In the main pane, click **Configure** to open **Fabric Connectivity Infra** screen and verify **External Subnet Pool** addresses.

   You can view the external subnet pools by selecting **General Settings** > **IPsec Tunnel Subnet Pools** tab of the **Fabric Connectivity Infra** screen and verify that the External Subnet Pools previously configured in Cloud Network Controller have been imported from the cloud sites.

   These subnets are used to address the IPsec tunnel interfaces and loopbacks of the Cloud Routers used for on-premises connectivity and had to be configured directly in the Cloud Network Controller in earlier Nexus Dashboard Orchestrator releases.

Resolve Configuration Drifts

In some cases you may run into a situation where the configuration actually deployed in the site's controller is different from the configuration defined in the Nexus Dashboard Orchestrator. These configuration discrepancies are referred to as **Configuration Drifts** and are indicated by an `Out of Sync` warning next to the site name in the template view page

After upgrading your Nexus Dashboard Orchestrator and restoring the previous configuration backup, we recommend that you check for and resolve any configuration drifts that were not automatically resolved by the upgrade process as described in this section.

Note Deploying any templates before resolving configuration drifts would push the configuration defined in the Orchestrator and overwrite the values defined in the fabrics' controllers.

Step 1 In your Nexus Dashboard Orchestrator, navigate to **Operate** > **Tenant Templates**.

Step 2 Choose the **Applications** tab.

Step 3 Select the first schema and check its templates for configuration drifts.

You will repeat the following steps for every schema and template in your deployment

You can check for configuration drifts in one of the following two ways:

- Check the template deployment status icon for each site to which the template is assigned.
- Select the template and click **Deploy to sites** to bring up the configuration comparison screen to check which objects contain configuration drifts.

Step 4 If the template contains a configuration drift, resolve the conflicts.

For more information about configuration drifts, check the "Configuration Drifts" chapter in the *Cisco Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics*.

a) Close the template deployment dialog to return to the Schema view.

   Deploying any templates at this point would push the values in the Orchestrator database and overwrite any existing settings in the fabrics.

b) From the template's **Actions** menu, select **Drift Reconciliation**.

The **Drift Reconciliation** wizard opens.

c) In the **Drift Reconciliation** screen, compare the template-level configurations for each site and choose the one you want.

Template-level properties are common across all sites associated to the template. You can compare the template level properties defined on Nexus Dashboard Orchestrator with the configuration rendered in each site and decide what should become the new configuration in the Nexus Dashboard Orchestrator template. Selecting the site configuration will modify those properties in the existing Nexus Dashboard Orchestrator template, whereas selecting the Nexus Dashboard Orchestrator configuration will keep the existing Nexus Dashboard Orchestrator template settings as is

d) Click **Go to Site Specific Properties** to switch to site-level configuration.

You can choose a site to compare that specific site's configuration. Unlike template-level configurations, you can choose either the Nexus Dashboard Orchestrator-defined or actual existing configurations for each site individually to be retained as the template's site-local properties for that site.

Even though in most scenarios you will make the same choice for both template-level and site-level configuration, the drift reconciliation wizard allows you to choose the configuration defined in the site's controller at the "Template Properties" level and the configuration defined in Nexus Dashboard Orchestrator at the "Site Local Properties" level or vice versa.

e) Click **Preview Changes** to verify your choices.

The preview will display full template configuration adjusted based on the choices picked in the **Drift Reconciliation** wizard

f) Save the schema.

g) Click **Deploy to sites** to deploy the configuration and finish reconciling the drift for that template

Step 5 Repeat the above steps for every schema and template in your Nexus Dashboard Orchestrator.

Step 6 Check audit logs to verify that all templates have been re-deployed.

You can view the audit logs in the **Operations** tab.

Audit Logs page and confirm that all templates show as `Redeployed` to ensure that full re-deployment successfully completed.

CHAPTER 12

Upgrading Manually Using Configuration Restore

Overview

There are two approaches when it comes to upgrading your Nexus Dashboard Orchestrator:

- Upgrading in-place by upgrading each component (such as the Nexus Dashboard platform and the Orchestrator service) in sequence.

  This approach is described in Upgrading Automatically Via Service Catalog, on page 69 and is recommended in the following cases:

  - If you are using a physical Nexus Dashboard cluster.
  - If you are running a recent release of Nexus Dashboard (2.2.2 or later) and Nexus Dashboard Orchestrator (3.7.1 or later).

    While you can use this approach to upgrade any Orchestrator release 3.3(1) or later, it may require upgrading the underlying Nexus Dashboard platform before you can upgrade the Orchestrator service. In those cases, an upgrade via configuration restore described below may be faster and simpler.

- Deploy a brand new Nexus Dashboard cluster, installing a new NDO service instance in it and transferring existing Orchestrator configuration via the configuration restore workflow

  This approach is described in this chapter and is recommended in the following cases:

  - If you are running any release of Nexus Dashboard Orchestrator or Multi-Site Orchestrator prior to release 3.3(1).

    In this case you must upgrade using configuration restore because in-place upgrade is not supported.

  - If you are using a virtual Nexus Dashboard cluster and running an older release of Nexus Dashboard Orchestrator.

Upgrading from an old Nexus Dashboard Orchestrator release requires upgrading the underlying Nexus Dashboard platform as well, in which case deploying a new cluster and restoring configuration may shorten the required maintenance window.

This also allows you to simply disconnect the existing cluster and keep the existing VMs until the upgrade is complete in case you want to revert to the previous version or the upgrade does not succeed.

Changes in Release 4.0(1) and Later

Beginning with Release 4.0(1), Nexus Dashboard Orchestrator will validate and enforce a number of best practices when it comes to template design and deployment:

- All policy objects must be **deployed** in order according to their dependencies.

  For example, when creating a bridge domain (BD), you must associate it with a VRF. In this case, the BD has a VRF dependency so the VRF must be deployed to the fabric before or together with the BD. If these two objects are defined in the same template, then the Orchestrator will ensure that during deployment, the VRF is created first and associate it with the bridge domain.

  However, if you define these two objects in separate templates and attempt to deploy the template with the BD first, the Orchestrator will return a validation error as the associated VRF is not yet deployed. In this case you must deploy the VRF template first, followed by the BD template.

- All policy objects must be **undeployed** in order according to their dependencies, or in other words in the opposite order in which they were deployed.

  As a corollary to the point above, when you undeploy templates, you must not undeploy objects on which other objects depend. For example, you cannot undeploy a VRF before undeploying the BD with which the VRF is associated.

- No cyclical dependencies are allowed across multiple templates.

  Consider a case of a VRF (`vrf1`) associated with a bridge domain (`bd1`), which is in turn associated with an EPG (`epg1`). If you create `vrf1` in `template1` and deploy that template, then create `bd1` in `template2` and deploy that template, there will be no validation errors since the objects are deployed in correct order. However, if you then attempt to create `epg1` in `template1`, it would create a circular dependency between the two template, so the Orchestrator will not allow you to save `template1` addition of the EPG.

Due to these additional rules and requirements, an upgrade to release 4.0(1) or later from an earlier release requires an analysis of all existing templates and conversion of any template that does not satisfy the new requirements. This is done automatically during the upgrade process described in the following sections and you will receive a detailed report of all the changes that had to be applied to your existing templates to make them compliant with the new best practices.

Note You must ensure that you complete all the requirements described in the following "Prerequisites and Guidelines" section before you back up your existing configuration for the upgrade. Failure to do so may result in template conversion to fail for one or more templates and require you to manually resolve the issues or restart the migration process.

Upgrade Workflow

The following list provides a high level overview of the migration process and the order of tasks you will need to perform.

1. Review the upgrade guidelines and complete all prerequisites.
2. Validate existing configuration using a Cisco-provided validation script, then create a backup of the existing Nexus Dashboard Orchestrator configuration and download the backup to your local machine.
3. Disconnect or bring down your existing cluster.

   If your existing cluster is virtual, you can simply disconnect it from the network until you've deployed a new cluster and restored the configuration backup in it. This allows you to preserve your existing cluster and easily bring it back in service in case of any issues with the migration procedure.
4. Deploy a brand new Nexus Dashboard cluster release 2.3(b) or later and install Nexus Dashboard Orchestrator release 4.1(2) or later.
5. Add a remote location for backups to the fresh Nexus Dashboard Orchestrator instance, upload the backup you took on your previous release, and restore the configuration backup in the new NDO installation.
6. Resolve any configuration drifts.

Prerequisites and Guidelines

Before you upgrade your Cisco Nexus Dashboard Orchestrator cluster:

- Note that downgrading from this release is not supported.

  If you ever want to downgrade, you can deploy a brand-new cluster using the earlier version and then restore configuration from the earlier release. Note that you cannot restore a backup created on a newer version in an older version, in other words restoring a backup from release 4.2(1) in release 3.7(1) is not supported.
- The backup/restore upgrade workflow supports upgrades from any Multi-Site Orchestrator (MSO) release 2.x and 3.x as well as Nexus Dashboard Orchestrator (NDO) release 3.x and 4.x to this release of NDO.
- Ensure that there are no configuration drifts before you back up your existing configuration.

  This applies to all template types available in your existing release, such as application, tenant policies, fabric policies, and fabric resource policies templates.

  If your existing Nexus Dashboard Orchestrator is release 3.7(1) or later, you can use the drift reconciliation workflow for application templates, as described in the "Configuration Drifts" section of the *Nexus Dashboard Orchestrator Configuration Guide*.
- Back up and download your existing Orchestrator configurations.

  Configuration backups are described in the "Backup and Restore" chapter of the *Nexus Dashboard Orchestrator Configuration Guide* for your release.
- Back up and download your existing fabrics' configurations.

  We recommend running configuration drift reconciliation after you upgrade your Nexus Dashboard Orchestrator, which may require you to redeploy configurations to your fabrics. As such, we recommend creating configuration backups of all fabrics managed by your Nexus Dashboard Orchestrator.

For more information on creating Cisco APIC configuration backups, see the "Management" chapter of the *Cisco APIC Basic Configuration Guide* for your release.

For more information on creating Cisco Cloud Network Controller configuration backups, see the "Configuring Cisco Cloud Network Controller Components" chapter of the *Cisco Cloud Network Controller User Guide* for your release.

For more information on creating Cisco Nexus Dashboard Fabric Controller configuration backups, see the "Backup and Restore" chapter of the *Cisco NDFC Fabric Controller Configuration Guide* for your release.

- Note that if versioning enabled (supported since release 3.4(1)), only the latest versions of the templates are preserved during the upgrade.

  All other existing versions of templates, including older versions that are tagged `Golden`, will not be transferred.

- Ensure that all templates are in a supported state before creating the configuration backup of the existing cluster:

  - Templates that are **undeployed** or were **never deployed** after they were created require no additional attention and will be migrated during the upgrade.

  - All **deployed** templates must have no pending configuration changes.

    If you have one or more templates that have been modified since they were last deployed, you must either deploy the latest version of the template or undo the changes to the template since it was deployed by reverting to the last-deployed version and re-deploying it.

Note Attempting to restore a backup that contains invalid templates will fail and you would need to revert to your existing release, restore your backup, resolve any existing issues, and then restart the migration process. So we strongly recommend that you validate your backup locally using the provided Python script before proceeding with the upgrade, as described in the Validate Existing Configuration and Create Backup, on page 93 section below. If for any reason you are unable to run the script, we recommend contacting Cisco support to have them validate your configuration backup before proceeding with the upgrade.

- When installing the Orchestrator service, you can do so in one of two ways:

  - Using the Nexus Dashboard's App Store, in which case, the Cisco DC App Center must be reachable from the Nexus Dashboard via the Management Network directly or using a proxy configuration. Nexus Dashboard proxy configuration is described in the *Nexus Dashboard User Guide*.

Note The App Store allows you to upgrade to the latest available version of the service only. If you want to upgrade to a different release, you must use the manual upgrade process as described below.

  - By manually uploading the new app image, which you can do if you are unable to establish the connection to the DC App Center or if you want to upgrade to a version of the service that is not the latest available release.

- SR-MPLS and SDA integration configurations are not transferred during the upgrade.

  If you have either of these integrations in your deployment, it will not affect the migration, but you will receive a notification and will need to reconfigure them after you complete the upgrade.

- If you plan to add and manage new Cloud Network Controller sites after you upgrade your Nexus Dashboard Orchestrator to this release, ensure that they are running Cloud Network Controller release 5.2(1) or later.

  On-boarding and managing Cloud Network Controller sites running earlier releases is not supported.

- Ensure that you have a remote location for backups that you can add to the Nexus Dashboard Orchestrator after the upgrade.

  Backing up and restoring configuration in release 4.1(2) and later requires the backup to be stored on a remote location, which must be configured in NDO UI. Detailed information about backups and remote locations is available in the **Operations** > **Backup and Restore** chapter of the *Cisco Nexus Dashboard Orchestrator Configuration Guide*.

  Note that if you had a remote location already configured in your existing installation, it is not preserved during configuration restore; so you will need to add the same remote location after you deploy this release in order to restore the configuration.

Validate Existing Configuration and Create Backup

This section describes how to create a backup of the existing configuration, which you will then restore after you re-deploy a fresh instance of Nexus Dashboard Orchestrator.

Before you begin

You must have the following completed:

- Familiarized yourself with the migration workflow described in the Overview, on page 89
- Reviewed and completed the prerequisites described in Prerequisites and Guidelines, on page 71.

Step 1 Download and verify the configuration validation script.

Note If you are upgrading from release 4.0(1) or later, you can skip this step.

You will use this script to validate your existing configuration before creating a backup and upgrading the Orchestrator service to this release.

a) Ensure that you have Python installed on your local machine.

The script requires Python 3 to run. You can check if Python is installed on your machine using the following command:

```
$ python3 --version
Python 3.9.6
```

b) Download and extract the validation script tarball.

Navigate to https://software.cisco.com/download/home/285968390/type/286317465, select the target NDO version to which you want to upgrade, download the upgrade validation script (`Final_ndo<version>-UpgradeValidationScript.tgz`), then extract it, for example:

```
$ tar -xzf Final_ndo<version>-UpgradeValidationScript.tgz
$ ls
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
Final_ndo4.1.2h-UpgradeValidationScript.tgz
UpgradeValidationScript.tgz
UpgradeValidationScript.tgz.signature
cisco_x509_verify_release.py3
```

c) Verify the validation script tarball signature.

You can use the following command to verify the Cisco signature on the configuration validation script.

```
$ ./cisco_x509_verify_release.py3 -e ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM \
  -i UpgradeValidationScript.tgz -s UpgradeValidationScript.tgz.signature -v dgst -sha512
Retrieving CA certificate from https://www.cisco.com/security/pki/certs/crcam2.cer ...
Successfully retrieved and verified crcam2.cer.
Retrieving SubCA certificate from https://www.cisco.com/security/pki/certs/innerspace.cer ...
Successfully retrieved and verified innerspace.cer.
Successfully verified root, subca and end-entity certificate chain.
Successfully fetched a public key from ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM.
Successfully verified the signature of UpgradeValidationScript.tgz using
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
```

Note If signature verification fails, you will receive the following error:

```
$ ./cisco_x509_verify_release.py3 -e ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM\
-i UpgradeValidationScript.tgz -s UpgradeValidationScript.tgz.signature.fail -v dgst
-sha512
Retrieving CA certificate from https://www.cisco.com/security/pki/certs/crcam2.cer ...
Successfully retrieved and verified crcam2.cer.
Retrieving SubCA certificate from https://www.cisco.com/security/pki/certs/innerspace.cer
 ...
Successfully retrieved and verified innerspace.cer.
Successfully verified root, subca and end-entity certificate chain.
Successfully fetched a public key from ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM.
Error log: Failed to verify dgst signature of UpgradeValidationScript.tgz.
Error log: Verification Failure
```

In this case, we recommend you re-download the `<ndo-version>-UpgradeValidationScript.tgz` tarball from the Cisco Software Download portal.

d) Once the validation script signature is verified, extract the script itself.

```
% tar -xzf UpgradeValidationScript.tgz
$ ls
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
Final_ndo4.1.2h-UpgradeValidationScript.tgz
README.md
UpgradeValidationScript.tgz
UpgradeValidationScript.tgz.signature
cisco_x509_verify_release.py3
ndo
ndoCmd.py
ndoCopy.py
requirements.txt
```

Step 2 Validate your existing configuration before you create the backup.

Note If you are upgrading from release 4.0(1) or later, you can skip this step.

You can verify that your configuration backup will be compatible with upgrade to this release by running the validation script you have downloaded in the previous step. If for any reason you are unable to run the script, we recommend contacting Cisco support to have them validate your configuration backup before proceeding with the upgrade.

a) Log in to your Nexus Dashboard and open the Nexus Dashboard Orchestrator service.

b) Download the tech support logs from your existing Orchestrator.

While for the migration you will create and download the configuration backup using the standard procedure, the validation is done on the tech support information. Note that it is normal for the tech support archive to be significantly larger than your typical configuration backup.

You can generate the tech support logs by navigating to **Admin** > **Tech Support** page in the Orchestrator UI. Then click the **Download** icon in the **System Logs** tile. This downloads the `msc_report_<date>.zip` archive to your machine.

c) Extract the tech support archive you downloaded.

The tech support archive comes in a standard `.zip` format, so you can use any tool of your choice to extract the contents, for example:

```
$ unzip msc_report_<date>.zip
```

After you extract the archive, copy the `msc-db-json-<date>_temp.tar.gz` file inside into the directory where you extracted the validation script.

d) Run the validation script.

The script requires a number of dependencies, which are all defined in the `requirements.txt` file that comes with the script, so we recommend creating a Python virtual environment before installing the dependencies and running the script:

```
$ python -m venv ndo-upgrade
$ source ndo-upgrade/bin/activate
$ pip install -r requirements.txt
```

After the virtual environment is set up and the required modules are installed, run the script using the tech support file you downloaded and extracted in a previous step, for example:

- `-f` allows you to provide the file on which to run the validation.
- `-N` specifies that no configuration will be deployed to any live system.
- `-C` generates the JSON-formatted output at the end of the script.

```
(ndo-upgrade)ndoCmd $ ./ndoCopy.py -f
msc_report_20220617_181529/msc-db-json-20220617181553_temp.tar.gz -N -C
11:49:56 Loading collection site2...4
11:49:56 Loading collection tenant...12
[...]
11:49:56 Checking template versions
11:49:56 Checking policy deployment dependencies
11:49:56 Fixing template policy flow loops
11:49:56 Fixing template dependency loops
11:49:56 Fixing policies for upgrade
11:49:56 Determine template ordering
11:49:56 Analysis completed
{
  "summaryStats": {
    "appTemplatePoliciesConverted": 139,
    "appTemplateSiteAssocMods": 7,
    "appTemplatePolicyEvictions": 2,
    "appTemplateSchemasConverted": 11,
    "appTemplatesConverted": 38,
    "appTemplatesCreated": 1,
    "tenantMods": 1
  },
[...]
}
```

After the output is generated:

- If there are no `errors` or `warnings` blocks at the end of the generated JSON, then your configuration is compliant with the migration requirements and you can proceed to the "Back up existing deployment configuration" step.
- If there is only a number of warnings but no errors, it means the migration will complete successfully, but there's a number of things that you may want to resolve before or after the upgrade. We recommend reviewing any warnings before continuing with the next step.

```
  "warnings": [
    "dropped DHCP Relay policy dhcp-tn-epgOnRL-policy: invalid provider ip address:
141.1.141.2/24",
    "dropped Route Map policy sameContract: fromPrefixLen and toPrefixLen must be larger than
 prefix",
    "dropped Mulicast Route Map policy mCastRt.map: invalid RP ip: 12.13.14.15/23",
    "dropped DHCP Option policy dhcpBdMso-option: duplicate option id: 1; duplicate option
id: 1",
    "removed dhcpLabels.0 from bd[tn-epgOnRL::Template 1::bdDhcpClient] for
unresolved policy ref key[dhcpRelayPolicies::tn-epgOnRL::dhcp-tn-epgOnRL-policy]",
    "removed dhcpLabels.0 from bd[dhcp-msite-mso::bd::bd-client-l3out] for
unresolved policy ref
key[dhcpRelayPolicies::dhcp-msite-mso::dhcp-msite-mso-relay-policy-epg-cleint-l3out]"
  ],
```

- If there is 1 or more errors listed in the JSON, the migration would fail if you continue with the current configuration.

Note You must resolve any existing errors before creating the backup and proceeding with the upgrade. We recommend re-running the validation script after you resolve any existing errors to ensure that the backup will be ready for the migration.

For example, the following sample shows 2 possible errors that can come up during validation:

```
"errors": [
  "template appTemplate[<template>] version 6 is in state EDIT_CONFIG",
  "deployed policy bd[<bd>] requires vrf[<vrf>] which is not deployed",
]
```

  - As mentioned in the Prerequisites and Guidelines, on page 71 section, any deployed templates must not have undeployed changes. You must either deploy the latest version of that template or revert to the deployed version (so it is the latest version) and re-deploy the template.
  - Objects must be deployed in order of their dependencies. In other words, you must not have a deployed bridge domain if the required VRF is not deployed.

e) Resolve any shown errors and repeat this step to re-validate the configuration.

Step 3 Back up existing deployment configuration.

a) From the left navigation pane, select **Operations** > **Backups & Restore**.
b) In the main window, click **New Backup**.

A **New Backup** window opens.

c) In the **Name** field, provide the name for the backup file.

The name can contain up to 10 alphanumeric characters, but no spaces or underscores (_).

d) From the **Remote Location** dropdown, choose the remote location you have configured previously.

e) In the **Remote Path** field, provide the path on the remote server where to store the backup.

f) Click **Save** to create the backup.

Step 4 Download the backup file.

In the main window, click the actions (⋮) icon next to the backup and select **Download**. This will download the backup file to your system.

Deploy Nexus Dashboard and Install Nexus Dashboard Orchestrator

Because you are deploying a fresh cluster, the steps are identical to the ones described in the Deploying Nexus Dashboard Orchestrator, on page 3 chapter of this guide. This section summarizes the steps and provides specific links for each one.

Before you begin

You must have the following completed:

- Backed up and downloaded the existing Nexus Dashboard Orchestrator configuration.

Step 1 Deploy a new Nexus Dashboard cluster.

Detailed information about deployment requirements, available form factors, and installation instructions are available from the *Cisco Nexus Dashboard Deployment Guide*.

Note You must deploy in Nexus Dashboard release 3.0.1 or later.

Step 2 After Nexus Dashboard and Nexus Dashboard Orchestrator have been successfully deployed, navigate to your Nexus Dashboard's **Admin Console**.

You can log in to Nexus Dashboard by opening your browser and navigating to the management IP address of any one of the Nexus Dashboard cluster's nodes, then selecting **Admin Console** from the drop down menu in the top navigation bar.

Step 3 Onboard the fabrics managed by the Orchestrator service to Nexus Dashboard.

Note You must on-board all the fabrics previously managed by your Nexus Dashboard Orchestrator (or Multi-Site Orchestrator) to Nexus Dashboard before you can proceed with restoring configuration and completing the upgrade process.

In this release, fabric on-boarding is done in the common Nexus Dashboard screen. The process is described in detail in the "Site Management" chapter of the *Cisco Nexus Dashboard User Guide*, but in short:

a) From the main navigation menu, choose **Sites**.

b) In the main pane, click **Add Site**.

c) Choose the type of site you want to on-board and provide the site's information, such as controller's IP address, username, and password.

- For on-premises sites managed by Cisco APIC, choose `ACI`.
- For cloud sites managed by Cloud Network Controller (previously Cloud APIC), choose `Cloud Network Controller`.

  Use this option for all Cloud APIC sites that were managed by your Orchestrator.
- For on-premises sites managed by NDFC (previously DCNM), choose `NDFC`.

  Use this option for all DCNM sites that were managed by your Orchestrator.

Note You must use the same site **Name** as you did when on-boarding the site to your Orchestrator in the past. Adding a site with a different name will cause configuration restore to fail.

d) Click **Save** to add the site.

e) Wait for site to come up and show as `Up` in the Nexus Dashboard UI.

f) Repeat this step for all sites that were previously managed by your Orchestrator.

Step 4 Install Nexus Dashboard Orchestrator service.

This is described in detail in the Deploying Nexus Dashboard Orchestrator, on page 3 chapter of this guide, but if you are already familiar with Orchestrator installation, the following steps summarize the process.

If you are installing the service using the App Store:

a) In the **Services** screen, select the **App Store** tab.

b) In the **Nexus Dashboard Orchestrator** tile, click **Upgrade**.

c) In the License Agreement window that opens, click **Agree and Download**.

If you are installing the service manually:

a) Browse to the Nexus Dashboard Orchestrator page on DC App Center:

https://dcappcenter.cisco.com/nexus-dashboard-orchestrator.html

b) From the **Version** dropdown, choose the version you want to install and click **Download**.

c) Click **Agree and download** to accept the license agreement and download the image.

d) From the left navigation menu in Nexus Dashboard, select **Services**.

e) In the Nexus Dashboard's **Services** screen, select the **Installed Services** tab.

f) From the **Actions** menu in the top right of the main pane, select **Upload Service**.

g) In the **Upload Service** window, choose the location of the image

If you downloaded the application image to your system, choose **Local**.

If you are hosting the image on a server, choose **Remote**.

h) Choose the file.

If you chose **Local** in the previous substep, click **Select File** and select the app image you downloaded.

If you chose **Remote**, provide the full URL to the image file, for example `http://<ip-address>:<port>/<full-path>/cisco-mso-<version>.nap`.

i) Click **Upload** to add the app to the cluster.

It may take a few minutes for the image to be uploaded to the cluster and initialized.

Step 5 Wait for the new image to initialize.

Step 6 In the Nexus Dashboard Orchestrator tile, click **Enable**.

It may take a few minutes for all the application services to start and the GUI to become available.

Step 7 Launch the Orchestrator service.

Simply click **Open** on the service tile.

The single sign-on (SSO) feature allows you to log in to the application using the same credentials as you used for the Nexus Dashboard.

What to do next

After Nexus Dashboard and Nexus Dashboard Orchestrator are installed and all sites are on-boarded, proceed to restore the configuration as described in Restore Configuration, on page 99.

Restore Configuration

This section describes how to deploy and configure the new Nexus Dashboard cluster and the NDO service, which you will use to restore your previous configuration.

Before you begin

You must have the following completed:

- Backed up and downloaded the existing Nexus Dashboard Orchestrator configuration.
- Installed the target Orchestrator release as described in Deploy Nexus Dashboard and Install Nexus Dashboard Orchestrator, on page 97.

Step 1 Ensure that the new Nexus dashboard cluster is up and running and the NDO service is installed.

The NDO service must be a fresh install with no configuration changes to the sites or policies.

Step 2 Open your new Nexus Dashboard Orchestrator service.

Step 3 Add remote location for configuration backups.

This release of Nexus Dashboard Orchestrator does not support configuration backups stored on the cluster's local disk. So before you can import the backup you saved before the migration, you need to configure a remote location in Nexus Dashboard Orchestrator to which you can then import your configuration backups.

a) From the left navigation pane, select **Admin** > **Backup & Restore**.

b) Choose the **Remote Locations** tab.

c) Choose **Create Remote Location**.

The **Create Remote Location** screen appears.

d) Provide the name for the remote location and an optional description.

Two protocols are currently supported for remote export of configuration backups:

- SCP
- SFTP

Note SCP is supported for non-Windows servers only. If your remote location is a Windows server, you must use the SFTP protocol

e) Specify the host name or IP address of the remote server.

Based on your **Protocol** selection, the server you specify must allow SCP or SFTP connections.

f) Provide the full path to a directory on the remote server where you will save the backups.

The path must start with a slash (/) characters and must not contain periods (.) or backslashes (\). For example, */backups/ndo*.

Note The directory must already exist on the remote server.

g) Specify the port used to connect to the remote server.

By default, port is set to 22.

h) Specify the authentication type used when connecting to the remote server.

You can configure one of the following two authentication methods:

- Password—provide the username and password used to log in to the remote server.
- SSH Private Files—provide the username and the SSH Key/Passphrase pair used to log in to the remote server.

i) Click **Save** to add the remote server.

Step 4 Import the backup file to your new Nexus Dashboard Orchestrator cluster.

a) From the left navigation pane, select **Operations** > **Backups & Restore**.
b) In the main pane, click **Upload**.
c) In the **Upload to Remote** window that opens, click **Select File** and choose the configuration backup file you created before the upgrade.
d) From the **Remote Location** dropdown menu, select the remote location.
e) (Optional) Update the remote location path.

The target directory on the remote server, which you configured when creating the remote backup location, will be displayed in the **Remote Path** field.

You can choose to append additional subdirectories to the path. However, the directories must be under the default configured path and must have been already created on the remote server.

f) Click **Import** to upload the file.

Importing a backup will add it to the list of the backups displayed the **Backups** page. Note that even though the backups are shown on the NDO UI, the files are stored only on the remote server and not directly on the cluster nodes.

Step 5 Restore the configuration.

a) In the main window, click the actions (**…**) icon next to the backup you want to restore and select **Rollback to this backup**.

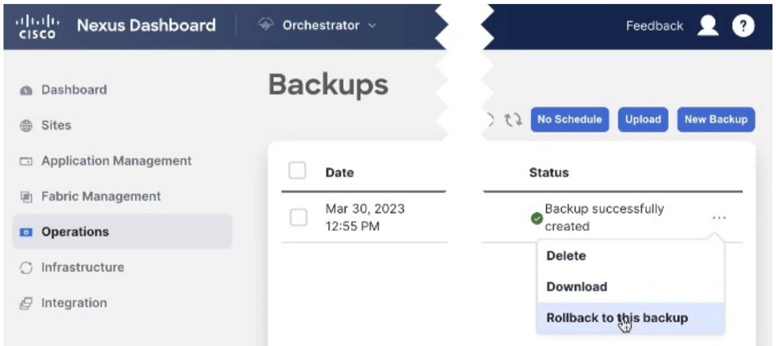

b) In the **Restore from this backup** dialog, read the warning and click **Restore** to confirm that you want to restore the backup you selected.

The restore process imports the backup and checks for any issues, which may take several minutes to complete. After the initial backup import, you will be prompted for additional validation in the next step, which is required for database upgrades from releases prior to release 4.0(1).

c) After the backup import is complete, ensure there are no failures listed in the report, then click **Restore Validation Required** to proceed.

Before the configuration database is updated for this release, the upgrade process performs a number of validations. The validation provides a summary of template and policy changes that will be performed during this final upgrade stage in the next step and includes the following:

- Implicit template stretching – if one or more objects are implicitly stretched, the upgrade process will create new explicitly-stretched templates and move the objects into those templates.

  For example, if you have a template (`t1`) that contains `vrf1` and is associated to `site1` and another template (`t2`) that contains a BD that references `vrf1` but is associated to two sites (`site1` and `site2`), then `vrf1` will be implicitly stretched between the two sites.

  This is no longer allowed starting with release 4.0(1) and the VRF must be explicitly stretched to both sites. In such cases during the upgrade, the VRF will be either moved to a different template which will be explicitly stretched between both sites or the original template will be associated with both sites, depending on whether the other policies in that template require stretching as well.

  Any templates that are created in this case will be named `UpgradeTemplate%d`, where `%d` is an incrementing number starting with 1 to ensure that all newly added templates are unique.

- Global policy migration – all global tenant policies (such as DHCP relay or route maps) and fabric policies (such as QoS) will be moved into the new tenant and fabric policy templates that have been added in release 4.0(1).

Note At this stage, all the tenants have been imported from the backup and created in NDO, but the schemas and templates will be created in the next step.

d) In the **Restore Validation Report** window, click **Restore and Continue** to proceed.

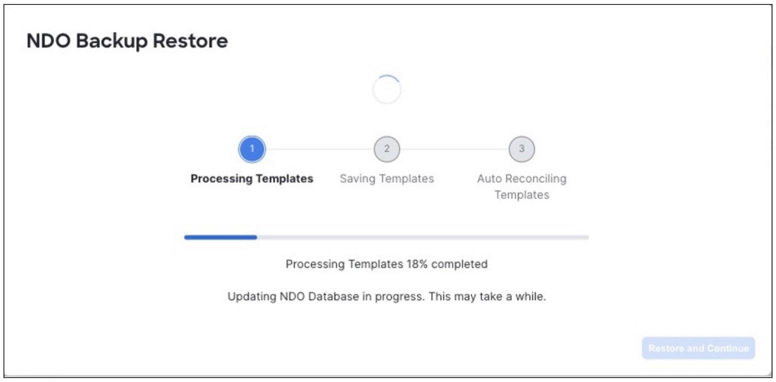

NDO Backup Restore

Processing Templates 18% completed

Updating NDO Database in progress. This may take a while.

This is the stage of the upgrade where the schemas and the templates present in the backup are imported and recreated in your NDO configuration database according to the current best practices. These schemas and templates are then posted to the local NDO database as if it were a greenfield schema/template creation. Then the newly saved templates are deployed in the correct order that conforms to the current deployment requirements and best practices. The template deployment in this step uses a "local deploy" option to calculate the deployment plan and update the database, but does not send any configuration payload to the sites' controllers.

The upgrade process also checks for any configuration drifts between the local NDO database (configuration that is correct from NDO's point of view) and what is actually deployed in the fabrics. If this release of NDO supports additional objects or properties compared to the release from which you are upgrading, the upgrade will automatically reconcile those drifts by importing the existing configuration from the site's controller.

Note that if a template is automatically reconciled, two template versions are created – one before the automatic reconciliation and one after:

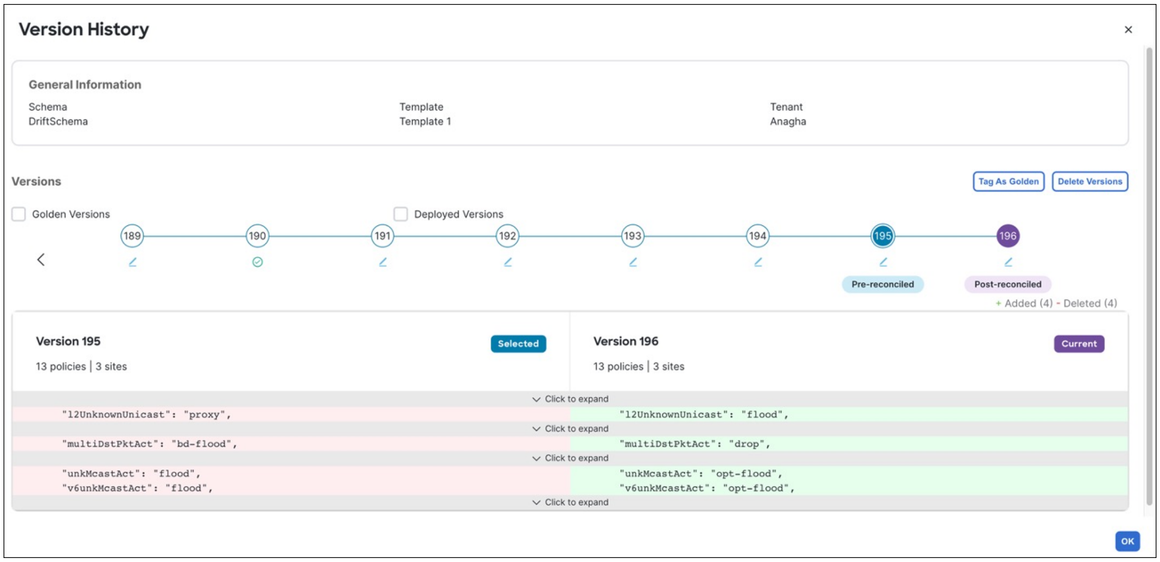

Version History

General Information

Schema	Template	Tenant
DriftSchema	Template 1	Anagha

Versions

Version 195

13 policies | 3 sites

Version 196

13 policies | 3 sites

```
"l2UnknownUnicast": "proxy",                "l2UnknownUnicast": "flood",
"multiDstPktAct": "bd-flood",               "multiDstPktAct": "drop",
"unkMcastAct": "flood",                     "unkMcastAct": "opt-flood",
"v6unkMcastAct": "flood",                   "v6unkMcastAct": "opt-flood",
```

e) Review the report from the previous substep and click **Ok** to finish.

The final stage of the database upgrade presents a full report of the performed actions for you to review. If you close the report but want to review it again, simply click the **View Restore Report** in the **Backups** page.

Step 6 Verify that backup was restored successfully and all objects and configurations are present.

a) In the **Sites** page, verify that all sites are listed as `Managed`.

b) In the **Tenants** and **Schemas** pages, confirm that all tenants and schemas from your previous Nexus Dashboard Orchestrator cluster are present.

c) Navigate to **Infrastructure** > **Site Connectivity** and confirm that intersite connectivity is intact.

In the main pane, click **Show Connectivity Status** next to each site and verify that the existing tunnels are up and connectivity was not interrupted.

d) In the main pane, click **Configure** to open **Fabric Connectivity Infra** screen and verify **External Subnet Pool** addresses.

You can view the external subnet pools by selecting **General Settings** > **IPSec Tunnel Subnet Pools** tab of the **Fabric Connectivity Infra** screen and verify that the External Subnet Pools previously configured in Cloud Network Controller have been imported from the cloud sites.

These subnets are used to address the IPsec tunnel interfaces and loopbacks of the Cloud Routers used for on-premises connectivity and had to be configured directly in the Cloud Network Controller in earlier Nexus Dashboard Orchestrator releases.

Resolve Configuration Drifts

In some cases you may run into a situation where the configuration actually deployed in the site's controller is different from the configuration defined in the Nexus Dashboard Orchestrator. These configuration discrepancies are referred to as **Configuration Drifts** and are indicated by an `Out of Sync` warning next to the site name in the template view page

After upgrading your Nexus Dashboard Orchestrator and restoring the previous configuration backup, we recommend that you check for and resolve any configuration drifts that were not automatically resolved by the upgrade process as described in this section.

Note Deploying any templates before resolving configuration drifts would push the configuration defined in the Orchestrator and overwrite the values defined in the fabrics' controllers.

Step 1 In your Nexus Dashboard Orchestrator, navigate to **Operate** > **Tenant Templates**.

Step 2 Choose the **Applications** tab.

Step 3 Select the first schema and check its templates for configuration drifts.

You will repeat the following steps for every schema and template in your deployment

You can check for configuration drifts in one of the following two ways:

- Check the template deployment status icon for each site to which the template is assigned.
- Select the template and click **Deploy to sites** to bring up the configuration comparison screen to check which objects contain configuration drifts.

Step 4 If the template contains a configuration drift, resolve the conflicts.

For more information about configuration drifts, check the "Configuration Drifts" chapter in the *Cisco Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics*.

a) Close the template deployment dialog to return to the Schema view.

   Deploying any templates at this point would push the values in the Orchestrator database and overwrite any existing settings in the fabrics.

b) From the template's **Actions** menu, select **Drift Reconciliation**.

   The **Drift Reconciliation** wizard opens.

c) In the **Drift Reconciliation** screen, compare the template-level configurations for each site and choose the one you want.

   Template-level properties are common across all sites associated to the template. You can compare the template level properties defined on Nexus Dashboard Orchestrator with the configuration rendered in each site and decide what should become the new configuration in the Nexus Dashboard Orchestrator template. Selecting the site configuration will modify those properties in the existing Nexus Dashboard Orchestrator template, whereas selecting the Nexus Dashboard Orchestrator configuration will keep the existing Nexus Dashboard Orchestrator template settings as is

d) Click **Go to Site Specific Properties** to switch to site-level configuration.

   You can choose a site to compare that specific site's configuration. Unlike template-level configurations, you can choose either the Nexus Dashboard Orchestrator-defined or actual existing configurations for each site individually to be retained as the template's site-local properties for that site.

   Even though in most scenarios you will make the same choice for both template-level and site-level configuration, the drift reconciliation wizard allows you to choose the configuration defined in the site's controller at the "Template Properties" level and the configuration defined in Nexus Dashboard Orchestrator at the "Site Local Properties" level or vice versa.

e) Click **Preview Changes** to verify your choices.

   The preview will display full template configuration adjusted based on the choices picked in the **Drift Reconciliation** wizard

f) Save the schema.

g) Click **Deploy to sites** to deploy the configuration and finish reconciling the drift for that template

Step 5 Repeat the above steps for every schema and template in your Nexus Dashboard Orchestrator.

Step 6 Check audit logs to verify that all templates have been re-deployed.

You can view the audit logs in the **Operations** tab.

Audit Logs page and confirm that all templates show as `Redeployed` to ensure that full re-deployment successfully completed.