

Cisco Nexus Dashboard Orchestrator Deployment Guide, Release 3.7(x)

First Published: 2022-03-14

Last Modified: 2022-12-19

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: https://www.cisco.com/c/en/us/about/legal/cloud-and-software/software-terms.html. Cisco product warranty information is available at https://www.cisco.com/c/en/us/products/warranty-listing.html. US Federal Communications Commission Notices are found here https://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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2022 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1 **New and Changed Information** 1

New and Changed Information 1

CHAPTER 2 **Deploying Nexus Dashboard Orchestrator** 3

Deployment Overview 3

Prerequisites and Guidelines 4

Hardware Requirements For ACI Fabrics 5

Hardware Requirements For DCNM Fabrics 7

Installing Nexus Dashboard Orchestrator Service Using App Store 8

Installing Nexus Dashboard Orchestrator Service Manually 9

PART I **Day-0 Operations for ACI Fabrics** 13

CHAPTER 3 **Configuring Cisco ACI Sites** 15

Pod Profile and Policy Group 15

Configuring Fabric Access Policies for All APIC Sites 16

Configuring Fabric Access Global Policies 16

Configuring Fabric Access Interface Policies 17

Configuring Sites That Contain Remote Leaf Switches 19

Remote Leaf Guidelines and Limitations 19

Configuring Routable Subnets for Remote Leaf Switches 19

Enabling Direct Communication for Remote Leaf Switches 20

Cisco Mini ACI Fabrics 20

CHAPTER 4 **Adding and Deleting Sites** 23

Cisco NDO and APIC Interoperability Support 23

Adding Cisco ACI Sites 25
Removing Sites 27
Cross Launch to Fabric Controllers 28

CHAPTER 5 **Configuring Infra General Settings** 31
Infra Configuration Dashboard 31
Partial Mesh Intersite Connectivity 32
Configuring Infra: General Settings 33

CHAPTER 6 **Configuring Infra for Cisco APIC Sites** 37
Refreshing Site Connectivity Information 37
Configuring Infra: On-Premises Site Settings 38
Configuring Infra: Pod Settings 40
Configuring Infra: Spine Switches 41

CHAPTER 7 **Configuring Infra for Cisco Cloud APIC Sites** 45
Refreshing Cloud Site Connectivity Information 45
Configuring Infra: Cloud Site Settings 45

CHAPTER 8 **Deploying Infra Configuration for ACI Sites** 49
Deploying Infra Configuration 49
Enabling Connectivity Between On-Premises and Cloud Sites 50

PART II **Day-0 Operations for DCNM Fabrics** 55

CHAPTER 9 **Adding and Deleting Sites** 57
Adding Cisco DCNM Sites 57
Removing Sites 60
Cross Launch to Fabric Controllers 61

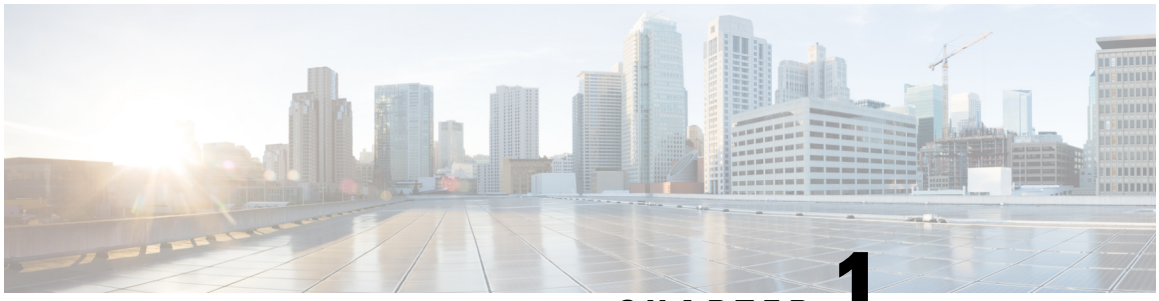
CHAPTER 10 **Configuring Infra for Cisco DCNM Sites** 63
Prerequisites and Guidelines 63
Configuring Infra: General Settings 63

Refreshing Site Connectivity Information 64
Configuring Infra: DCNM Site Settings 65
Deploying Infra Configuration 67

PART III **Upgrading Nexus Dashboard Orchestrator** 69

CHAPTER 11 **Upgrading NDO Service in Nexus Dashboard** 71
Overview 71
Prerequisites and Guidelines 71
Upgrading NDO Service Manually 73
Restore Existing Configuration for Database Optimization 75
Resolve Configuration Drifts 76

CHAPTER 12 **Migrating Existing MSO Cluster to Nexus Dashboard** 83
Overview 83
Prerequisites and Guidelines 84
Back Up Existing Cluster Configuration 86
Prepare New Cluster 87
Restore Configuration in the New Cluster 90
Upgrade Cloud Sites 94
Update NDO Infra Configuration for Cloud Sites 96
Resolve Configuration Drifts 98



CHAPTER 1

New and Changed Information

• New and Changed Information, on page 1

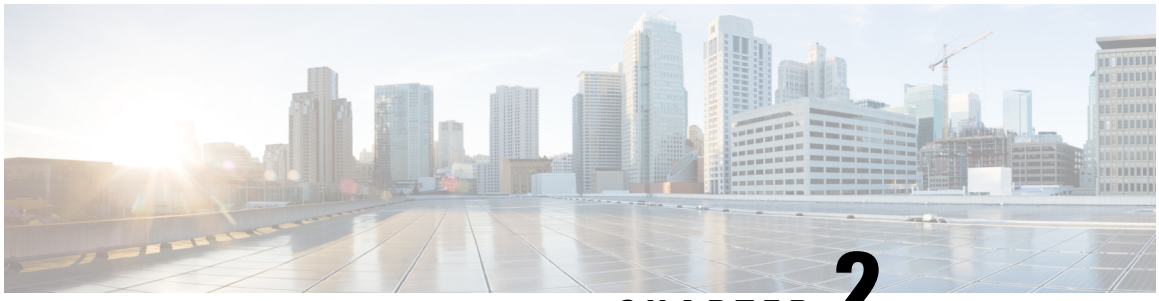
New and Changed Information

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
3.7(1)	First release of this document.	--

New and Changed Information



CHAPTER 2

Deploying Nexus Dashboard Orchestrator

- Deployment Overview, on page 3
- Prerequisites and Guidelines, on page 4
- Installing Nexus Dashboard Orchestrator Service Using App Store, on page 8
- Installing Nexus Dashboard Orchestrator Service Manually, on page 9

Deployment Overview

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

Note If you are upgrading from a release prior to release 3.2(1), familiarize yourself with deployment overview described in this section, then follow the instructions in Migrating Existing MSO Cluster to Nexus Dashboard, on page 83.

Cisco Nexus Dashboard is a central management console for multiple data center sites and a common platform for hosting Cisco data center services, such as the Nexus Dashboard Orchestrator or Nexus Insights. Nexus Dashboard provides a common platform and modern technology stack for these micro-services-based services, simplifying the life cycle management of the different modern services and reducing the operational overhead to run and maintain those services.

Each Nexus Dashboard cluster consists of 3 `master` nodes. You can also deploy additional `worker` nodes to enable horizontal scaling and a `standby` node for easy cluster recovery in case of a master node failure.

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 DCNM* for your release, 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 3.7(1) and later	Cisco Nexus Dashboard, Release 2.1.1 or later We recommend deploying in Nexus Dashboard, Release 2.2.1 or later for all new installations.

Nexus Dashboard Orchestrator Image Format

Starting with Nexus Dashboard Orchestrator, Release 3.7(1), the Orchestrator services is delivered using a new `.nap` image format which allows the service to provide additional features and greatly reduces initial deployment time. For all new Nexus Dashboard Orchestrator deployments and upgrades to release 3.7(1) or later, you must deploy in Nexus Dashboard release 2.1.1 or later using the new image format.

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 DCNM • Any other remote devices or controllers	Data network
Intra-cluster communication	Data network

NDO Traffic Type	Nexus Dashboard 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)

Nexus Dashboard Orchestrator uses NTP for clock synchronization, so you must have an NTP server configured in your environment.

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 Multi-Site 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).

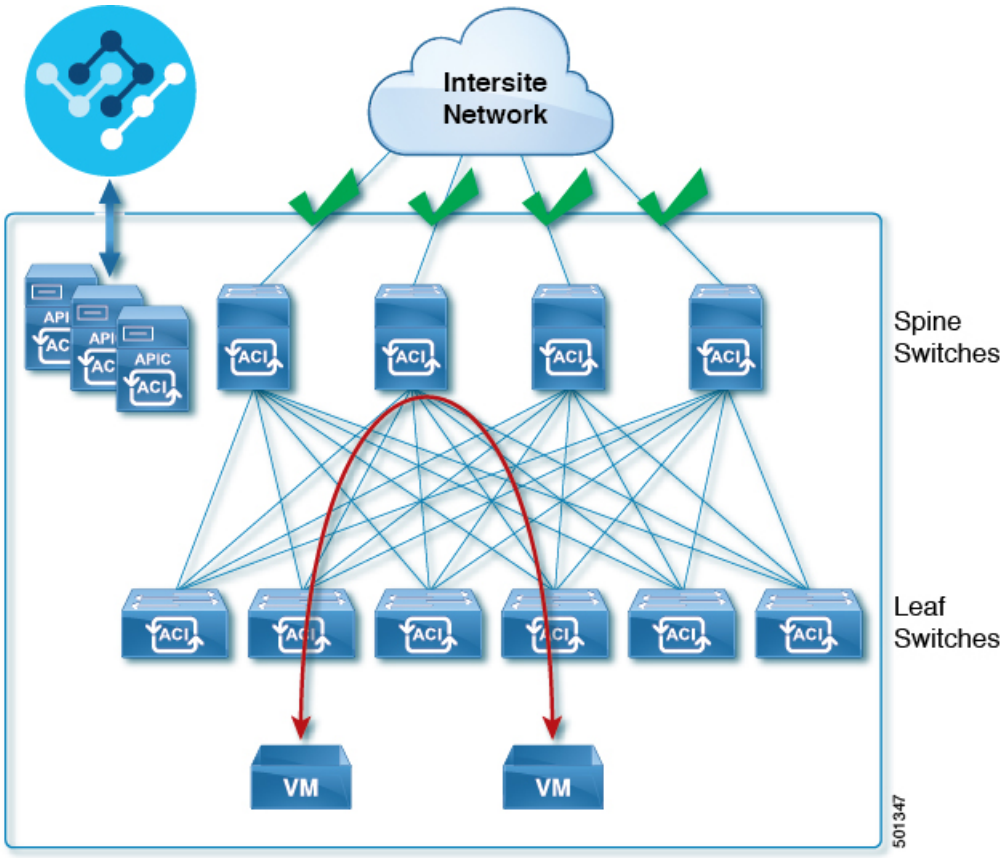
Refer to the ACI-mode Switches Hardware Support Matrix for the complete list of supported spines for each release.

Leaf Switch Requirements

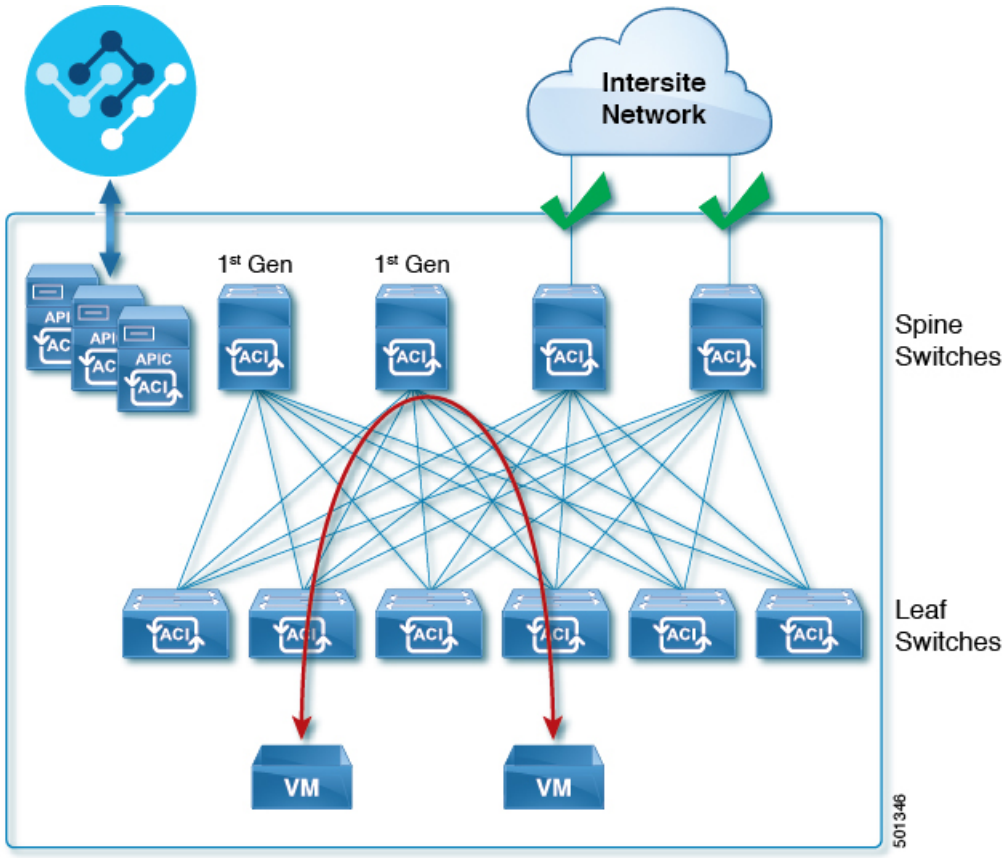
Multi-Site has no dependency on the fabrics' leaf switches and as such supports the same leaf switch models as the Cisco APIC. The full list of supported hardware is available in the ACI-mode Switches Hardware Support Matrix.

IPN Connectivity Across Sites

The following figure shows how spine switches supported with ACI 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 DCNM 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 Nexus Dashboard Orchestrator Service Using App Store

This section describes how to 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.
- 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*.

  If you are unable to establish the connection to the DC App Center, skip this section and follow the steps 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 prior to Release 3.3(1), see the *Nexus Dashboard Orchestrator Installation Guide* specific to that release for the available deployment options and procedures.

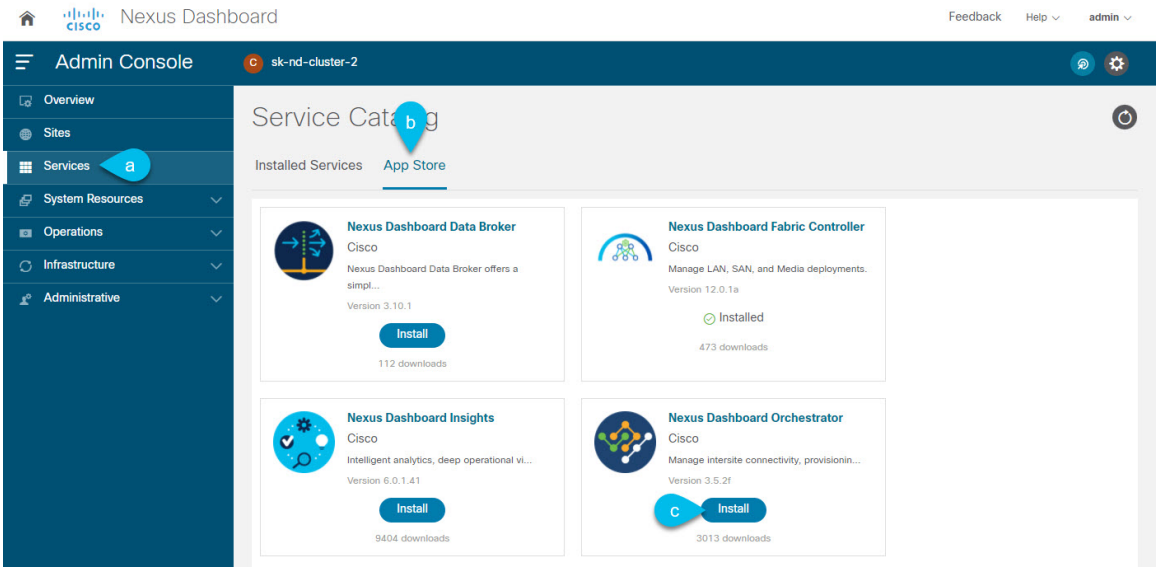
Procedure

Step 1 Log in to the Nexus Dashboard GUI.

Step 2 From the left navigation menu, select **Admin Console**.

You must have `admin` privileges to deploy services.

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



a) From the left navigation menu, select **Service Catalog**.
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 the Nexus Dashboard and deployed.

Step 6 Enable the app.

After installation is complete, the service will remain in the `Disabled` state by default and you must enable it.

To enable the app, click the **...** menu on the app and select **Enable**.

Step 7 Launch the app.

To launch the app, simply click **Open** on the service tile in the Nexus Dashboard's **Service Catalog** page.

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

Procedure

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

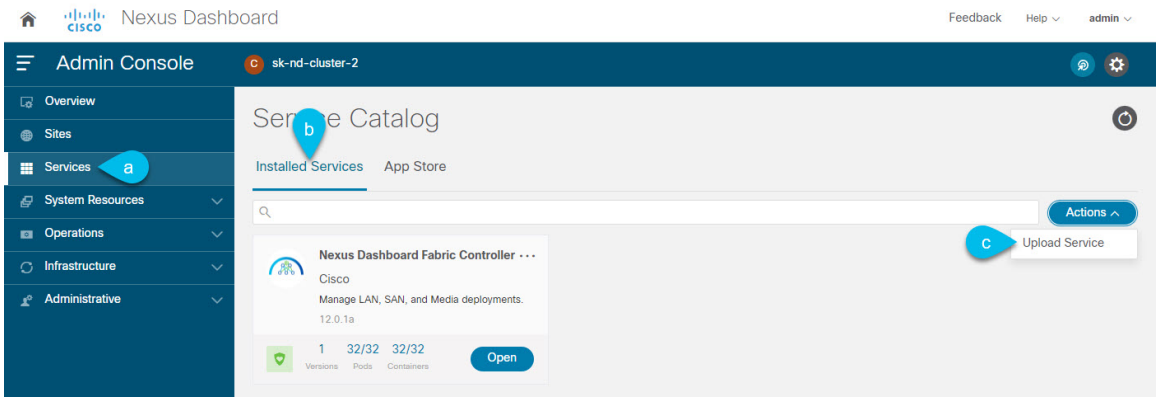
b) From the **Version** dropdown, choose the version you want to install and click **Download**.

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

Step 2 Log in to your Cisco Nexus Dashboard dashboard.

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

Step 3 Choose to manually upload the image.



a) In the left navigation bar, click **Service Catalog**.

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

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

Step 4 Select the image file to upload.

a) 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**.

b) Choose the file.

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

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

Step 5 Wait for the service to be downloaded to the Nexus Dashboard and deployed.

Step 6 Enable the service.

After installation is complete, the service will remain in the `Disabled` state by default and you must enable it.

To enable the service, click the **...** menu and select **Enable**.

Step 7 Launch the service.

To launch the service, simply click **Open** on the service tile in the Nexus Dashboard's **Service Catalog** page.

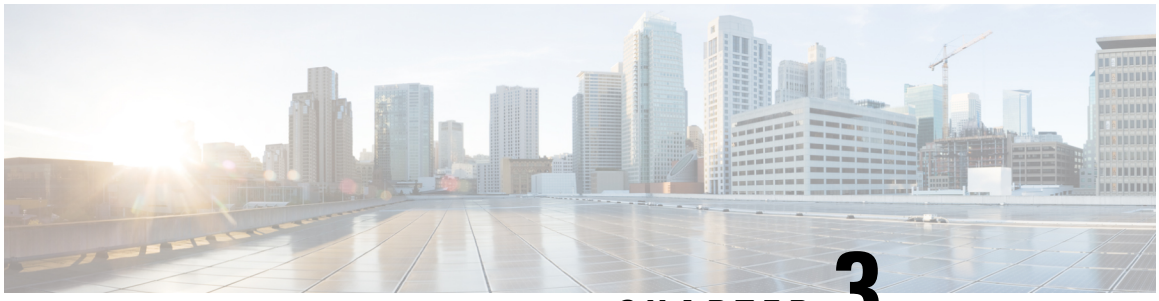
The single sign-on (SSO) feature allows you to log in to the service using the same credentials as you used for the Nexus Dashboard.



PART I

Day-0 Operations for ACI Fabrics

- Configuring Cisco ACI Sites, on page 15
- Adding and Deleting Sites, on page 23
- Configuring Infra General Settings, on page 31
- Configuring Infra for Cisco APIC Sites, on page 37
- Configuring Infra for Cisco Cloud APIC Sites, on page 45
- Deploying Infra Configuration for ACI Sites, on page 49



CHAPTER 3

Configuring Cisco ACI Sites

- Pod Profile and Policy Group, on page 15
- Configuring Fabric Access Policies for All APIC Sites, on page 16
- Configuring Sites That Contain Remote Leaf Switches, on page 19
- Cisco Mini ACI Fabrics, on page 20

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.

Procedure

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.

Procedure

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 17.

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 16.

Procedure

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.

Procedure

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

Procedure

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.

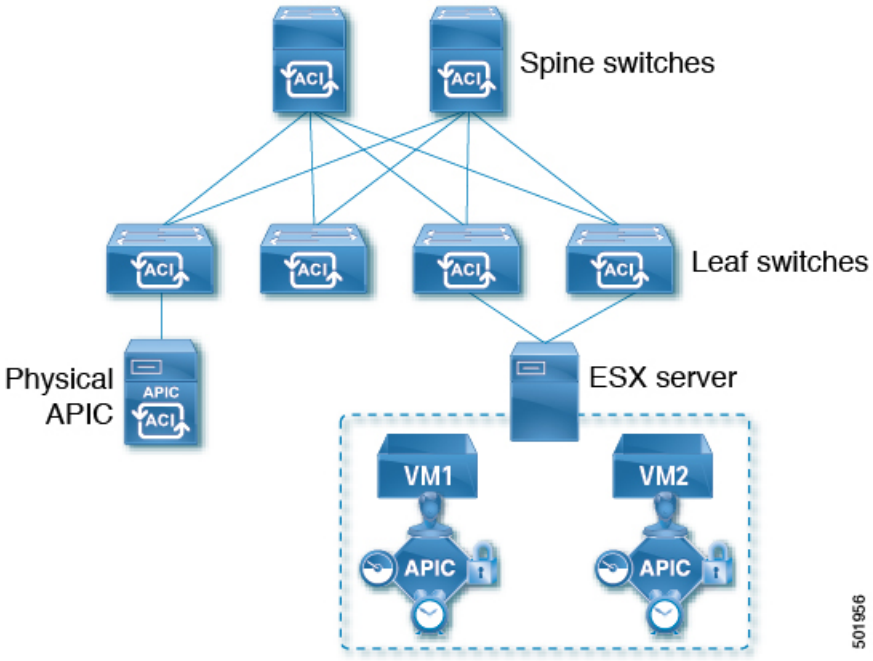
Cisco Mini ACI Fabrics

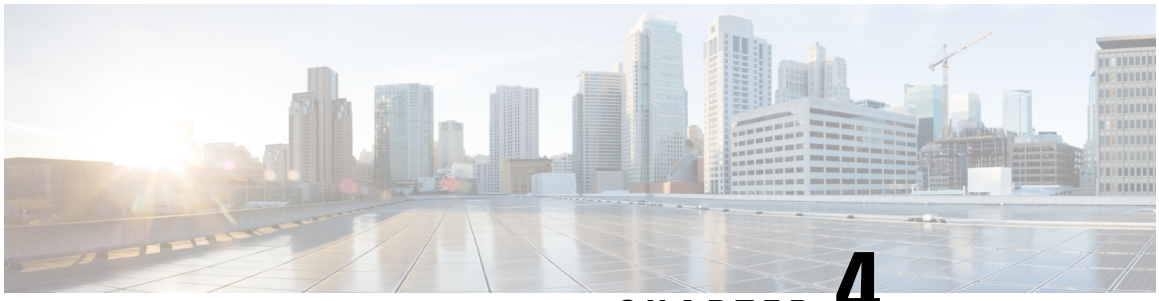
Cisco Multi-Site supports Cisco Mini ACI fabrics as typical on-premises sites without requiring any additional configuration. This section provides a brief overview of Mini ACI fabrics, detailed info on deploying and configuring this type of fabrics is available in *Cisco Mini ACI Fabric and Virtual 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, on page 23
- Adding Cisco ACI Sites, on page 25
- Removing Sites, on page 27
- Cross Launch to Fabric Controllers, on page 28

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 APIC site using the Nexus Dashboard GUI and then enable that site to be managed by Nexus Dashboard Orchestrator.

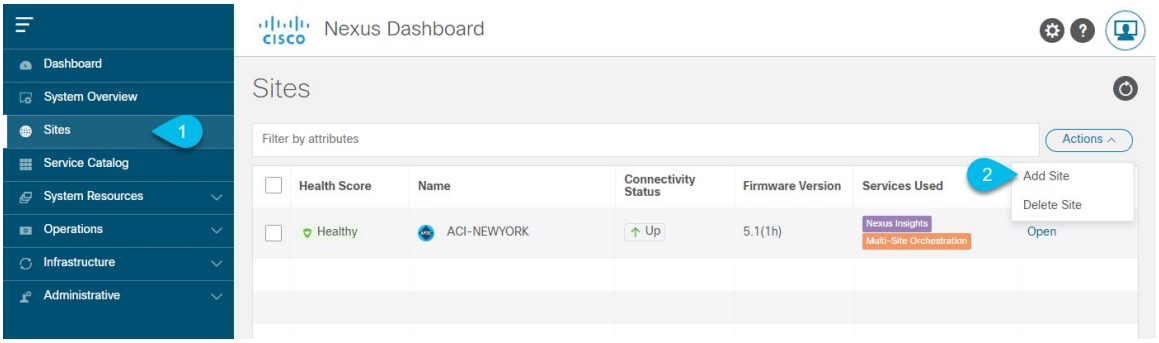
Before you begin

- 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 the site(s) you are adding are running Release 4.2(4) or later.

Procedure

Step 1 Log in to the Nexus Dashboard GUI

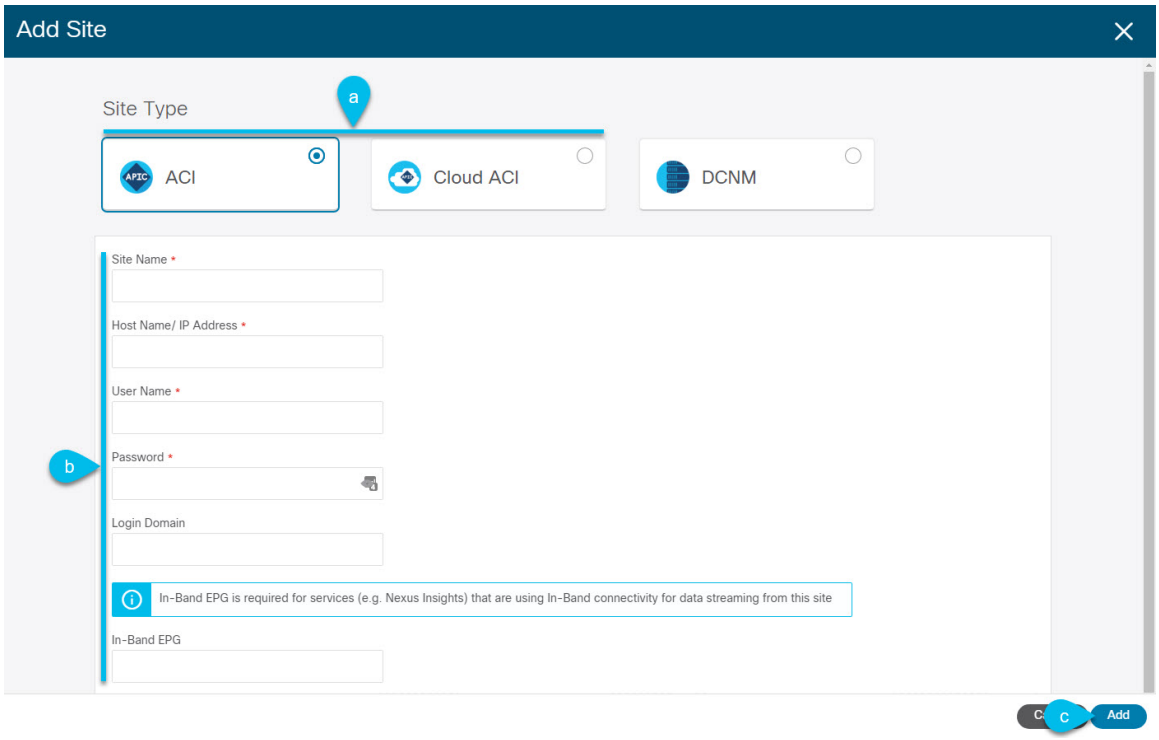
Step 2 Add a new site.



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

b) In the top right of the main pane, select **Actions** > **Add Site**.

Step 3 Provide site information.



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

b) Provide the controller information.

- You need to 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 will use the site with Nexus Dashboard Orchestrator service only, you can provide either the in-band or out-of-band IP address of the APIC. If you will use the site with 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 Nexus Insights, you must also provide the **In-Band EPG** name used to connect the Nexus Dashboard to the fabric you are adding. Otherwise, if you will use this site with Nexus Dashboard Orchestrator only, you can leave this field blank.

- For cloud ACI sites, **Enable Proxy** if your cloud site is reachable via a proxy.

  Proxy must be already configured in your Nexus Dashboard's cluster settings. If the proxy is reachable via 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 **Add** to finish adding the site.

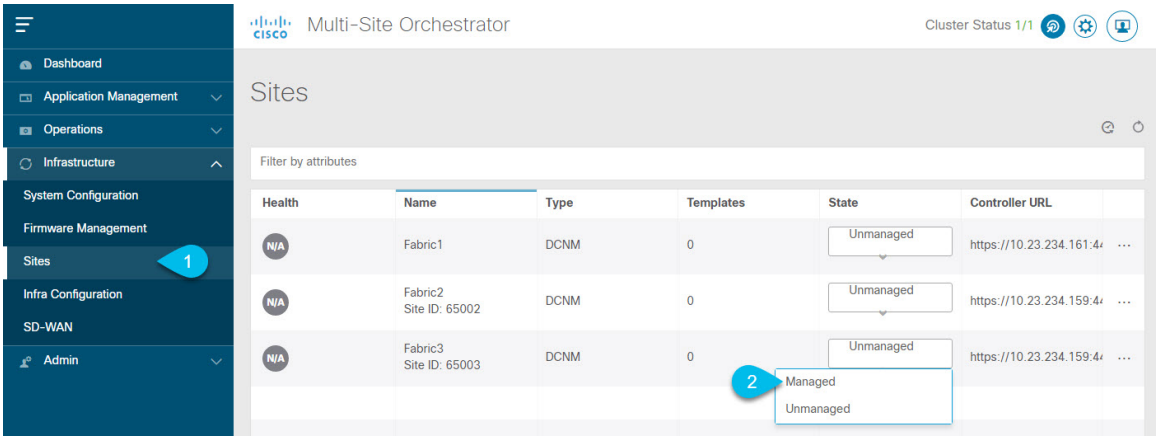
At this time, the sites will be available in the Nexus Dashboard, but you still need to enable them for Nexus Dashboard Orchestrator management as described in the following steps.

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

Step 5 From the Nexus Dashboard's **Service Catalog**, 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.



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.

Note
Ensure that ACI site names are limited to 125 characters or less to avoid any issues when enabling orchestration.

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.

Procedure

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) In the left sidebar's **Sites** area, select a template associated with the site, click the options menu (**...**) next to the template, and choose **Undeploy Template**.

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

Note
For non-stretched templates, you can choose to preserve the configuration by selecting **Dissociate Template** instead of **Undeploy Template**, but you must undeploy any stretched templates.

d) 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 **Infrastructure** > **Site Connectivity**.
b) In the main pane, click **Configure**.
c) In the left sidebar, select the site you want to unmanage.
d) In right sidebar's **Inter-Site Connectivity** tab, disable the **Multi-Site** checkbox.
e) Click **Deploy** to deploy the changes to the site.

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

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

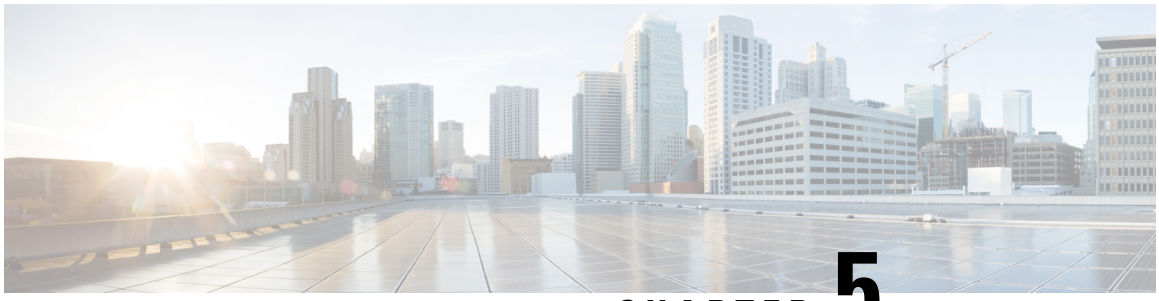
The site will be removed from the Nexus Dashboard.

Cross Launch to Fabric Controllers

Nexus Dashboard Orchestrator currently supports a number of configuration options for each type of fabrics. For many additional configuration options, you may need to log in directly into the fabric's controller.

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

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



CHAPTER 5

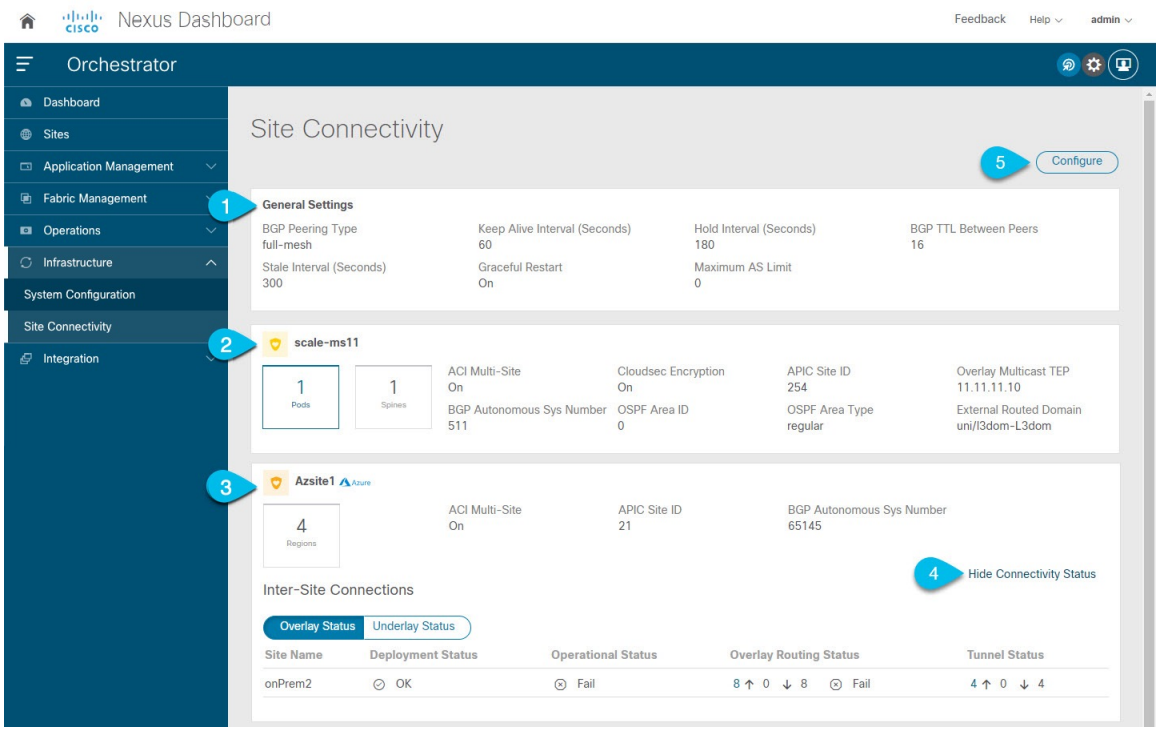
Configuring Infra General Settings

- Infra Configuration Dashboard, on page 31
- Partial Mesh Intersite Connectivity, on page 32
- Configuring Infra: General Settings, on page 33

Infra Configuration Dashboard

The **Infra Configuration** page displays an overview of all sites and inter-site connectivity in your Nexus Dashboard Orchestrator deployment and contains the following information:

Figure 2: Infra Configuration Overview



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 on 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 37.

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 APIC Sites, on page 45.

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 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 Nexus Dashboard Orchestrator fabric connectivity information refresh described in the Refreshing Site Connectivity Information, on page 37 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 APIC 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 APIC sites). Ensure that you complete all the required configurations in infra general settings before proceeding to the site-local settings specific to each site.

Procedure

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

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

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

Step 4 In the left sidebar, select **General Settings**.

Step 5 Provide **Control Plane Configuration**.

a) Select the **Control Plane Configuration** tab.

b) Choose **BGP Peering Type**.

- `full-mesh`—All border gateway switches in each site will establish peer connectivity with remote sites' border gateway switches.

  In `full-mesh` configuration, Nexus Dashboard Orchestrator uses the spine switches for ACI managed fabrics and border gateways for DCNM 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 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 keep alive 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 APIC sites, this field will not be present in the UI.

This is OSPF area ID used by cloud sites for on-premises IPN peering, which you previously configured in the Cloud APIC for inter-site connectivity in earlier Nexus Dashboard Orchestrator releases.

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

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

When you configure inter-site underlay connectivity between on-premises and cloud sites as described in later sections, you will need to select an on-premises IPN device which will establish connectivity to the cloud CSRs. These IPN devices must first be defined here before they are available in the on-premises site configuration screen, which is described in more detail in Configuring Infra: On-Premises Site Settings, on page 38.

a) Select the **IPN Devices** tab.
b) Click **Add IPN Device**.
c) Provide the **Name** and the **IP Address** of the IPN device.

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

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

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

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

If you do not have any Cloud APIC 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 will open.

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

The IP address you provide will be used as the tunnel peer address from the Cloud APIC'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 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 APIC 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 Nexus Dashboard Orchestrator. The Orchestrator, creates smaller subnets from these pools and allocates them to sites to be used for inter-site IPsec tunnels and external connectivity IPsec tunnels.

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

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

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

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

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

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

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

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

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

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

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

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

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

c) Provide the subnet **Name**.

   You will be able to use the subnet pool's name to choose the pool from which to allocate the IP addresses later on.

d) Click **+Add IP Address** to add one or more subnet pools.

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

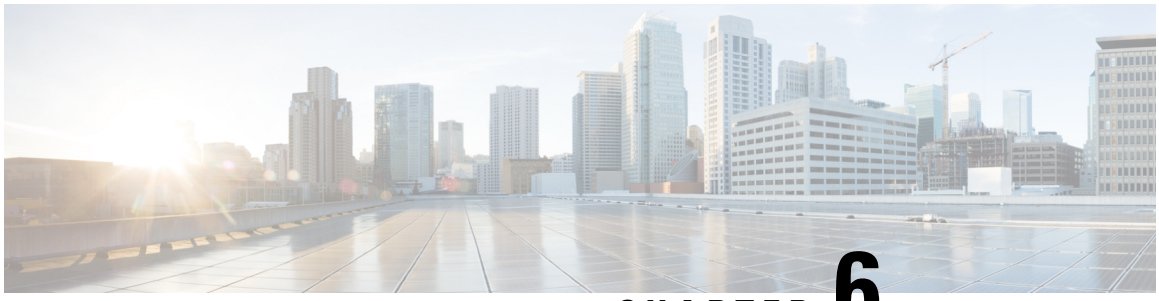
e) Click the check mark icon to save the subnet information.

Repeat the steps if you want to add multiple subnets to the same named subnet pool.

f) Click **Save** to save the named subnet pool.
g) Repeat these substeps for any additional named subnet pools you want to add.

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 (on-premises ACI, cloud ACI, or on-premises fabric managed by DCNM) 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, on page 37
- Configuring Infra: On-Premises Site Settings, on page 38
- Configuring Infra: Pod Settings, on page 40
- Configuring Infra: Spine Switches, on page 41

Refreshing Site Connectivity Information

Any infrastructure changes, such as adding and removing spines or changing spine node IDs, require a Multi-Site fabric connectivity site refresh. This section describes how to pull up-to-date connectivity information directly from each site's APIC.

Procedure

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

Step 2 In the left navigation menu, select **Infrastructure** > **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 checkbox, all configuration info for any currently decommissioned spine switches will be removed from the database.

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

This will discover any new or removed spines and all site-related fabric connectivity will be re-imported from the APIC.

Configuring Infra: On-Premises Site Settings

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

Procedure

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

Step 2 In the left navigation menu, select **Infrastructure** > **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 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 inter-site 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 inter-site L2 BUM and L3 multicast traffic. This IP address is deployed on all spine switches that are part of the same fabric, regardless of whether it is a single pod or Multi-Pod fabric.

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

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 41.

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

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 41.

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 41.

You can either click an existing policy (for example, `msc-ospf-policy-default`) to modify it or click **+Add Policy** to add a new OSPF policy. Then in the **Add/Update Policy** window, specify the following:

- In the **Policy Name** field, enter the policy name.
- In the **Network Type** field, choose either `broadcast`, `point-to-point`, or `unspecified`.

  The default is `broadcast`.
- In the **Priority** field, enter the priority number.

  The default is `1`.
- In the **Cost of Interface** field, enter the cost of interface.

  The default is `0`.
- From the **Interface Controls** dropdown menu, choose one of the following:
  - **advertise-subnet**
  - **bfd**
  - **mtu-ignore**
  - **passive-participation**
- In the **Hello Interval (Seconds)** field, enter the hello interval in seconds.

  The default is `10`.
- In the **Dead Interval (Seconds)** field, enter the dead interval in seconds.

  The default is `40`.
- In the **Retransmit Interval (Seconds)** field, enter the retransmit interval in seconds.

  The default is `5`.
- In the **Transmit Delay (Seconds)** field, enter the transmit delay in seconds.

  The default is `1`.

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

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

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 will need to configure additional 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 via an MPLS network, enable the **SR-MPLS Connectivity** knob and provide the Segment Routing global block (SRGB) range.

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

The default range is `16000-23999`.

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

Step 6 Configure inter-site connectivity between on-premises and cloud sites.

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

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

a) Click **+Add IPN Device** to specify an IPN device.
b) From the dropdown, select one of the IPN devices you defined previously.

The IPN devices must be already defined in the **General Settings** > **IPN Devices** list, as described in Configuring Infra: General Settings, on page 33

c) Configure inter-site connectivity for cloud sites.

Any previously configured connectivity from the cloud sites to this on-premises site will be displayed here, but any additional configuration must be done from the cloud site's side as described in Configuring Infra for Cisco Cloud APIC Sites, on page 45.

What to do next

While you have configured all the required inter-site connectivity information, it has not been pushed to the sites yet. You need to deploy the configuration as described in Deploying Infra Configuration, on page 49

Configuring Infra: Pod Settings

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

Procedure

Step 1 Log in to the Cisco Nexus Dashboard Orchestrator GUI.
Step 2 In the left navigation menu, select **Infrastructure** > **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.

Prior to 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 inter-site underlay connectivity. We recommend configuring either OSPF or BGP and not both, however if you configure both protocols, BGP will take precedence and OSPF will not be installed in the route table.

Procedure

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

Step 2 In the left navigation menu, select **Infrastructure** > **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 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 sub-interface with VLAN 4 with the specified IP ADDRESS under the specified PORT.

- In the **MTU** field, enter the MTU. You can specify either `inherit`, which would configure an MTU of 9150B, or choose a value between `576` and `9000`.

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

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 required in the next substep.

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

  The OSPF settings will become available.

- From the **OSPF Policy** dropdown, select the OSPF policy for the switch that you have configured in Configuring Infra: On-Premises Site Settings, on page 38.

  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 APIC 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 will 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.

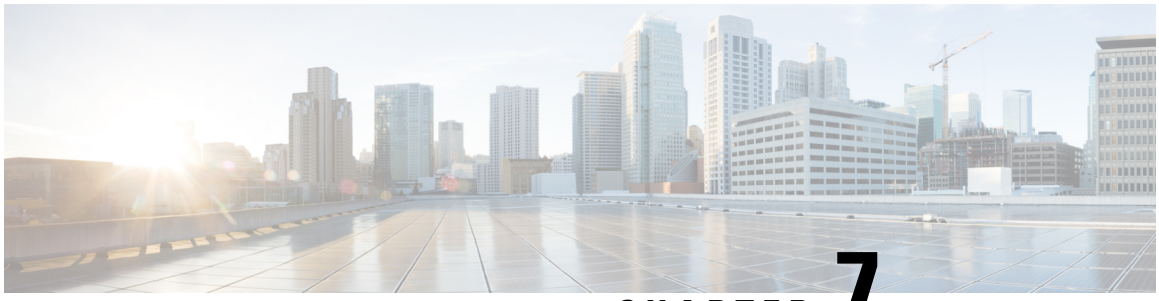
- In the **Peer AS Number** field, provide the Autonomous System (AS) number of the BGP neighbor.

  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 APIC Sites

- Refreshing Cloud Site Connectivity Information, on page 45
- Configuring Infra: Cloud Site Settings, on page 45

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.

Procedure

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

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

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

Step 4 In the left 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 will discover 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 need to 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 APIC sites.

Procedure

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

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

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

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

Step 5 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.

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

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

Step 6 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`) .

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

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

The following options are available:

- `Public Internet`—connectivity between the two sites is established via 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 inter-site 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 inter-site 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 inter-site connectivity from `Site1` to `Site3`, you must repeat this step for that site as well.

Step 7 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 will open.

c) From the **VRF** dropdown, 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 will display the cloud regions that contain the CSRs to which this configuration be applied.

d) From the **Name** dropdown in the **External Devices** section, select the external device.

This is the external device 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 33.

e) From the **Tunnel IKE Version** dropdown, 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** dropdown, choose one of the named subnet pools.

Named subnet pool 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 33.

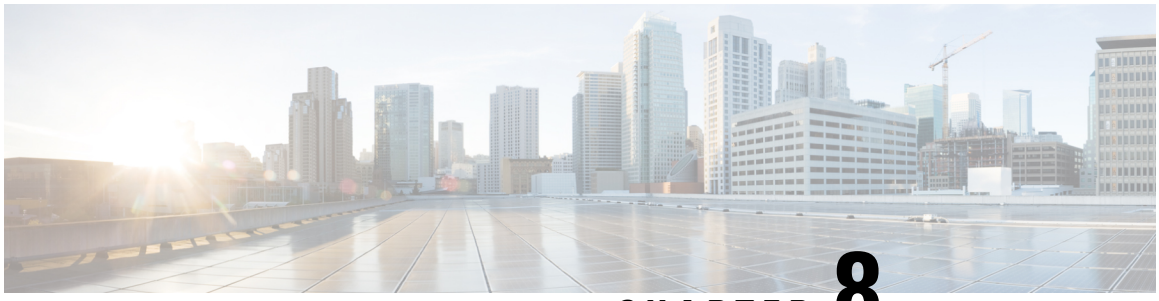
g) (Optional) In the **Pre-Shared Key** field, provide the custom keys you want to use to establish the tunnel.

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).

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

What to do next

While you have configured all the required inter-site connectivity information, it has not been pushed to the sites yet. You need to deploy the configuration as described in Deploying Infra Configuration, on page 49



CHAPTER 8

Deploying Infra Configuration for ACI Sites

- Deploying Infra Configuration, on page 49
- Enabling Connectivity Between On-Premises and Cloud Sites, on page 50

Deploying Infra Configuration

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

Procedure

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 APIC 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 APIC 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 37.

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 APIC sites.

By default, the Cisco Cloud APIC 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.

Procedure

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 49.

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 APIC, 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 144904725321902286651389219409672714611O
    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.

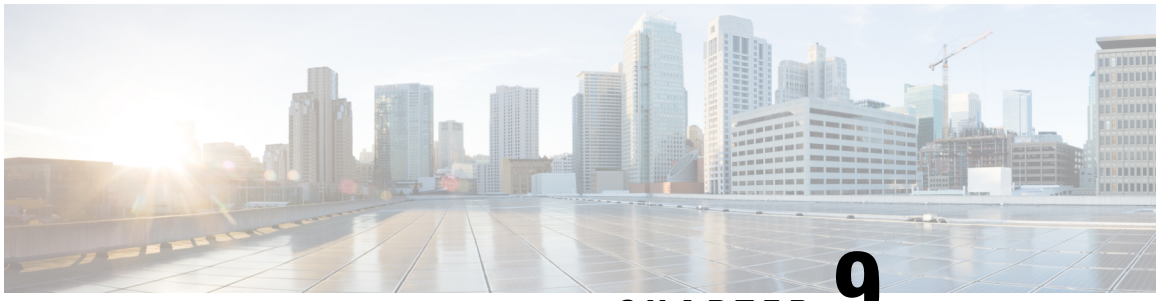
```
ISN_CSR# show ip interface brief | include Tunnel
Interface              IP-Address      OK? Method Status                Protocol
Tunnel1000             30.29.1.2         YES manual up                    up
Tunnel1001             30.29.1.4         YES manual up                    up
```



PART II

Day-0 Operations for DCNM Fabrics

- Adding and Deleting Sites, on page 57
- Configuring Infra for Cisco DCNM Sites, on page 63



CHAPTER 9

Adding and Deleting Sites

- Adding Cisco DCNM Sites, on page 57
- Removing Sites, on page 60
- Cross Launch to Fabric Controllers, on page 61

Adding Cisco DCNM Sites

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

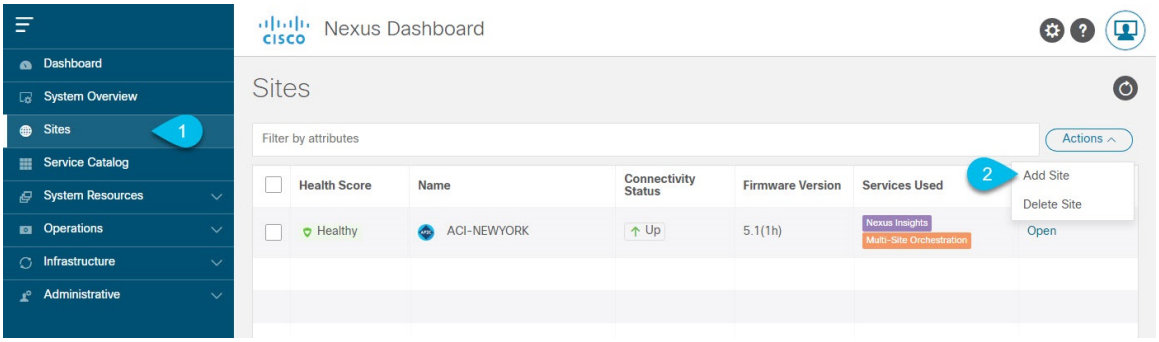
Before you begin

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

Procedure

Step 1 Log in to the Nexus Dashboard GUI

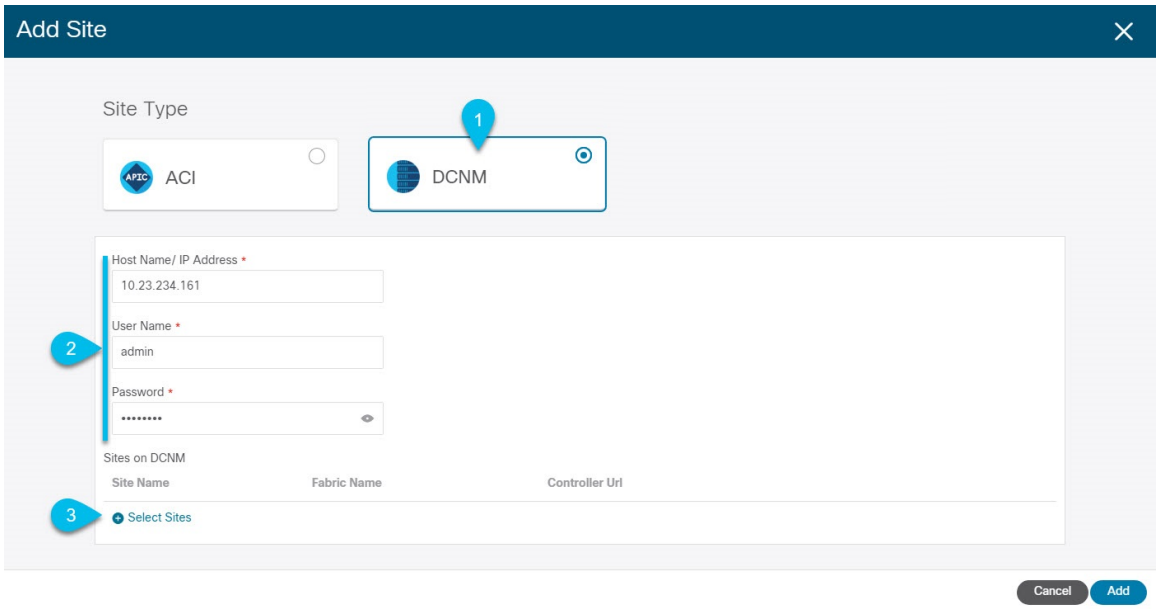
Step 2 Add a new site.



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

b) In the top right of the main pane, select **Actions** > **Add Site**.

Step 3 Provide site information.



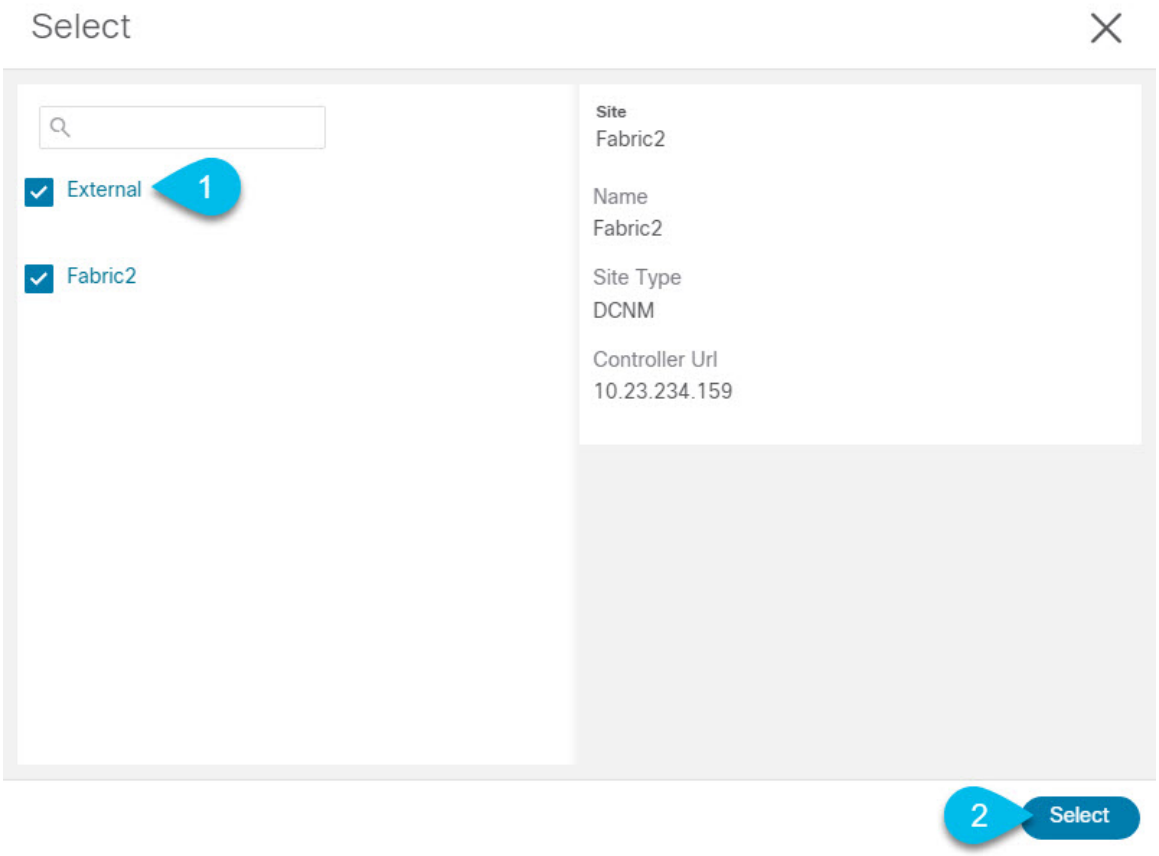
a) For **Site Type**, select **DCNM**.
b) Provide the DCNM controller information.

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

c) Click **Select Sites** to select the specific fabrics managed by the DCNM controller.

The fabric selection window will open.

Step 4 Select the fabrics you want to add to the Nexus Dashboard.



a) Check one or more fabrics that you want to be available to the applications running in your Nexus Dashboard.
b) Click **Select**.

Step 5 In the **Add Site** window, click **Add** to finish adding the sites.

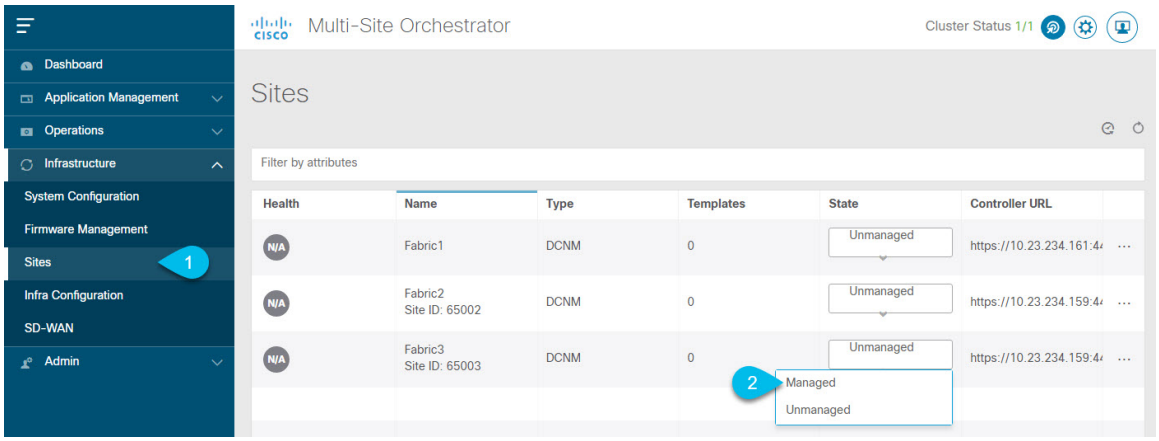
At this time, the sites will be available in the Nexus Dashboard, but you still need to enable them for Nexus Dashboard Orchestrator management as described in the folowing steps.

Step 6 Repeat the previous steps for any additional DCNM controllers.

Step 7 From the Nexus Dashboard's **Service Catalog**, open the Nexus Dashboard Orchestrator service.

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

Step 8 In the Nexus Dashboard Orchestrator GUI, manage the sites.



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 DCNM 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 a DCNM 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.

Procedure

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) In the left sidebar's **Sites** area, select a template associated with the site, click the options menu (**...**) next to the template, and choose **Undeploy Template**.

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

Note
For non-stretched templates, you can choose to preserve the configuration by selecting **Dissociate Template** instead of **Undeploy Template**, but you must undeploy any stretched templates.

d) 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 **Infrastructure** > **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**.

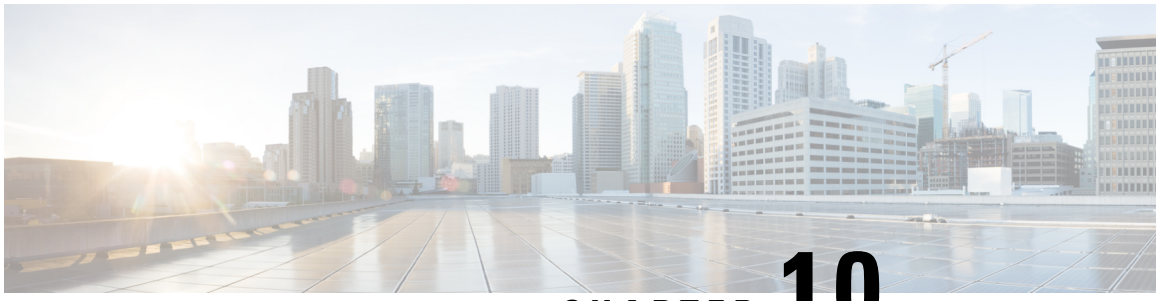
The site will be removed from the Nexus Dashboard.

Cross Launch to Fabric Controllers

Nexus Dashboard Orchestrator currently supports a number of configuration options for each type of fabrics. For many additional configuration options, you may need to log in directly into the fabric's controller.

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

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



CHAPTER 10

Configuring Infra for Cisco DCNM Sites

- Prerequisites and Guidelines, on page 63
- Configuring Infra: General Settings, on page 63
- Refreshing Site Connectivity Information, on page 64
- Configuring Infra: DCNM Site Settings, on page 65
- Deploying Infra Configuration, on page 67

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 64 as part of the general Infra configuration procedures.

Configuring Infra: General Settings

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

Procedure

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

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

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

Step 4 In the left sidebar, select **General Settings**.

Step 5 Configure **Control Plane Configuration**.

a) Select the **Control Plane BGP** tab.

b) Choose **BGP Peering Type**.

- `full-mesh`—All border gateway switches in each site will establish peer connectivity with remote sites' border gateway switches.
- `route-server`—The route-server option allows you to specify one or more control-plane nodes to which each site establishes MP-BGP EVPN sessions. The route-server nodes perform a function similar to traditional BGP route-reflectors, but for EBGP (and not iBGP) sessions. The use of route-server nodes avoids creating MP-BGP EVPN full mesh adjacencies between all the VXLAN EVPN sites managed by NDO.

c) If you set the **BGP Peering Type** to `route-server`, click **+Add Route Server** to add one or more route servers.

In the **Add Route Server** window that opens:

- From the **Site** dropdown, select the site you want to connect to the route server.
- The **ASN** field will be auto-populated with the site's ASN.
- From the **Core Router Device** dropdown, select the route server to which you want to connect.
- From the **Interface** dropdown, select the interface on the core router device.

You can add up to 4 route servers. If you add multiple route servers, every site will establish MP-BGP EVPAN adjacencies to every route server.

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

e) Skip the **OSPF Area ID** and **External Subnet Pool** fields at default values as they are relevant for Cisco Cloud ACI fabrics only.

Step 6 Skip the **IPN Devices** tab settings.

The settings under the **IPN Devices** tab are for Cisco ACI inter-site connectivity between on-premises APIC and Cloud APIC sites. You can skip these settings when managing Cisco DCNM sites only.

Step 7 Configure **DCNM Settings**.

a) Select the **DCNM Settings** tab.

b) Provide the **L2 VXLAN VNI Range**.

c) Provide the **L3 VXLAN VNI Range**.

d) Provide the **Multi-Site Routing Loopback IP Range**.

This field is used to auto-populate the **Multi-Site TEP** field for each fabric, which is described in Configuring Infra: DCNM Site Settings, on page 65.

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

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

Refreshing Site Connectivity Information

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

Procedure

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

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

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

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

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

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

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

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

This will discover any new or removed spines and all site-related fabric connectivity will be re-imported from the site's controller.

Configuring Infra: DCNM Site Settings

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

Procedure

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

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

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

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

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

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

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

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

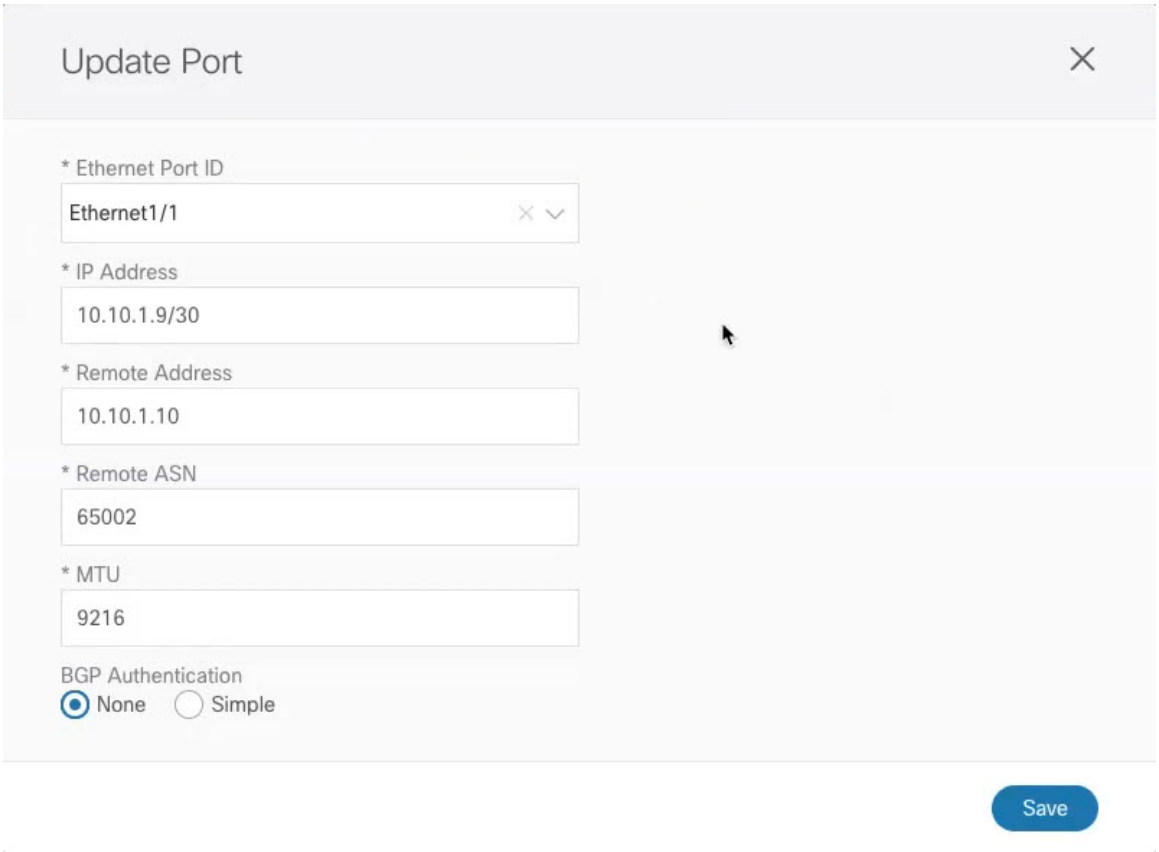
Step 6 Within the **<fabric-name>** tile, select the border gateway.

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

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

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

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



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

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

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

  You can specify either `inherit` or a value between `576` and `9000`.
- For **BGP Authentication**, you can pick either `None` or `Simple` (MD5).

If you select `Simple`, you must also provide the **Authentication Key**.

Deploying Infra Configuration

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

Before you begin

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

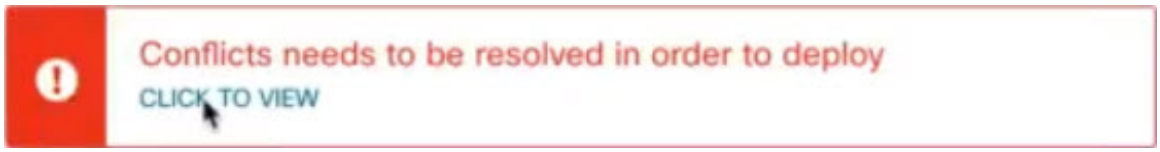
Procedure

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

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

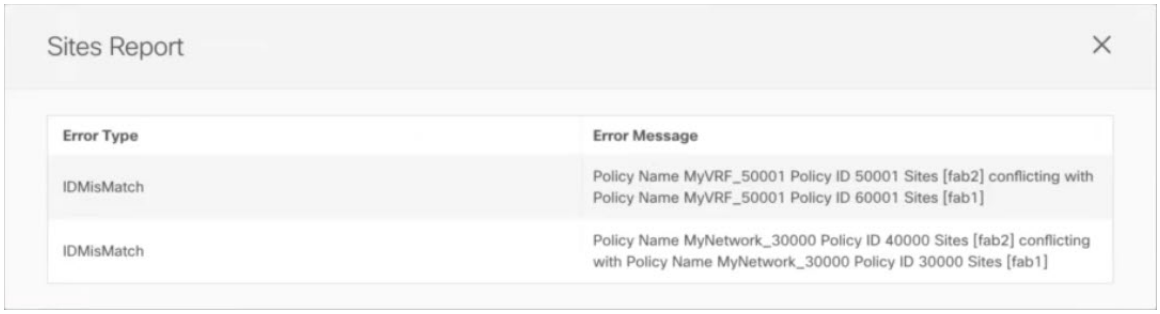
In case of configuration conflicts:

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



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.



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 60.

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

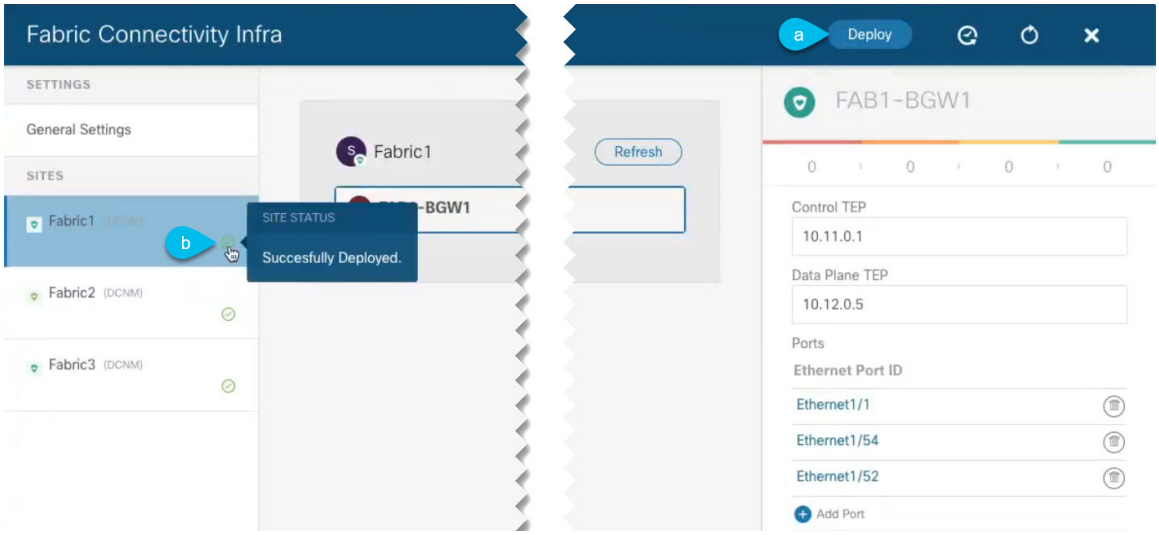
e) Resolve the existing configuration conflicts.

f) Manage the site again, as described in Adding Cisco DCNM Sites, on page 57.

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

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

Step 2 Deploy configuration.



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

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

b) Wait for configuration to be deployed.

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

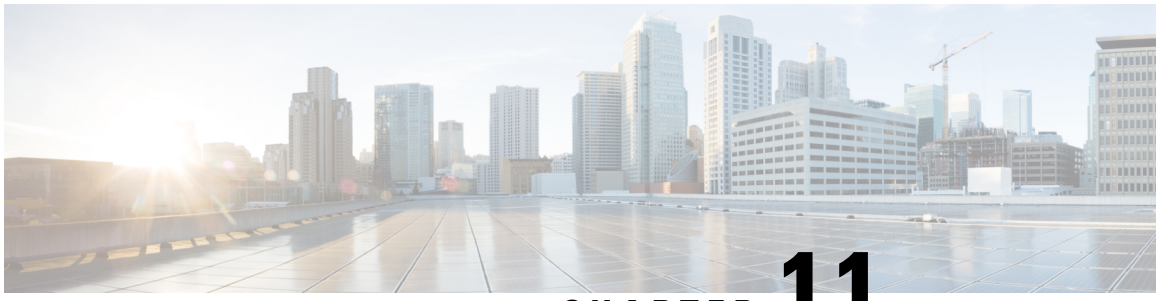
When configuration is successfully deployed, you will see a green checkmark next to the site in the **Fabric Connectivity Infra** screen:



PART III

Upgrading Nexus Dashboard Orchestrator

- Upgrading NDO Service in Nexus Dashboard, on page 71
- Migrating Existing MSO Cluster to Nexus Dashboard, on page 83



CHAPTER 11

Upgrading NDO Service in Nexus Dashboard

- Overview, on page 71
- Prerequisites and Guidelines, on page 71
- Upgrading NDO Service Manually, on page 73
- Restore Existing Configuration for Database Optimization, on page 75
- Resolve Configuration Drifts, on page 76

Overview

The following sections describe how to upgrade Cisco Nexus Dashboard Orchestrator that is deployed in Cisco Nexus Dashboard.

If you are running an earlier release deployed in VMware ESX VMs or Cisco Application Services Engine, you must deploy a brand new cluster and then transfer the configuration from your existing cluster, as described in the "Migrating Existing MSO Cluster to Nexus Dashboard" chapter of the *Cisco Nexus Dashboard Orchestrator Deployment Guide* instead.

Prerequisites and Guidelines

Before you upgrade your Cisco Nexus Dashboard Orchestrator cluster:

- When upgrading an existing Nexus Dashboard Orchestrator release 3.2(1) or later, we recommend upgrading to release 3.7(2).

  At this time, stateful upgrades from release 3.2(1) or later to release 4.x are not supported. You can migrate to a 4.x release as described in *Cisco Nexus Dashboard Orchestrator Deployment Guide, Release 4.0(x)*, however we recommend upgrading to release 3.7(2) instead.

- We recommend that you first familiarize yourself with the Nexus Dashboard platform and overall deployment overview and guidelines described in the *Cisco Nexus Dashboard Deployment Guide* and the *Cisco Nexus Dashboard Orchestrator Deployment Guide* for your release.

Note Ensure that you have followed the Nexus Dashboard deployment prerequisites and guidelines (such as CPU, RAM, and disk requirements) for the cluster where you deploy your Nexus Dashboard Orchestrator. Specifically, if you have a virtual cluster, the CPU and RAM system requirements must be available with physical reservation.

- Stateful upgrades from releases prior to Release 3.2(1) are not supported.

  If you are upgrading from an earlier release, skip the rest of this chapter and follow the instructions described in the "Migrating Existing Cluster to Nexus Dashboard" chapter of the *Cisco Nexus Dashboard Orchestrator Deployment Guide*.

- 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.
- When upgrading to this release, you will manually download the upgrade image and install it as described in Upgrading NDO Service Manually, on page 73.

  You must manually download the upgrade image because the DC App Center includes only the latest release of NDO and stateful upgrades from release 3.2(1) or later to release 4.x are not supported.

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

  On-boarding and managing Cloud APIC sites running earlier releases is not support.

- Ensure that there are no configuration drifts between the Orchestrator's configuration and what is actually deployed in the fabrics before you upgrade.

Note Any templates that have configuration changes that are not yet deployed to the sites may cause the upgrade to fail.

More information on resolving configuration drifts is available in the "Schemas" chapter of the *Nexus Dashboard Orchestrator Configuration Guide* for your current release.

- Back up 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 your existing fabrics' configurations.

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 for AWS 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.

- Once you upgrade to this release, downgrading to an earlier release is not supported.

  If you want to revert to an earlier release, you will need to re-install the NDO service and restore a configuration backup from that release.

Upgrading NDO 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.

Procedure

Step 1 Download the target 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**.

Note
We recommend upgrading to release 3.7(2).

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

Step 2 Log in to your Nexus Dashboard.

Step 3 Upload the image to your Nexus Dashboard.

a) From the left navigation menu, select **Service Catalog**.
b) In the Nexus Dashboard's **Service Catalog** 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`.

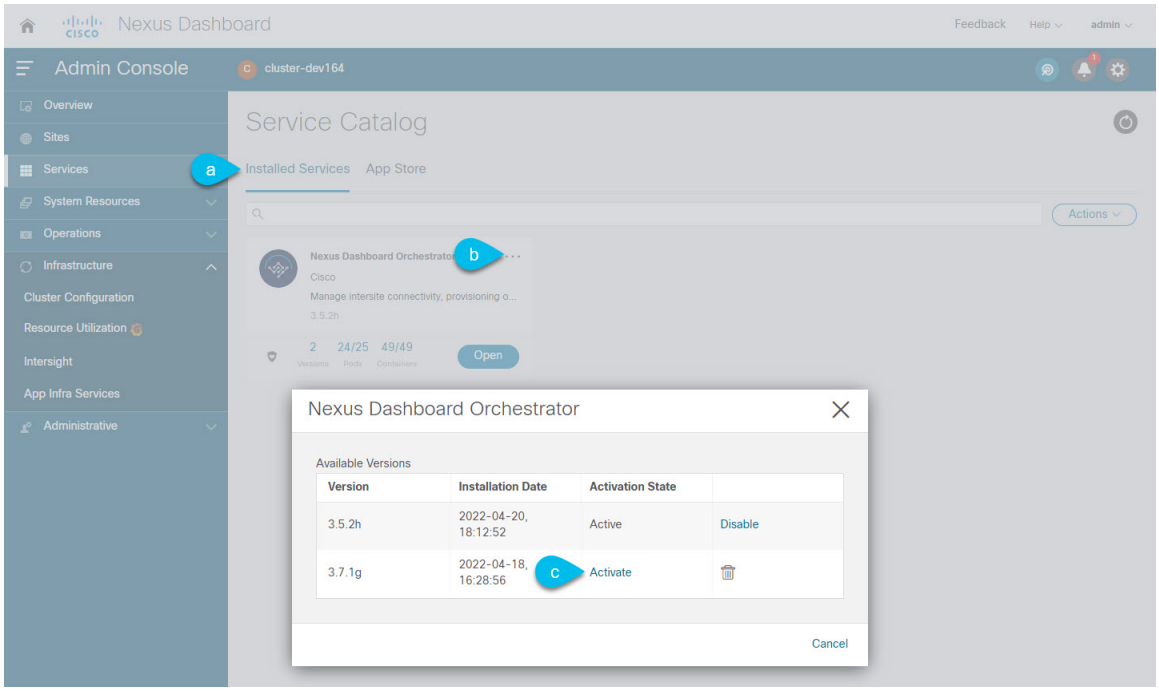
f) Click **Upload** to add the app to the cluster.

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 4 Wait for the new image to initialize.

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

Step 5 Activate the new image.



a) In the **Service Catalog** screen, select the **Installed Services** tab.
b) In the top right 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 6 (Optional) Delete the old application image.

Downgrading from this release is not supported so we recommend delete the old Orchestrator release image as described in this step.

a) In the **Service Catalog** screen, select the **Installed Services** tab.
b) In the top right 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 previous image.

Step 7 Launch the app.

To launch the app, simply click **Open** on the application tile in the Nexus Dashboard's **Service Catalog** page.

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 you have upgraded the NDO service, we recommend you go through the configuration restore workflow to optimize your database, as described in Restore Existing Configuration for Database Optimization, on page 75 and then you must resolve any configuration drifts and redeploy the templates as described in Resolve Configuration Drifts, on page 76.

Restore Existing Configuration for Database Optimization

Release 3.7(2) added database optimization functionality to the configuration restore workflow. After your upgrade is complete, we strongly recommend going through configuration restore process in order to update your existing configuration databases.

Note Skipping this procedure may result in stale values from older configuration changes to remain in the database.

Before you begin

You must have:

- Upgraded your Nexus Dashboard Orchestrator as described in Upgrading NDO Service Manually, on page 73.
- A backup of the existing configuration taken right before the upgrade to Release 3.7(2)

Procedure

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

Step 2 Restore the configuration.

a) In the main window, click the actions (**…**) icon next to the backup you created prior to the upgrade and select **Rollback to this backup**.

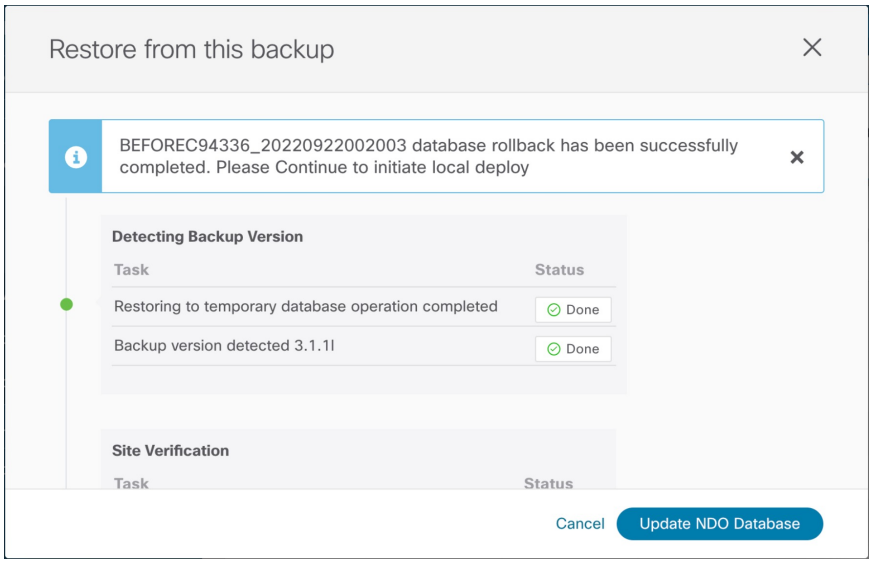
This opens the **Restore from this backup** warning dialog.

b) In the **Restore from this backup** dialog window, click **Restore** to confirm that you want to restore the backup you selected.

The time required for the database rollback and optimization to complete depends on the size of your configuration. Very large configurations may take up to an hour to finish.

c) After the database is restored, click **Update NDO Database** to complete database optimization.

Release 3.7(2) added database optimization functionality to the configuration restore workflow. So you will get an additional prompt for database optimization workflow.



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 version's configuration 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 underlay and overlay connectivity is still successfully established.

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 APIC 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 APIC 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 a yellow warning sign next to the template name in the schema view as shown in the following figure.

When migrating to NDO release 3.7(2) or later, enhancements have been introduced in the configuration rollback procedure to ensure that the content of the NDO database can be fully rebuilt based on the configuration information present in the backup file. This means that if some of the templates in your existing configuration were not fully deployed when the backup file was originally created (for example, left in the "edit" state), the NDO configuration for those templates would be based on that state and may differ from the configuration actually deployed on the fabrics' controllers resulting in a configuration drift.

Before you begin

You must have upgraded your Nexus Dashboard Orchestrator as described in Upgrading NDO Service Manually, on page 73.

Procedure

Step 1 Check for configuration drifts using the API.

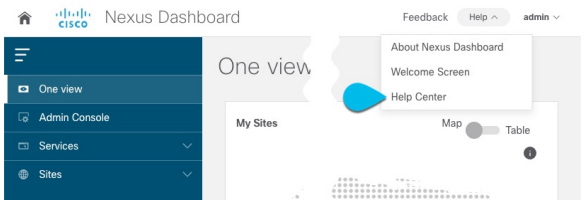
Beginning with release 3.7(2), you can generate a list of all templates that contain configuration drifts by using the `/api/v1/schemas/template-modified-policy-states` API call directly from your Nexus Dashboard Orchestrator's GUI as described in this step.

Alternatively, you can manually check every schema and template individually as described in the next step.

a) Ensure that you are logged in to you Orchestrator UI.

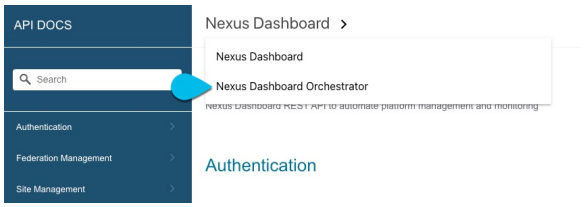
The API uses the authentication token from the Orchestrator UI login.

b) From the **Help** menu in the top right corner of the window, choose **Help Center**.

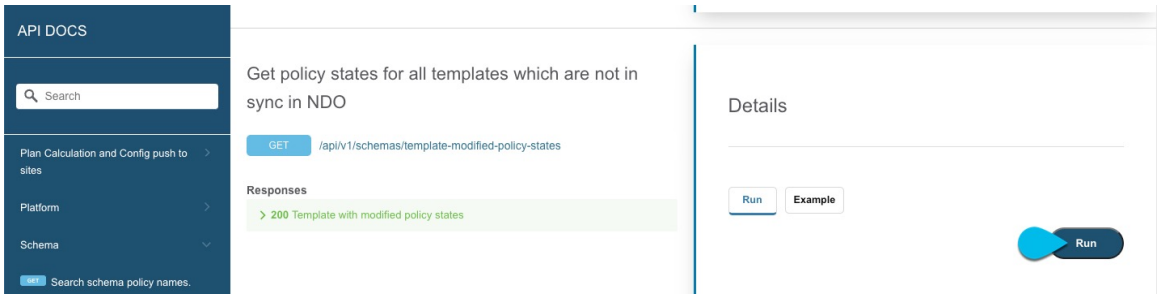


c) In the **Help Center**'s **Programming** tile, click **REST API**.

d) From the dropdown at the top of the page, select **Nexus Dashboard Orchestrator** to show NDO APIs.

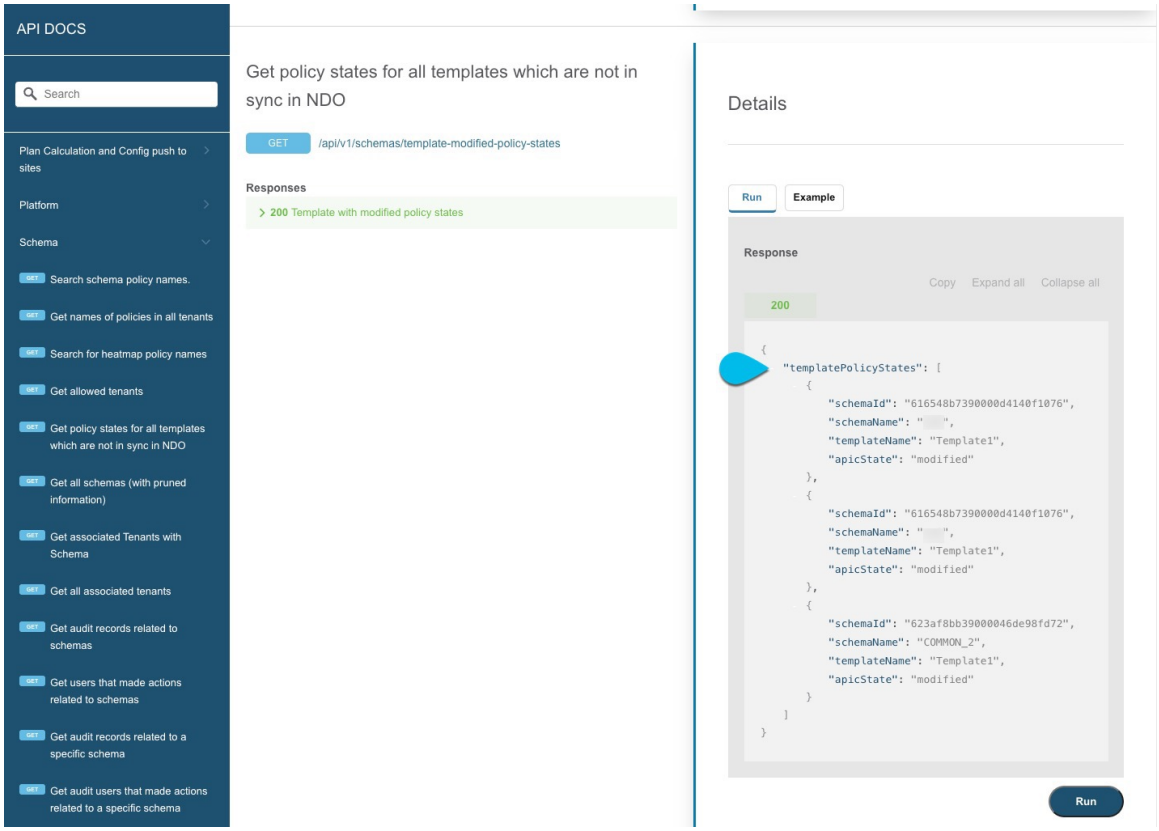


e) Scroll down to the `/api/v1/schemas/template-modified-policy-states` API and click **Run**.



Depending on the number of templates and the size of the configuration, this may take a few minutes, and the **Run** button will be grayed out during this process.

f) Note down all the templates returned by the API call.



Step 2 Check for configuration drifts using the GUI.

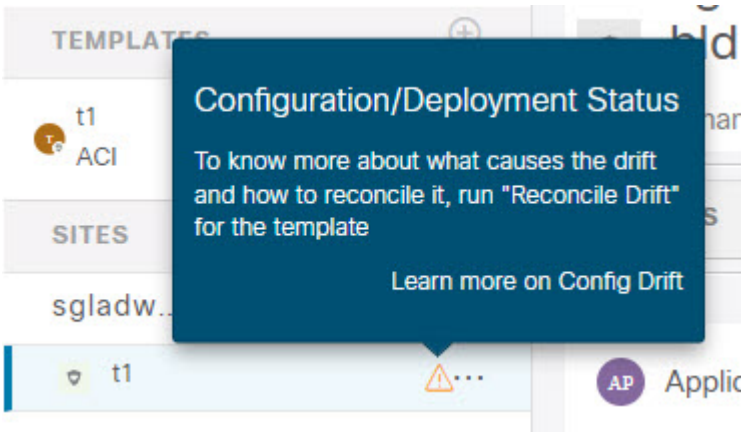
a) In your Nexus Dashboard Orchestrator, navigate to **Application Management** > **Schemas**.

b) Select the first schema and check its templates for configuration drifts.

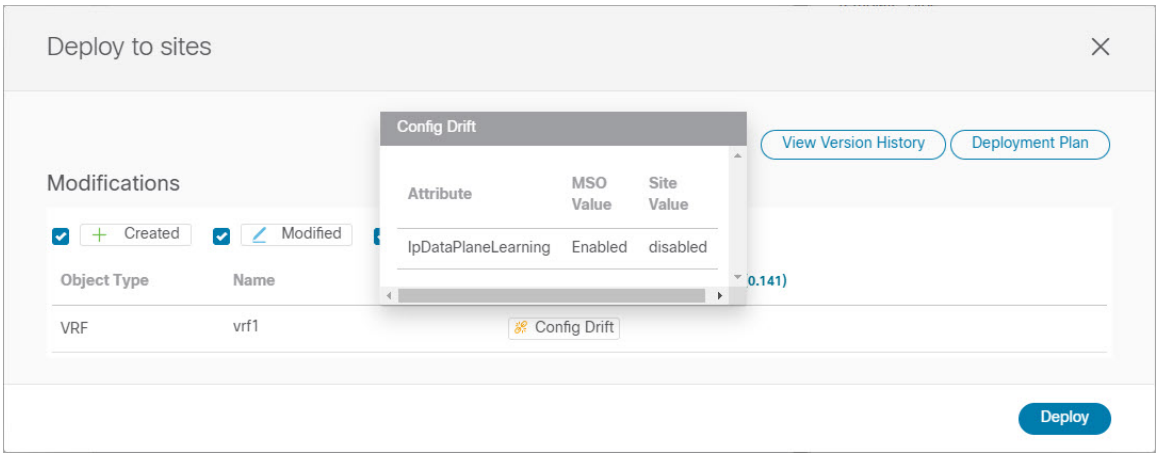
You will repeat the following steps for every schema and template in your deployment

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:



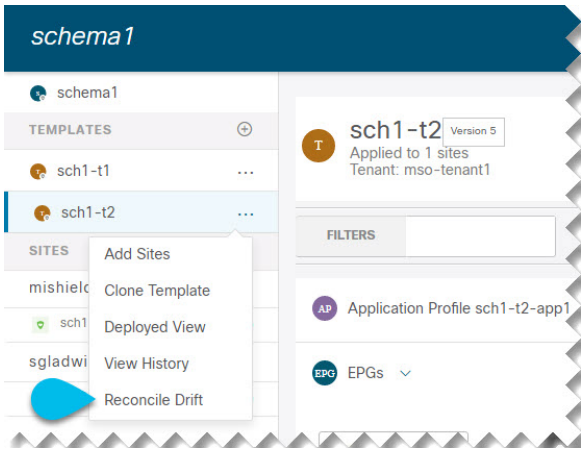
Step 3 For eveyr template that contains a configuration drift, resolve the conflicts.

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.

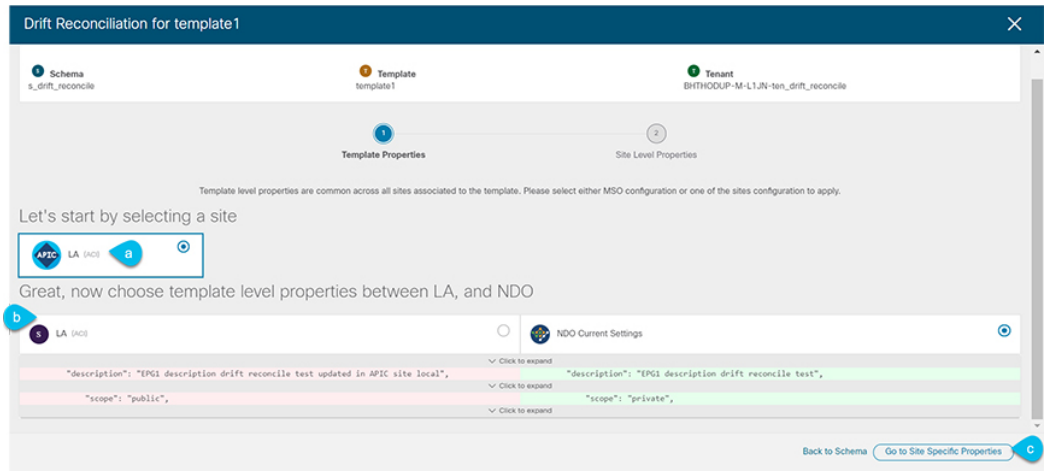
Deploying any templates at this point would push the values in the Orchestrator database and overwrite any existing settings in the fabrics.

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



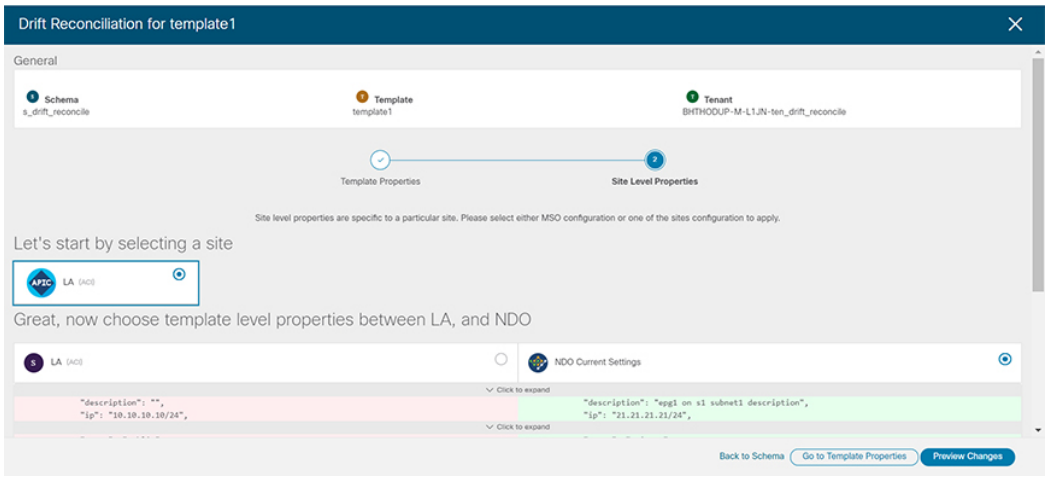
The **Drift Reconciliation** wizard opens.

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.

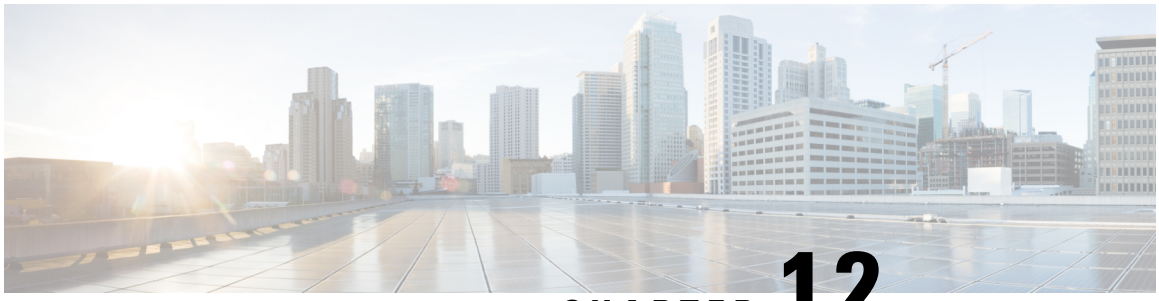


You can choose a site to compare that specific site's configuration. Unlike template-level configurations, you can choose either the Nexus Dashboard Orchestrator-defined or actual existing configurations for each site individually to be retained as the template's site-local properties for that site.

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. You can then click **Deploy to sites** to deploy the configuration and reconcile the drift for that template.



CHAPTER 12

Migrating Existing MSO Cluster to Nexus Dashboard

- Overview, on page 83
- Prerequisites and Guidelines, on page 84
- Back Up Existing Cluster Configuration, on page 86
- Prepare New Cluster, on page 87
- Restore Configuration in the New Cluster, on page 90
- Upgrade Cloud Sites, on page 94
- Update NDO Infra Configuration for Cloud Sites, on page 96
- Resolve Configuration Drifts, on page 98

Overview

This release of Nexus Dashboard Orchestrator (previously known as Multi-Site Orchestrator) must be deployed as a service in Cisco Nexus Dashboard. The previously supported VMware ESX virtual appliance and Cisco Application Services Engine form factors are now deprecated.

The following sections describe how to migrate an earlier release of Cisco Multi-Site Orchestrator to Nexus Dashboard Orchestrator on Nexus Dashboard platform.

If your NDO cluster is already deployed in Nexus Dashboard, follow the steps described in Upgrading NDO Service in Nexus Dashboard, on page 71 instead.

Migration Workflow

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

1. Back up existing Multi-Site Orchestrator configuration and disconnect the existing Multi-Site Orchestrator cluster.

   If you deploy a brand new Nexus Dashboard cluster rather than upgrade an existing cluster, we recommend preserving the existing Multi-Site Orchestrator cluster until the new Nexus Dashboard Orchestrator service is deployed and configuration is restored.

2. Deploy a Nexus Dashboard cluster using physical, virtual, or cloud form factor.

   During new cluster deployment, you will also complete the following:

a. (Optional) Configure the Nexus Dashboard cluster with additional nodes if required for service co-hosting.

b. (Optional) Configure remote authentication servers in the Nexus Dashboard if required by your existing Multi-Site Orchestrator deployment.

c. On-board the APIC, Cloud APIC, or DCNM sites that you currently manage from the Multi-Site Orchestrator to the Nexus Dashboard.

Note When on-boarding the fabrics in the new cluster, you must use the same exact name for each fabric as in the original cluster.

d. Install the Nexus Dashboard Orchestrator service in the Nexus Dashboard.

3. Restore the configuration backup in the new NDO service installed in the Nexus Dashboard.
4. Upgrade cloud sites to Cloud APIC release 5.2(x) one site at a time.

   You will upgrade a site's Cloud APIC, then that site's CSRs, then repeat the procedure for each additional site.
5. Update Infra configuration settings in Nexus Dashboard Orchestrator.
6. Resolve any configuration drifts and redeploy those templates.

   Resolving configuration drifts may require importing objects from the on-boarded fabrics or deploying the configuration from the Orchestrator.

Prerequisites and Guidelines

Because the new platform is vastly different in how it implements clustering and infrastructure, site management, and user management, the migration process involves parallel deployment of a new Nexus Dashboard platform and manual transfer of the current configuration database from your existing Multi-Site Orchestrator (MSO) cluster.

Before you migrate your existing cluster to Nexus Dashboard:

- When upgrading an existing Nexus Dashboard Orchestrator release 3.2(1) or later, we recommend upgrading to release 3.7(2).
- We recommend that you first familiarize yourself with the Nexus Dashboard platform and overall deployment overview and guidelines described in the *Cisco Nexus Dashboard Deployment Guide* and the Deploying Nexus Dashboard Orchestrator, on page 3 chapter of this document.

Note Ensure that you have followed the Nexus Dashboard deployment prerequisites and guidelines (such as CPU, RAM, and disk requirements) for the cluster where you deploy your Nexus Dashboard Orchestrator. Specifically, if you have a virtual cluster, the CPU and RAM system requirements must be available with physical reservation.

- Ensure that your current Multi-Site Orchestrator cluster is healthy.

  You will create a backup of your existing configuration and then import it into the newly deployed NDO service in Nexus Dashboard.

  Ensure that the cluster is healthy and existing IPsec intersite connectivity between cloud and on-premises sites is up.

- Ensure that your on-premises sites are running Cisco APIC release 4.2(4) or later.

  Site management has moved from the Multi-Site Orchestrator UI to the Nexus Dashboard common site management, which supports releases 4.2(4) or later. Fabric upgrades are described in detail in *Cisco APIC Installation, Upgrade, and Downgrade Guide*

- Ensure that your cloud sites are running Cisco Cloud APIC release 5.1(1).

  Site management has moved from the Multi-Site Orchestrator UI to the Nexus Dashboard common site management, which supports on-boarding cloud site releases 5.1(1) or later. Fabric upgrades are described in detail in *Cisco APIC Installation, Upgrade, and Downgrade Guide*



Note However, you must not upgrade to Cloud APIC 5.2(1) release or later before Nexus Dashboard Orchestrator is migrated to this release. If your cloud sites are running Cloud APIC 4.x or 5.0(x) releases, you must upgrade to a Cloud APIC 5.1(x) release before following the instructions in this chapter.

- Ensure that there are no configuration drifts between the Orchestrator's configuration and what is actually deployed in the fabrics before you upgrade.



Note Any templates that have configuration changes that are not yet deployed to the sites may cause the upgrade to fail.

  More information on resolving configuration drifts is available in the "Schemas" chapter of the *Nexus Dashboard Orchestrator Configuration Guide* for your current release.

- Back up 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 your existing fabrics' configurations.

  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 for AWS 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.

- Once you upgrade to this release, downgrading to an earlier release is not supported.

  If you want to revert to an earlier release, you will need to re-install the NDO service and restore a configuration backup from that release.

Back Up Existing Cluster Configuration

The migration process includes creating a backup of current configuration from your existing Multi-Site Orchestrator cluster and then restoring that in the new Nexus Dashboard Orchestrator service running in Nexus Dashboard.

This section describes how to back up your existing cluster configuration.

Note Ensure that there are no configuration drifts between the Orchestrator's configuration and what is actually deployed in the fabrics before you upgrade. This includes any templates that are in edit mode and contain changes that have not been deployed to the fabrics yet. More information on resolving configuration drifts is available in the "Schemas" chapter of the *Nexus Dashboard Orchestrator Configuration Guide* for your current release.

Before you begin

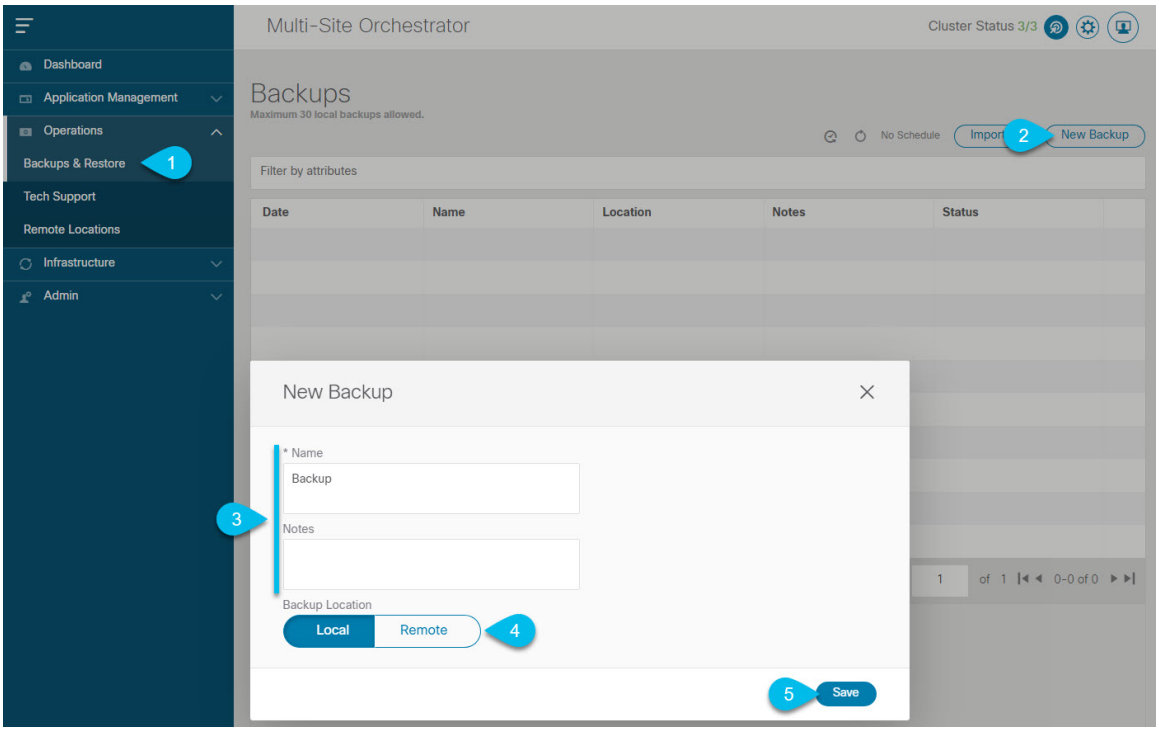
You must have the following completed:

- Familiarized yourself with the migration workflow order described in the Overview, on page 83
- Reviewed and completed general prerequisites described in Prerequisites and Guidelines, on page 84.

Procedure

Step 1 Log in to your existing Multi-Site Orchestrator.

Step 2 Backup 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) Choose `Local` for the **Backup Location**.
e) Click **Save** to create the backup.

Step 3 Download the backup file from the existing Orchestrator.

If you created the backup using a remote location, you can skip this step.

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

Prepare New Cluster

This section describes how to prepare a Nexus Dashboard cluster for installing the Nexus Dashboard Orchestrator service.

It includes choosing and deploying an appropriate form factor of Nexus Dashboard cluster and establishing network connectivity from the cluster to each site you plan to manage from the Nexus Dashboard Orchestrator.

Before you begin

You must have the following completed:

- Familiarized yourself with the migration workflow order described in the Overview, on page 83
- Reviewed and completed general prerequisites described in Prerequisites and Guidelines, on page 84.
- Existing configuration backed up as described in Back Up Existing Cluster Configuration, on page 86.

Procedure

Step 1 Deploy a Nexus Dashboard release 2.1.1e or later cluster and configure fabric connectivity.

How you deploy or upgrade to Nexus Dashboard depends on the deployment type of your existing cluster:

- If your existing Multi-Site Orchestrator is deployed directly in VMware ESX or in a **virtual** Cisco Application Services Engine cluster, you must deploy a brand new virtual or cloud Nexus Dashboard cluster as described in the *Cisco Nexus Dashboard Deployment Guide*.

  We also recommend completing the entire migration process before deleting the existing cluster.

- If you have an existing **physical** Cisco Application Services Engine cluster with Multi-Site Orchestrator service release 3.1(x), you must uninstall the existing service, then upgrade the cluster to Nexus Dashboard release 2.1.1e or later as described in the "Upgrading" chapter of the *Cisco Nexus Dashboard Deployment Guide*.
- If you have an existing **physical** Nexus Dashboard cluster with Nexus Dashboard Orchestrator service release 3.2(x), you can upgrade the cluster as described in the "Upgrading" chapter of the *Cisco Nexus Dashboard Deployment Guide* and then upgrade the Nexus Dashboard Orchestrator service as described in Upgrading Nexus Dashboard Orchestrator, on page 69 and skip the rest of this chapter.

  Note
  If you plan to add any Cloud APIC sites after the upgrade, ensure that they are running Cloud APIC release 5.2(1) or later.

Step 2 Ensure that your Nexus Dashboard cluster is appropriately scaled based on the fabric sizes and number of applications.

If you deployed a virtual or cloud form factor of the Nexus Dashboard, Nexus Dashboard Orchestrator is the only application supported and the base 3-node cluster is sufficient, so you can skip this step.

If you deployed a physical Nexus Dashboard cluster and Nexus Dashboard Orchestrator is the only application you plan to host, the base 3-node cluster is sufficient and you can skip this step.

However, if you deployed a physical Nexus Dashboard cluster and plan to co-host multiple applications, use the Cisco Nexus Dashboard Capacity Planning tool to determine the required cluster size for your specific use case. If you need to extend your cluster to support all required services, see the *Cisco Nexus Dashboard User Guide* for information on deploying additional worker nodes.

Step 3 Install the NDO service in your Nexus Dashboard.

This process is described in detail in the Deploying Nexus Dashboard Orchestrator, on page 3 chapter.

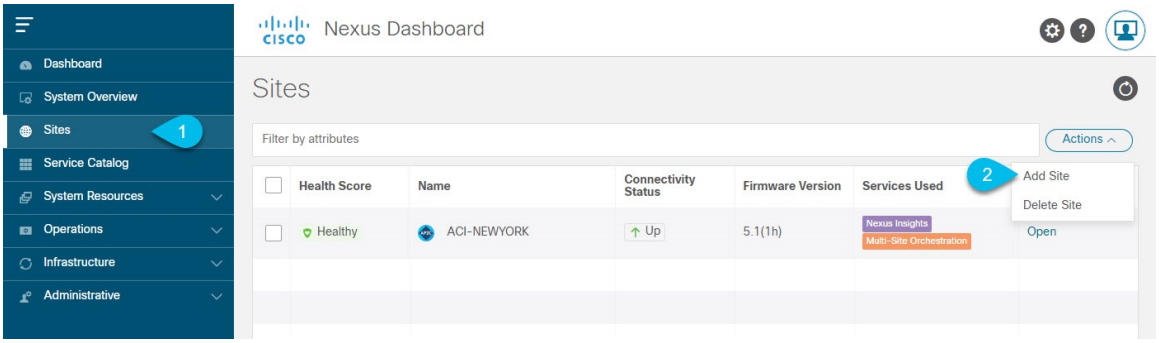
Step 4 On-board all sites to the Nexus Dashboard.

Site management has moved from the Multi-Site Orchestrator UI to the Nexus Dashboard common site management. As such, you must on-board the same sites using the same names that were assigned to the sites when on-boarded on the

original Multi-Site Orchestrator cluster to the Nexus Dashboard GUI before migrating your existing configuration to the new cluster, as described in Adding and Deleting Sites, on page 23. If any site that exists in you current deployment is not present in Nexus Dashboard (or it exists with a different name), the configuration restore during migration will fail with a `Pre-restore check failed` error message.

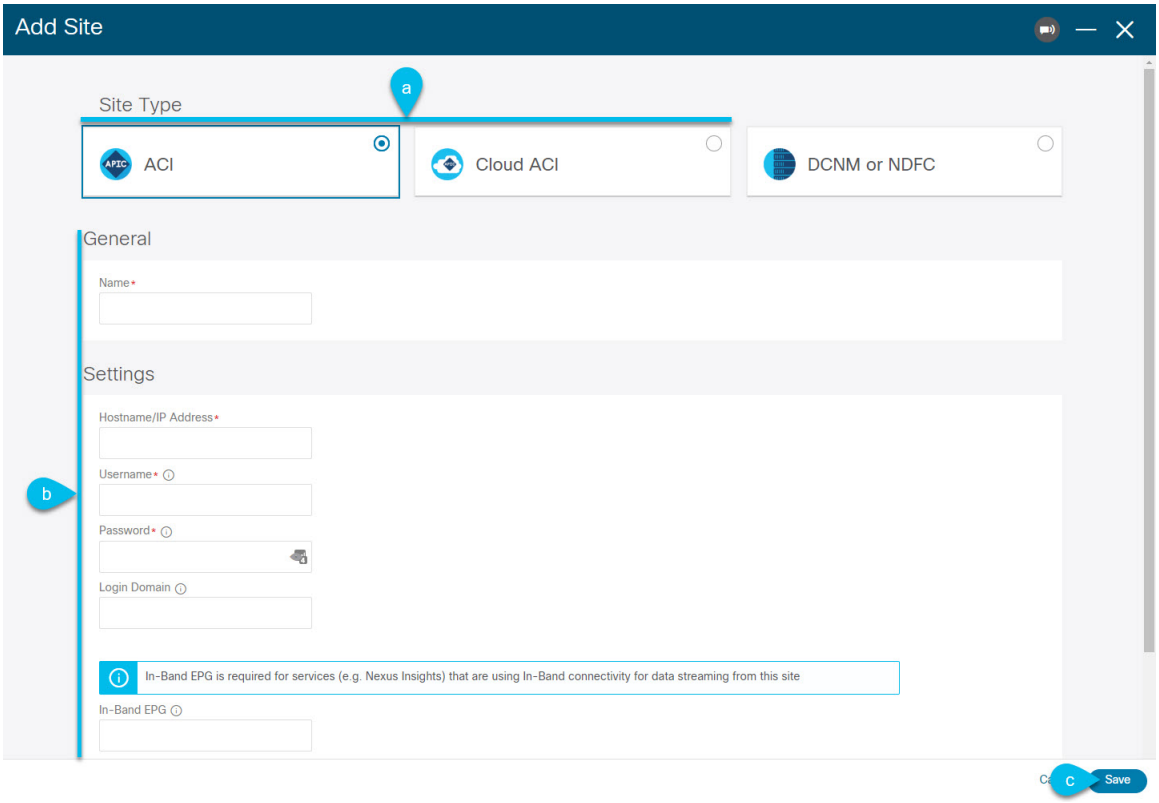
Note
After you add the sites to the Nexus Dashboard, you must not set them to `Managed` in the NDO service. The sites will be enabled for management automatically when you restore your configuration from backup.

Add a site:



a) From the left navigation menu, select **Sites**.
b) In the top right of the main pane, select **Actions** > **Add Site**.

If adding an ACI site, provide the following information:



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

b) Provide the controller information.

You need to provide the **Host Name/IP Address**, **User Name**, and **Password.** for the APIC controller currently managing your ACI fabrics. If NDO is the only application you plan to host, you can specify either the in-band or out-of-band address of the on-premises APIC; however, if you plan to host other applications, such as Nexus Insights, you must specify the in-band address.

Note
By default, the in-band or out-of-band address of the on-premises APIC that you use to on-board the fabric must be reachable from the Nexus Dashboard's data interface.

If you want to use the Nexus Dashboard's management interface for NDO traffic, you must configure a static route from the Nexus Dashboard cluster to the fabric's IP from the management interface. For more information, see the **Infrastructure Management** > **Cluster Configuration** chapter of the *Nexus Dashboard User Guide*.

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

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

At this time, the sites will be available in the Nexus Dashboard, but you still need to enable them for Nexus Dashboard Orchestrator management as described in the following steps.

d) Repeat this step to add all the sites from your existing Multi-Site deployment.

Step 5 Add any remote authentication servers you had configured in your Multi-Site Orchestrator to the Nexus Dashboard.

User management has moved from the Multi-Site Orchestrator UI to the Nexus Dashboard common user management. As such, you must add the same remote users and authentication servers to the Nexus Dashboard, as described in the *Cisco Nexus Dashboard User Guide*.

Any local users you had previously configured directly in Multi-Site Orchestrator will be added into the Nexus Dashboard automatically when you import the existing configuration backup.

Step 6 Add any proxy configuration you had configured in your Multi-Site Orchestrator to the Nexus Dashboard.

Proxy configuration has moved from the Multi-Site Orchestrator UI to the Nexus Dashboard common cluster configuration. As such, you must add the proxy server to the Nexus Dashboard, as described in the *Cisco Nexus Dashboard User Guide*.

Any existing proxy configuration will not be migrated automatically and you must manually re-add it in Nexus Dashboard after the migration.

Restore Configuration in the New Cluster

This section describes how 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:

- Existing configuration backed up as described in Back Up Existing Cluster Configuration, on page 86.

- Deployed Nexus Dashboard cluster and installed Nexus Dashboard Orchestrator service as described in Prepare New Cluster, on page 87.

Procedure

Step 1 Disconnect the existing Multi-Site Orchestrator cluster.

You must disconnect the existing Multi-Site Orchestrator cluster so it does not communicate with the sites during migration.

Note
We recommend preserving the existing Multi-Site Orchestrator cluster until the new cluster is deployed and configuration is restored.

Step 2 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 3 Log in to your Nexus Dashboard GUI.

Step 4 Ensure that all the sites are on-boarded to Nexus Dashboard.

When you restore the backup, NDO will validate that every site in the backup is present in the Nexus Dashboard with matching site name and type. If validation is unsuccessful, for example if a site is not on-boarded in Nexus Dashboard, configuration restore will fail and you will need to on-board the site before retrying, as described in the previous section.

Step 5 Open your new Nexus Dashboard Orchestrator service.

Step 6 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 **Operations** > **Remote Locations**.

b) In the top right of the main window, click **Add Remote Location**.

An **Add New Remote Location** screen appears.

c) 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

d) 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.

e) 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.

f) Specify the port used to connect to the remote server.

By default, port is set to `22`.

g) 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.

h) Click **Save** to add the remote server.

Step 7 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 from file** window that opens, click **Select File** and choose the backup file you want to import.

This is the backup of your existing MSO configuration that you created and downloaded in previous section.

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 **Upload** to import 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 8 Restore the configuration.

a) In the main window, click the actions (**…**) icon next to the backup you created prior to the upgrade and select **Rollback to this backup**.

This opens the **Restore from this backup** warning dialog.

b) In the **Restore from this backup** dialog window, click **Restore** to confirm that you want to restore the backup you selected.

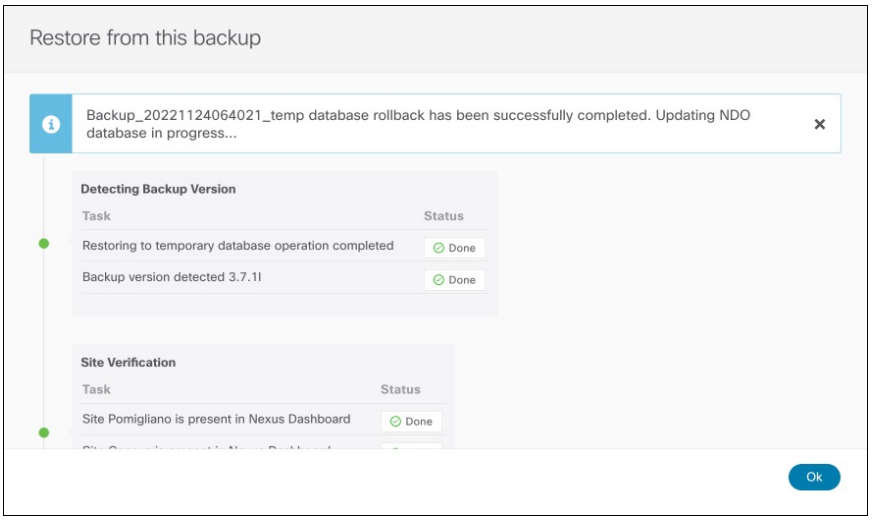
When the configuration is restored, any sites previously managed by Multi-Site Orchestrator and on-boarded to the Nexus Dashboard will be enabled for NDO management in the GUI. If the configuration backup contains sites that are not on-boarded to your Nexus Dashboard, backup restore will fail with a `Pre-restore check failed` error and you will need to repeat the procedure after on-boarding any missing sites.

Depending on the size of your configuration, the database rollback may take several minutes to complete.

c) After the database is restored, click **Ok** to proceed.

Release 3.7(2) added database optimization to the configuration rollback workflow, which is automatically triggered as the final stage of restoring configuration.

Simply click **Ok** to view the database update progress:

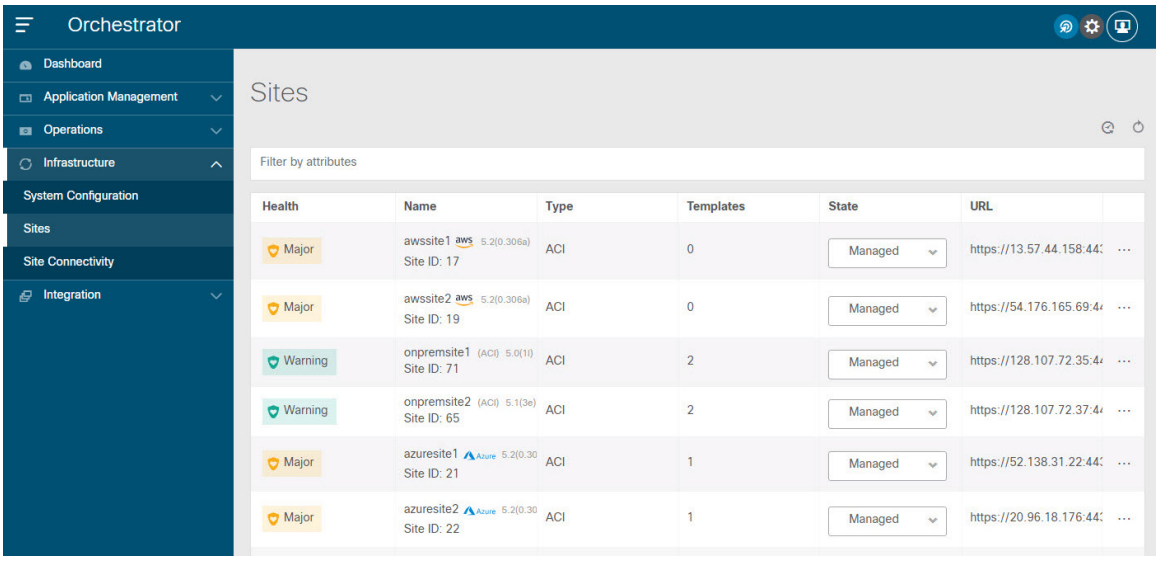


Step 9 Update the password.

Due to CSDL (Cisco Secure Development Lifecycle) requirements, you may be required to update the `admin` user password after configuration restore is completed.

Step 10 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 Multi-Site 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 underlay and overlay connectivity is still successfully established.

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 APIC 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 APIC in earlier Nexus Dashboard Orchestrator releases.

Note
You must not make any changes or deploy any configurations at this stage until the cloud sites are upgraded to Cloud APIC release 5.2(1) as described in following sections.

Upgrade Cloud Sites

After Nexus Dashboard Orchestrator is migrated to this release, you must upgrade any Cloud APIC sites managed by the NDO to release 5.2(1) or later. While existing intersite connectivity will remain intact, you will not be able to change or deploy any cloud site Infra configurations to sites running Cloud APIC releases prior to release 5.2(1).

Before you begin

You must have the following completed:

- Deployed Nexus Dashboard cluster and installed Nexus Dashboard Orchestrator service as described in Prepare New Cluster, on page 87.
- Existing configuration backup restored to the new cluster as described in Restore Configuration in the New Cluster, on page 90.

Procedure

Step 1 Upgrade cloud sites.

For each cloud site, you must upgrade its Cloud APIC and then its CSRs before proceed to upgrading the next site. After a site is upgraded and healthy, you can repeat the same steps to upgrade any additional sites.

a) Upgrade a site's Cloud APIC.

You can upgrade Cloud APIC as you typically would.

Note that after the Cloud APIC upgrade, any existing public IP tunnels will remain intact and intersite connectivity via public IPsec will not be interrupted .

b) Upgrade that site's CSR.

Starting with Cloud APIC release 5.2(1) , CSRs upgrade does not happen automatically as it used to in earlier releases, so you must manually trigger CSR upgrade after Cloud APIC is upgraded. You must upgrade the site's CSRs before moving on to upgrading the next site.

As you upgrade CSRs in each site, the following will occur:

- As each CSR is upgraded, its existing `/30` tunnels will be recreated and the traffic will continue to flow.
- Tunnel-management and all Infra configuration changes from Nexus Dashboard Orchestrator are disabled for as long as any of the cloud sites are still running any Cloud APIC or CSR releases prior to 5.2(1).
- If the last site you upgrade is an AWS cloud site, the following will occur for that site's CSRs only:
  - The last cloud site's tunnel endpoints will be deleted by Cloud APIC and NDO will delete the corresponding tunnels that use the endpoint
  - NDO will delete the tunnels originating from CSRs in the last cloud site
  - New `hcloudInterCloudSiteTunnel` MO will be created and Nexus Dashboard Orchestrator's tunnel management will allocate `/31` addresses for the new tunnels
  - The CSRs in this site and the CSRs in another cloud site peering with it will establish `/31` tunnels.

  If the last upgraded site is an Azure site, the same `/30` tunnel will be created on the CSRs and the above four bullet points are not relevant.

For any CSRs you add or any underlay configuration changes to existing CSRs after the migration process is completed, all new tunnels created by NDO will be `/31` tunnel.

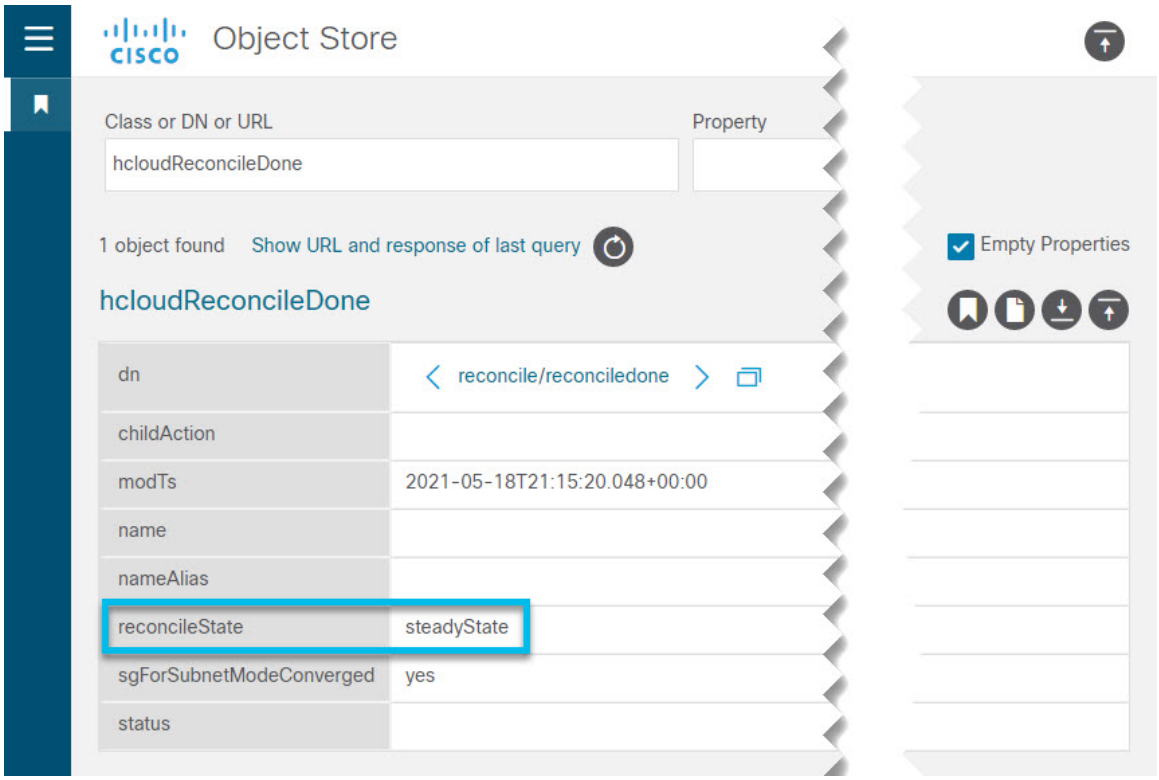
Note
If you do not see BGP sessions within 5 minutes of CSRs upgrade finishing and CSRs coming up, refresh the site's infra connectivity in the Nexus Dashboard Orchestrator **Infra Configuration** screen.

c) Repeat this step for each cloud site one at a time.

Step 2 Verify Cloud APIC and CSR upgrades have completed.

a) In each site's Cloud APIC, check that the `hcloudReconcileDone` MO shows `reconcileState=steadyState`.

You can check the MO by navigating to `https://<cloud-apic-ip>/visore.html` and searching for `hcloudReconcileDone` in the **Class or DN or URL** field.



b) In Nexus Dashboard Orchestrator, 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 connectivity is healthy in the `Overlay Status` and `Underlay Status` tabs.

c) In Nexus Dashboard Orchestrator's **Site Connectivity** page, click **Configure** and confirm that the External Subnet Pools previously configured in Cloud APIC have been imported and are present.

You can view the external subnet pools by selecting **General Settings** > **IPSec Tunnel Subnet Pools** tab of the **Fabric Connectivity Infra** screen.

d) In Nexus Dashboard Orchestrator's **Fabric Connectivity Infra** screen, select a cloud site, click the **Inter-Site Connectivity** tab in the right-hand sidebar, and confirm that underlay connectivity using public IPs is preserved for existing sites.

Update NDO Infra Configuration for Cloud Sites

In order to make subsequent changes to Infra configuration, you must first provide the following information immediately after the cloud sites are upgraded to Cloud APIC release 5.2(1):

- OSPF area ID
- IPN configuration

Note If you have no cloud sites, you can skip this section.

Before you begin

You must have the following completed:

- Deployed Nexus Dashboard cluster and installed Nexus Dashboard Orchestrator service as described in Prepare New Cluster, on page 87.
- Existing configuration backup restored to the new cluster as described in Restore Configuration in the New Cluster, on page 90.
- Upgraded cloud sites as described in Upgrade Cloud Sites, on page 94.

Procedure

Step 1 Log in to your new Nexus Dashboard Orchestrator.

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

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

Step 4 In the left sidebar, select **General Settings**.

Step 5 Provide the **OSPF Area ID** field.

This is OSPF area ID used by cloud sites for on-premises ISN peering, which you previously configured in the Cloud APIC for inter-site connectivity in earlier Nexus Dashboard Orchestrator releases.

Step 6 Add **IPN Devices** information.

a) Select the **IPN Devices** tab.
b) Click **Add IPN Device**.
c) Provide the **Name** and the **IP Address** of the on-premises IPN devices.

You must provide the IP addresses of the devices in your on-premises sites that are used as the tunnel peer address from the Cloud APIC's CSRs, not the IPN device's management IP address.

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

Step 7 Update **Underlay Configuration** for inter-site connectivity between on-premises and cloud sites.

For each on-premises site that connects to cloud sites, you need to provide at least one IPN device IP address from the ones you added in the previous step, to which the Cloud APIC's CSRs establish a tunnel.

a) In the left pane, under **Sites**, select the on-premises site.
b) In the right **<*Site*> Settings** pane, select the **Underlay Configuration** tab.
c) Click **+Add IPN Device** to specify an IPN device.
d) From the dropdown, select one of the IPN devices you defined previously.

The IPN devices must be already defined in the **General Settings** > **IPN Devices** list, as described in the previous step.

Step 8 From the dropdown at the top of the screen, select **Deploy** to re-deploy the Infra configuration.

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 a yellow warning sign next to the template name in the schema view as shown in the following figure.

After restoring an MSO backup configuration in NDO, some templates may show configuration drifts, which can occur for one of the following reasons:

- Nexus Dashboard Orchestrator added support for managing more objects' properties compared to the previous Multi-Site Orchestrator versions. As a result, the MSO configuration backup would not contain any information about or values for the new properties and NDO will assign default values to them. If you had modified those properties directly in APIC managed by MSO, the NDO templates containing those objects would show a drift.

Note Deploying any templates before resolving these drifts would push the configuration defined on the NDO templates and overwrite the non-default values defined in the fabrics' controllers.

- When migrating to NDO release 3.7(2) or later, enhancements have been introduced in the configuration rollback procedure to ensure that the content of the NDO database can be fully rebuilt based on the configuration information present in the backup file. This means that if some of the MSO templates were not fully deployed when the backup file was originally created (for example, left in the “edit” state), the NDO configuration for those templates would be based on that state and may differ from the configuration actually deployed on the fabrics' controllers resulting in a configuration drift.

Before you begin

You must have the following completed:

- Deployed Nexus Dashboard cluster and installed Nexus Dashboard Orchestrator service as described in Prepare New Cluster, on page 87.
- Existing configuration backup restored to the new cluster as described in Restore Configuration in the New Cluster, on page 90.
- Upgraded cloud sites as described in Upgrade Cloud Sites, on page 94.
- Updated Nexus Dashboard Orchestrator Infra configuration for the cloud sites as described in Update NDO Infra Configuration for Cloud Sites, on page 96.

Procedure

Step 1 Check for configuration drifts using the API.

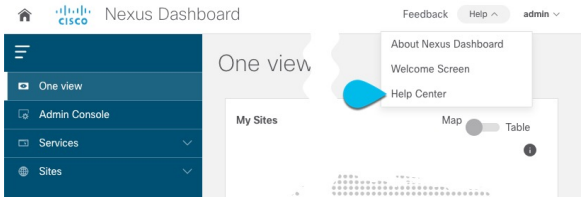
Beginning with release 3.7(2), you can generate a list of all templates that contain configuration drifts by using the `/api/v1/schemas/template-modified-policy-states` API call directly from your Nexus Dashboard Orchestrator's GUI as described in this step.

Alternatively, you can manually check every schema and template individually as described in the next step.

a) Ensure that you are logged in to you Orchestrator UI.

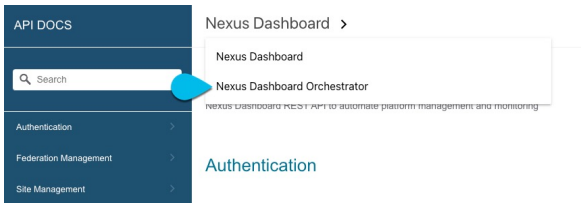
The API uses the authentication token from the Orchestrator UI login.

b) From the **Help** menu in the top right corner of the window, choose **Help Center**.



c) In the **Help Center**'s **Programming** tile, click **REST API**.

d) From the dropdown at the top of the page, select **Nexus Dashboard Orchestrator** to show NDO APIs.

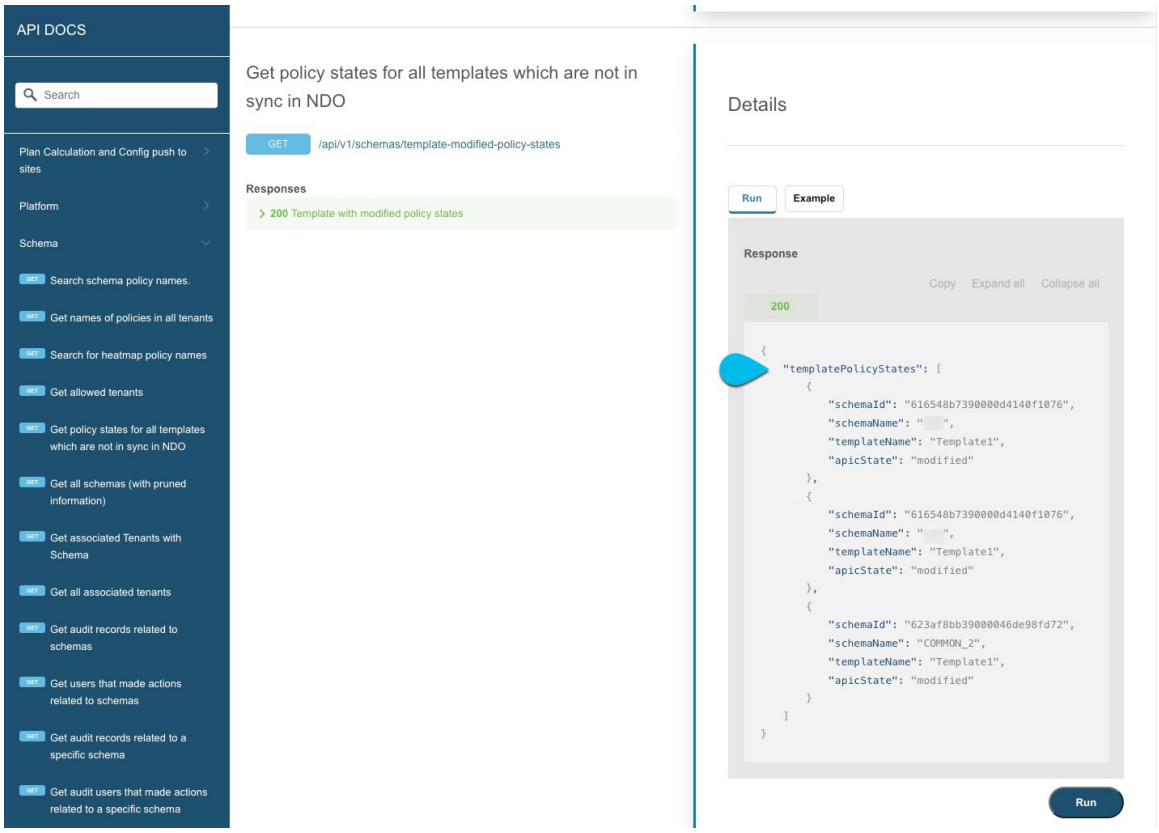


e) Scroll down to the `/api/v1/schemas/template-modified-policy-states` API and click **Run**.



Depending on the number of templates and the size of the configuration, this may take a few minutes, and the **Run** button will be grayed out during this process.

f) Note down all the templates returned by the API call.



Step 2 Check for configuration drifts using the GUI.

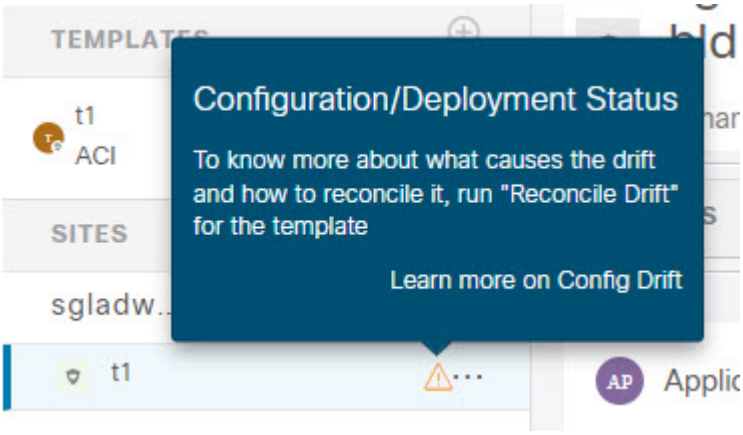
a) In your Nexus Dashboard Orchestrator, navigate to **Application Management** > **Schemas**.

b) Select the first schema and check its templates for configuration drifts.

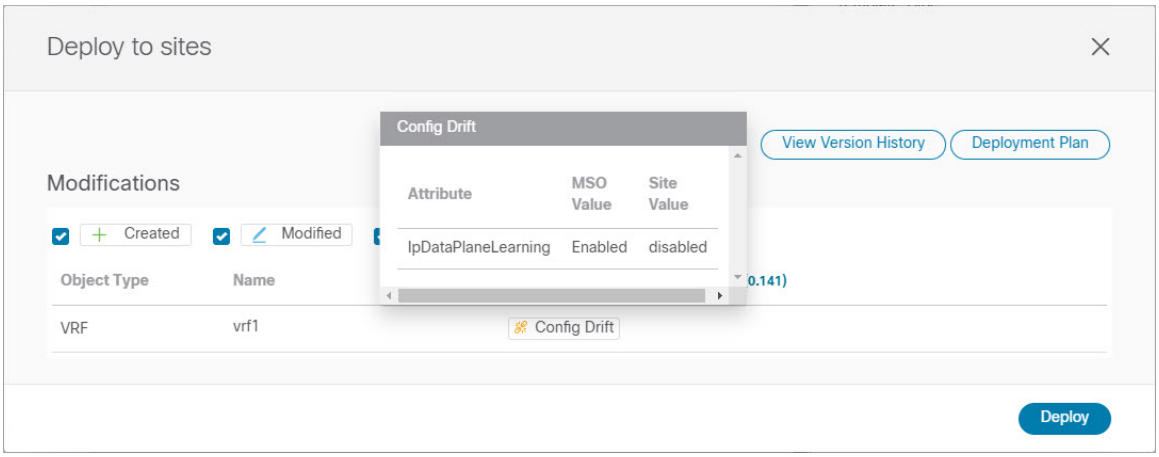
You will repeat the following steps for every schema and template in your deployment

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:



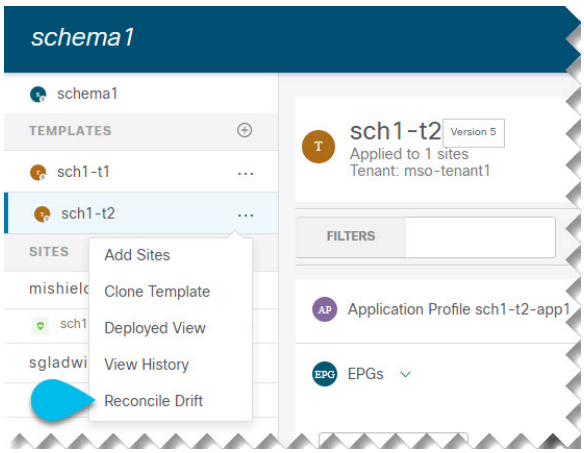
Step 3 For every template that contains a configuration drift, resolve the conflicts.

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.

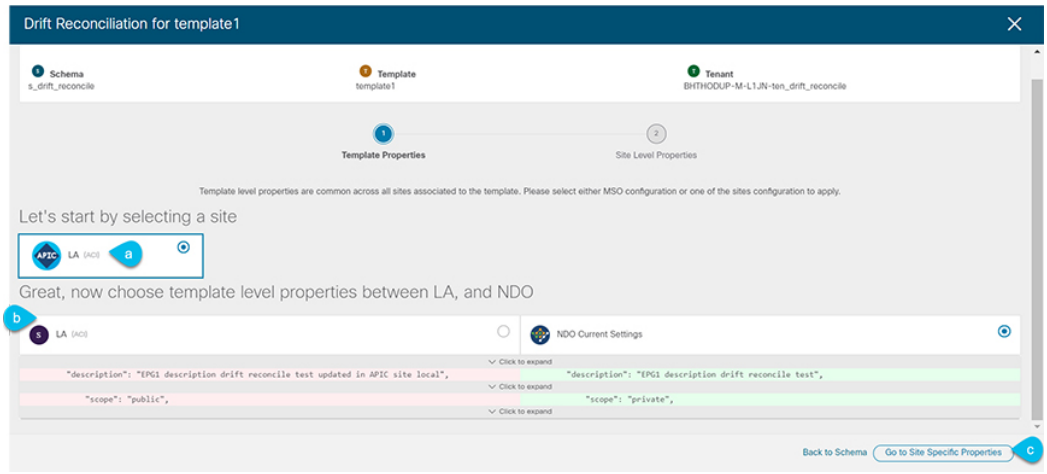
Deploying any templates at this point would push the values in the Orchestrator database and overwrite any existing settings in the fabrics.

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



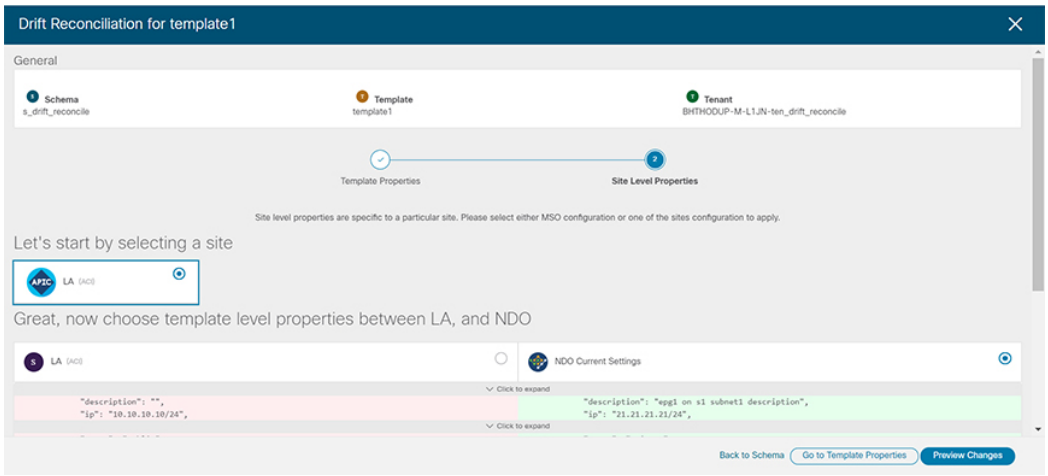
The **Drift Reconciliation** wizard opens.

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. You can then click **Deploy to sites** to deploy the configuration and reconcile the drift for that template.