



Creating LAN and ACI Fabrics and Fabric Groups, Release 4.1.1

Table of Contents

New and changed information	1
Understanding LAN and ACI fabrics and fabric groups	2
Understanding LAN and ACI fabrics	2
Grouping fabrics and clusters	3
Grouping fabrics	3
Grouping clusters	3
Mapping fabric types	3
Guidelines and limitations	4
Viewing LAN fabric information	5
View information on local online fabrics	5
View information on local snapshot fabrics	7
View information on remote fabrics	8
Creating or onboarding local online LAN fabrics	9
Guidelines and limitations: Onboarding LAN fabrics	9
Create or onboard a local online LAN fabric	9
Select a category	9
Select a type	10
Settings	11
Advanced settings	15
Fabric summary	15
Fabric creation	16
Add switches to the fabric	16
Editing fabric settings	17
Onboard ACI fabrics	18
Onboard snapshot LAN fabrics	19
Select a category	19
Basic settings	20
Fabric summary	20
Create fabric groups	21
Settings	21
Advanced settings	22
Fabric group summary	22
Fabric group creation	22
Add child fabrics to the fabric group	22
Delete fabric groups	23
Create multi-cluster fabric groups	24
Guidelines and limitations for creating a multi-cluster fabric group	24
Configure a multi-cluster fabric group	24
Add child fabrics to a multi-cluster fabric group	25
Remove child fabrics from a multi-cluster fabric group	25
Back up and restore multi-cluster fabric group configurations	26

Back up multi-cluster fabric group configurations	26
Restore multi-cluster fabric group configurations	26
Migrate a Nexus Dashboard 3.2.x Orchestration-managed fabric to a multi-cluster fabric group . . .	28
Prerequisites for migrating a Nexus Dashboard 3.2.x Orchestration-managed fabric to a multi-cluster fabric group.	28
Guidelines and limitations for migrating a Nexus Dashboard 3.2.x Orchestration-managed fabric to a multi-cluster fabric group	28
Upgrade from Nexus Dashboard 3.2.x to Nexus Dashboard 4.1.1	28
How to migrate a Nexus Dashboard 3.2.x Orchestration-managed fabric to a multi-cluster fabric group	29
View the migrated multi-cluster fabric group on the Topology page.	29
Additional settings	31
Understanding the Fabric Summary page	31
Prerequisites to creating a fabric	31
Change persistent IP address	31
Configuring overlay mode.	32
Configuring Netflow support.	32
Netflow support for brownfield deployments.	33
VXLAN OAM	33
AI QoS classification and queuing policies	35
Understanding AI QoS classification and queuing policies	35
Guidelines and limitations for AI QoS classification and queuing policies	36
Configure AI QoS classification and queuing policies	36
Create a policy using the custom QoS templates	38
Configuring downstream VNI	38
Benefits of downstream VNI	40
Use cases for downstream VNI	40
Supported platforms	41
Guidelines and limitations for downstream VNI	41

New and changed information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
Nexus Dashboard 4.1.1	Improved navigation and workflow when creating LAN and ACI fabrics and fabric groups	Beginning with Nexus Dashboard 4.1.1, Nexus Dashboard enhanced the navigation and workflow when creating LAN and ACI fabrics and fabric groups.
Nexus Dashboard 4.1.1	Support for migrating from a Nexus Dashboard 3.2.x Orchestration-managed fabric to a Nexus Dashboard 4.1.1 multi-cluster fabric group	With this release, Nexus Dashboard added support for migrating from a Nexus Dashboard 3.2.x Orchestration-managed fabric to a Nexus Dashboard 4.1.1 multi-cluster fabric group after upgrading to Nexus Dashboard 4.1.1. For more information, see Migrate a Nexus Dashboard 3.2.x Orchestration-managed fabric to a multi-cluster fabric group .

Understanding LAN and ACI fabrics and fabric groups

Before you begin creating fabrics or fabric groups, it's helpful to understand more about each.

- [Understanding LAN and ACI fabrics](#)
- [Grouping fabrics and clusters](#)
- [Mapping fabric types](#)
- [Guidelines and limitations](#)

Understanding LAN and ACI fabrics

Fabrics are on-premises network regions that include a group of switches and other networking devices that provide connectivity to your applications and endpoints. They may be split in different availability zones (such as pods) that are analyzed and managed by Nexus Dashboard.

There are two types of fabrics:

- **LAN**—This type of fabric contains either of these switch types:
 - **NX-OS switches**—These are a group of Nexus switches running NX-OS software. Nexus Dashboard manages and monitors NX-OS switches using best-practice templates. You can add a fresh set of NX-OS switches for greenfield deployments or onboard existing NX-OS switches to an existing fabric for incremental management and monitoring.
 - **Non-Nexus switches and devices**—Nexus Dashboard supports onboarding a variety of non-Nexus switches and service appliances such as IOS-XE, IOS-XR, firewalls, load balancers for a Campus VXLAN EVPN fabric, IPN, ISN, core, edge, and Layer 4 to Layer 7 service deployments.
- **ACI**—This type of fabric consists of multiple Nexus 9000 switches running ACI software managed by an Application Policy Infrastructure Controller (APIC) cluster. You can onboard existing ACI fabrics to Nexus Dashboard for continuously streaming telemetry. You can also onboard these ACI switches as an offline snapshot to provide a point-in-time analysis. An added benefit is that you can use ACI fabrics in air-gapped environments where Nexus Dashboard may not have direct connectivity to the ACI fabric.

Fabrics can also be broken down as either **local** or **remote**:

- **Local**: A fabric that is present in this cluster.
- **Remote**: A fabric that is not present in this cluster, but rather is present in another cluster that is part of this multi-cluster connectivity.

And finally, local fabrics can be broken down in the following manner:

- **Online fabrics**: Fabrics that are connected to Nexus Dashboard over the network.
- **Snapshot fabrics**: Fabrics that are referenced by a snapshot for use in one-time analysis or demonstrations. They may or may not be connected to Nexus Dashboard over the network.

Grouping fabrics and clusters

There are several ways to group fabrics and clusters together in Nexus Dashboard:

- [Grouping fabrics](#)
- [Grouping clusters](#)

Grouping fabrics

The method that you use to group fabrics together differs depending on the type of fabric:

- **NX-OS fabrics:**
 - You can use **fabric groups** to create groups of VXLAN fabrics to form a VXLAN fabric group, or to support logical groups of LAN or IPFM fabrics for simplified management. For more information, see [Create fabric groups](#).



You cannot group ACI fabrics together into fabric groups. See the **ACI fabrics** bullet below for information on grouping ACI fabrics.

- You can also establish inter-fabric connectivity using an **Inter-Fabric** link type through **Connectivity > Links** in your NX-OS fabric. You can then choose how you want to establish inter-fabric connectivity, such as connecting two NX-OS fabrics together using inter-fabric links with MACsec or establishing inter-fabric connectivity using VRF Lite, where you would use VRF Lite to establish external connectivity from a LAN fabric to an external Layer 3 domain. For more information, see [Create inter-fabric links](#).
- **ACI fabrics:** You can use the **Orchestration** feature through Nexus Dashboard to connect multiple ACI fabrics together, and consolidate and deploy tenants, along with network and policy configurations, across multiple ACI fabrics. For more information, see [Connecting Multiple ACI Fabrics and Working with Orchestration](#).

Grouping clusters

You can use either of these methods to group clusters together:

- **Multi-cluster connectivity:** You can establish connectivity between multiple Nexus Dashboard and APIC clusters for ease of access to all the clusters, as well as access to any of the fabrics running on any of the connected clusters. For more information, see [Connecting Clusters](#).
- **Multi-cluster fabric groups:** You can create groups of VXLAN fabrics to form a multi-cluster fabric group where VXLAN fabrics span across clusters for simplified management. For more information, see [Create multi-cluster fabric groups](#).

Mapping fabric types

This table provides mapping information between the fabric types that existed in releases prior to Nexus Dashboard release 4.1.1 and the new fabric types available in release 4.1.1.

Pre-4.1.1 fabrics		4.1.1 fabric types
Fabric technologies	Fabric types	

LAN		
VXLAN EVPN	Data Center VXLAN EVPN	Data Center VXLAN EVPN - iBGP
eBGP VXLAN EVPN	BGP fabric	Data Center VXLAN EVPN - eBGP
VXLAN EVPN	Campus VXLAN EVPN	Campus VXLAN EVPN
eBGP Routed	BGP fabric	BGP fabric
Classic LAN	Enhanced Classic LAN	Enhanced Classic LAN
Classic LAN	Classic LAN	Legacy Classic LAN
Custom	External connectivity network	External and inter-fabric connectivity network
Custom	Custom network	External and inter-fabric connectivity network
Custom	Multi-site external network	External and inter-fabric connectivity network
LAN Monitor	LAN Monitor	External and inter-fabric connectivity network
VXLAN EVPN	VXLAN EVPN Multi-Site	VXLAN (fabric group)
Multi-Fabric Domain	Fabric Group	Classic (fabric group)
IPFM		
IPFM	IPFM	IPFM
IPFM	IPFM Classic	IPFM classic
Generic Multicast	IPFM Classic	IPFM classic
Multi-Fabric Domain	Fabric Group	IPFM (fabric group)

Guidelines and limitations

- MTU requirements when onboarding ACI fabrics: Nexus Dashboard, when used with remote leaf switches and Multi-Site architectures in Cisco ACI, requires jumbo MTU (Maximum Transmission Unit) settings on the network switches and on the Nexus Dashboard data interfaces.

Viewing LAN fabric information

These sections describe how to view fabric information.

- [View information on local online fabrics](#)
- [View information on local snapshot fabrics](#)
- [View information on remote fabrics](#)

View information on local online fabrics

Follow these steps to view information on local online fabrics.

1. Navigate to the **Fabrics** page.

Manage > Fabrics

2. Click the **Fabrics** tab, then click the **Local** subtab.
3. In the drop-down list underneath the **Local** tab, choose **Online fabrics**.

The table displays this information on already-configured local online fabrics.

Field	Description
Name	Displays the name of the fabric.
Type	Displays the type of the fabric.
Anomaly level	<p>Displays the highest level of anomalies currently detected in the fabric. Anomalies are classified into these levels.</p> <ul style="list-style-type: none">▪ Critical: Shown when the network is down, such as when a fabric is not operational.▪ Major: Shown when connectivity to a given prefix or endpoint could be compromised, such as overlapping IP addresses or subnets.▪ Warning: Shown when the network is impacted, such as when connectivity to a given prefix or endpoint is degraded.

Field	Description
Advisory level	<p>Displays the highest level of advisories currently detected in the fabric. Advisories are classified into these levels.</p> <p>Advisory levels display only if you have enabled telemetry. Otherwise, Nexus Dashboard displays NA as the advisory level.</p> <ul style="list-style-type: none"> ▪ Critical: Shown when there are unsupported infrastructure and the severity of the bugs associated with notices is Severity1, such as when switches in a fabric are running under End-of-Life conditions or when a critical (Severity1) field notice or PSIRT has been issued for a switch or software version currently running in your network. ▪ Major: Shown when the severity of the bugs associated with notices is Severity2, such as when a critical (Severity2) field notice or PSIRT has been issued for a switch or software version currently running in your network. ▪ Warning: Shown when there is support for potentially at-risk infrastructure and the severity of the bugs associated with notices is Severity3, such as when switches in a fabric are approaching end-of-life conditions, or when a Severity3 field notice or PSIRT has been issued for a switch or software version currently running in your network.
License tier	Displays the license tier for the software features that is being used in the fabric.
ASN	Displays the ASN for the fabric.
Connectivity status	Shows whether the fabric is reachable from the Nexus Dashboard cluster.
Fabric group	Shows when a fabric is a member of a fabric group.
Features	Shows the features that are enabled on the fabric.

4. If you want to modify the columns shown in the table, click the gear icon at the top right of the table, then choose the columns that you want to display in the table.

You can also perform these actions on this page.

Action	Description
Actions > Edit fabric settings	<p>Choose a fabric to edit, then click Actions > Edit fabric settings.</p> <ul style="list-style-type: none"> ▪ Make the necessary changes and click Save. ▪ Click Close to discard the updates without saving any changes.
Actions > Delete fabric	Choose a fabric to delete, then click Actions > Delete Fabric . Click Confirm to delete the fabric.
Actions > Re-register	Click Re-register to re-register the connection between Nexus Dashboard and a fabric.

View information on local snapshot fabrics

Follow these steps to view information on local snapshot fabrics.

1. Navigate to the **Fabrics** page. **Manage > Fabrics**
2. Click the **Fabrics** tab, then click the **Local** subtab.
3. In the drop-down list underneath the **Local** tab, choose **Snapshot fabrics**.

The table displays this information on already-configured local snapshot fabrics.

Field	Description
Name	Displays the name of the fabric.
Anomaly level	<p>Displays the anomaly level of the fabric. Anomalies are classified into these levels.</p> <ul style="list-style-type: none">▪ Critical: Shown when the network is down, such as when a fabric is not operational.▪ Major: Shown when connectivity to a given prefix or endpoint could be compromised, such as overlapping IP addresses or subnets.▪ Warning: Shown when the network is impacted, such as when connectivity to a given prefix or endpoint is degraded.
Advisory level	<p>Displays the advisory level of the fabric. Advisories are classified into these levels.</p> <p>Advisory levels display only if you have enabled telemetry. Otherwise, Nexus Dashboard displays NA as the advisory level.</p> <ul style="list-style-type: none">▪ Critical: Shown when there are unsupported infrastructure and the severity of the bugs associated with notices is Severity1, such as when switches in a fabric are running under End-of-Life conditions or when a critical (Severity1) field notice or PSIRT has been issued for a switch or software version currently running in your network.▪ Major: Shown when the severity of the bugs associated with notices is Severity2, such as when a critical (Severity2) field notice or PSIRT has been issued for a switch or software version currently running in your network.▪ Warning: Shown when there is support for potentially at-risk infrastructure and the severity of the bugs associated with notices is Severity3, such as when switches in a fabric are approaching end-of-life conditions, or when a Severity3 field notice or PSIRT has been issued for a switch or software version currently running in your network.
Type	Displays the type of the fabric.
Connectivity to Nexus Dashboard Insights	Shows the connectivity status for this snapshot fabric.
Features	Shows the features that are enabled on the fabric.

Field	Description
Onboarding Time	Shows the date and time when this snapshot fabric was onboarded.

- If you want to modify the columns shown in the table, click the gear icon at the top right of the table, then select the columns that you want to display in the table.

You can also perform this action on this page.

Action	Description
Actions > Delete fabric	Choose a fabric to delete, then click Actions > Delete Fabric . Click Confirm to delete the fabric.

View information on remote fabrics

Follow these steps to view information on remote fabrics.

- Navigate to the **Fabrics** page. **Manage > Fabrics**
- Click the **Fabrics** tab, then click the **Remote** subtab.

The table displays this information on already-configured remote fabrics.

Field	Description
Name	Displays the name of the fabric.
Type	Displays the type of the fabric.
Owner	Shows the owner of a Nexus Dashboard cluster.
License tier	Displays the license tier for the software features enabled on the fabric.
Features	Shows the features that are enabled on the fabric.

- If you want to modify the columns shown in the table, click the gear icon at the top right of the table, then select the columns that you want to display in the table.

You can also perform this action on this page.

Action	Description
Actions > Edit fabric settings	Choose a fabric to edit, then click Actions > Edit fabric settings .

Creating or onboarding local online LAN fabrics

- [Guidelines and limitations: Onboarding LAN fabrics](#)
- [Create or onboard a local online LAN fabric](#)

Guidelines and limitations: Onboarding LAN fabrics

Following are the guidelines and limitations when creating or onboarding local online LAN fabrics:

- NAT is not supported between the cluster IP addresses and switch management IP addresses
- The enhanced classic LAN fabric does not support flow collection nor flow telemetry.

Create or onboard a local online LAN fabric

Follow these steps to create a local online LAN fabric:

1. Click **Manage > Fabrics** to navigate to the **Fabrics** page.

You can view, create, delete, and modify fabrics and fabric groups in this page.

2. Click the **Fabrics** tab, then click the **Local** subtab.

See [Understanding LAN and ACI fabrics and fabric groups](#) for more information on the different types of local fabrics.

3. Click **Create fabric**.

The **Create/Onboard Fabric** page appears. Navigate through the **Create/Onboard Fabric** wizard to create a local online fabric.

- [Select a category](#)
- [Select a type](#)
- [Settings](#)
- [Advanced settings](#)
- [Fabric summary](#)
- [Fabric creation](#)

Select a category

Follow these steps to select a category.

1. Determine what sort of fabric you want to create.
 - **Create new LAN fabric:** Choose this option to provision a new network comprising of Cisco NX-OS, IOS-XE, or IOS-XR devices through Nexus Dashboard.
 - **Onboard existing LAN fabric:** Choose this option to preserve an existing Cisco NX-OS, IOS-XE, or IOS-XR devices network's configuration, and to monitor and automate the deployment

of VXLAN, IP media, and Ethernet fabrics through Nexus Dashboard.


2. Click **Next**.

You advance to [Select a type](#).

Select a type

Follow these steps to select a type. See [Mapping fabric types](#) for mapping information between pre-4.1.1 fabric types and the fabric types that are available in Nexus Dashboard release 4.1.1.

1. Choose the type of fabric that you want to create.

Fabric type	Description
VXLAN	<p>Used to automate a VXLAN BGP EVPN fabric for Cisco Nexus (NX-OS) or Catalyst (IOS-XE) switches.</p> <p>Choose which type of VXLAN fabric you want to create:</p> <ul style="list-style-type: none">▪ Data Center VXLAN EVPN: Used for a VXLAN EVPN deployment with Nexus 3000 and 9000 switches.▪ Campus VXLAN EVPN: Used for VXLAN EVPN campus deployments with Catalyst 9000 and Nexus 9000 switches as border gateways.
Classic LAN	<p>Used to automate the provisioning of a two- or three-tier traditional classic Ethernet network. This is a fabric for Nexus 3000/7000/9000-based Access-Aggregation-Core Classic LAN architectures. This type is also known as enhanced classic LAN.</p>
AI	<p>Not shown when onboarding an existing LAN fabric. Used to automate the provisioning of a Nexus (NX-OS) fabric for top performance artificial intelligence and machine learning (AI) networks using RoCEv2.</p> <p>Choose which type of AI fabric that you want to create:</p> <ul style="list-style-type: none">▪ AI Routed: Used for eBGP-based CLOS fabrics using Nexus 9000 switches optimized for AI deployments.▪ AI VXLAN EVPN: Used for a VXLAN EVPN deployment with Nexus 3000 and 9000 switches optimized for AI deployments. <div> The new AI Fabric type deployment options are only available for greenfield deployments.</div>
External Inter-fabric connectivity and	<p>Used to automate provisioning of a network that might include Cisco NX-OS, IOS-XE, IOS-XR, or third-party devices for monitoring or provisioning. This includes use cases for external connectivity and multi-site inter-connectivity (IPNs or ISNs).</p>
Routed	<p>Not shown when onboarding an existing LAN fabric. Used to automate provisioning of BGP-based CLOS fabric on Cisco Nexus (NX-OS) switches.</p>
IP Fabric for Media	<p>Used to automate the creation of IP-based broadcast production networks on Cisco Nexus (NX-OS) switches.</p>

2. Click **Next**.

You advance to [Settings](#).

Settings


Follow these steps to configure the settings.

1. Configure the parameters and capabilities of the new fabric.

Field	Description
Configuration Mode	<p>Determine which type of configuration mode that you want to use to configure this fabric.</p> <ul style="list-style-type: none">▪ Default: If you use the Default (basic) mode to create the fabric, you will provide a minimal number of configuration entries during this process so that you are able to create a fabric quickly and easily, and default recommended entries, based on Cisco best practices, are used for the remaining configuration entries that are available for this fabric type. You can then modify those remaining default entries at any point after you've created the fabric using the information provided in Editing fabric settings.▪ Advanced: If you use the Advanced mode to create the fabric, you see several more advanced configuration settings in this page. In addition, another step appears in the Create/Onboard Fabric workflow (Advanced settings), where you can create the fabric using more advanced configuration settings.
Name	Enter a unique name for the fabric. In Nexus Dashboard 4.1.1, fabric name length is limited to 64 characters.
Location	Choose the location for the fabric.
Overlay routing protocol	<p>Shown in the following situations:</p> <ul style="list-style-type: none">▪ You chose one of the Data Center VXLAN EVPN fabric types in 2. Select a type (either Data Center VXLAN EVPN or AI Data Center VXLAN EVPN).▪ You clicked Advanced in the Configuration Mode field above. <p>Choose the type of overlay routing protocol:</p> <ul style="list-style-type: none">▪ iBGP: Interior Border Gateway Protocol. Used to set up a link between the same Autonomous Systems (AS).▪ eBGP: Exterior Border Gateway Protocol. Used to establish a connection between two distinct Autonomous Systems.

Field	Description
VRF-Lite protocol	<p>Shown in the following situations:</p> <ul style="list-style-type: none"> • You chose Classic as the fabric type in 2. Select a type. • You clicked Advanced in the Configuration Mode field above. <p>Choose the VRF-Lite Agg-Core/Edge or Collapsed Core-WAN peering protocol:</p> <ul style="list-style-type: none"> • EBGp • OSPF • None: Nexus Dashboard does not configure the peering protocol if the None option is selected. You must manually configure the peering protocol with this option, if necessary.

Field	Description
BGP ASN/BGP ASN for spines	<p>Available for these fabric types.</p> <ul style="list-style-type: none"> ▪ VXLAN (BGP ASN for spines) ▪ Classic (BGP ASN) ▪ AI (BGP ASN for spines) ▪ External and inter-fabric connectivity (BGP ASN) ▪ Routed (BGP ASN for spines) <p>Enter the BGP autonomous system number (ASN) for the fabric's spine switches. Valid entries are:</p> <p>1-4294967295 1-65535[.0-65535]</p> <p>where:</p> <ul style="list-style-type: none"> ▪ 1-4294967295 denotes an integer (whole) value from 1 to 4294967295 (inclusive), or ▪ 1-65535[.0-65535] denotes a decimal value, where the left side of the decimal is a number from 1 to 65535 (inclusive) and the right side of the decimal is a value from 0 to 65535 (inclusive). <p>Following are examples of valid entries that would fall into either category above:</p> <ul style="list-style-type: none"> ▪ 1 ▪ 31 ▪ 7654321 ▪ 1.1 ▪ 1.65535 ▪ 2.999 ▪ 65535.1 <p>Following are examples of <i>invalid</i> entries that would violate the guidelines shown above:</p> <ul style="list-style-type: none"> ▪ -5 ▪ 1.300000 ▪ 65536.1

Field	Description
AI QoS & queuing policy	<p>Shown in the following situations:</p> <ul style="list-style-type: none"> You chose one of the AI fabric types in 2. Select a type (either AI Routed or AI Data Center VXLAN EVPN). You clicked Advanced in the Configuration Mode field above. <p>Choose the queuing policy from the drop-down list based on the predominant fabric link speed for certain switches in the fabric:</p> <ul style="list-style-type: none"> AI_Fabric_QOS_400G: Enable QoS queuing policies for an interface speed of 400 Gb. This is the default value. AI_Fabric_QOS_100G: Enable QoS queuing policies for an interface speed of 100 Gb. AI_Fabric_QOS_25G: Enable QoS queuing policies for an interface speed of 25 Gb. <p>For more information, see AI QoS classification and queuing policies.</p>
License tier for fabric	<p>Choose the licensing tier for the fabric:</p> <ul style="list-style-type: none"> Essentials Advantage Premier <p>Click on the information icon (i) next to License tier to see what functionality is enabled for each license tier.</p>
Enabled features	<p>Check the Telemetry check box to enable telemetry for the fabric. This is the equivalent of enabling the Nexus Dashboard Insights service in previous releases.</p> <div>  <p>The Telemetry option is only available if you check the Telemetry check box in Settings page while configuring the parameters during fabric creation.</p> </div>
<p>The following fields appear if:</p> <ul style="list-style-type: none"> You enabled Telemetry in the Enabled features field, and You clicked Advanced in the Configuration Mode field. 	
Telemetry collection	<p>This option becomes available if you choose to enable Telemetry in the Enabled features field above.</p> <p>Choose either Out-of-band or In-band for telemetry collection.</p>
Telemetry streaming via	<p>This option becomes available if you choose to enable Telemetry in the Enabled features field above.</p> <p>Choose either IPv4 or IPv6 for telemetry streaming.</p>

Field	Description
Telemetry source interface	This option becomes available if you choose to enable Telemetry in the Enabled features field above. Enter the source interface for telemetry streaming.
Telemetry VRF	This option becomes available if you choose to enable Telemetry in the Enabled features field above. Enter the appropriate VRF instance in this field.
Security domain	Choose the security domain for the fabric.

- Click **Next** to advance to the next step in the fabric creation process.
 - If you chose **Default** in the **Configuration Mode** field above, the next step in the **Create/Onboard Fabric** workflow is [Fabric summary](#).
 - If you chose **Advanced** in the **Configuration Mode** field above, the next step in the **Create/Onboard Fabric** workflow is [Advanced settings](#).

Advanced settings

When you create a fabric using these procedures, the standard workflow allows you to create a fabric using the bare minimum settings so that you are able to create a fabric quickly and easily. However, you can make more advanced configurations on this fabric, either by clicking **Advanced** in the **Configuration Mode** field as part of this fabric creation workflow or after you have completed this fabric configuration workflow.

Follow these steps to configure the advanced settings.

- Locate the article that provides information on each of the available fields for the configuration settings for your fabric type.

See [Editing fabric settings](#) for more information on these advanced configuration settings for different types of fabrics.

- Make the necessary advanced configuration changes for your fabric using the Editing Fabric Settings article for your fabric type.
- Return to these procedures after you have completed the advanced configurations for your fabric type, then click **Next**.

You advance to [Fabric summary](#).

Fabric summary

Follow these steps to view the fabric summary.

- Verify all of the information that is shown in the **Fabric summary** page is correct.
- If all of the information shown in the page looks correct, click **Submit**.

You advance to [Fabric creation](#).

Fabric creation

Follow these steps to monitor the fabric creation.

1. Monitor the creation of the fabric.

You can see the fabric creation progress in this page. If you close the session, your fabric will still be created.

2. When the fabric creation is completed, determine your next step.
 - o To bring up the **Overview** page for the fabric that you just created, click **View fabric details**.
 - o To go back to the main **Fabrics** page with all of the configured fabrics in your cluster listed, click **View all fabrics**.
 - o If you want to create another fabric at this time, click **Create another fabric**, then repeat these procedures.

Add switches to the fabric

Follow these procedures to add switches to the fabric:

1. Navigate to the **Overview** page for the fabric.
 - a. Click **Manage > Fabrics > Fabrics**.
 - b. Click on the fabric where you want to add a switch.
2. Click the **Inventory** tab.
3. Click the **Switches** subtab.
4. Click **Actions > Add Switches**, then enter the necessary information to add the switch to the fabric. Refer to the "Adding Switches to Your Fabric" article for more information.

Editing fabric settings

The following table provides pointers to articles that describe how to edit fabric settings for each type of fabric.

Type of Fabric	Detailed Procedures
ACI fabric	Editing ACI Fabric Settings
AI Data Center Routed fabric	Editing AI Data Center Routed Fabric Settings
AI Data Center VXLAN fabric	Editing AI Data Center VXLAN Fabric Settings
Campus VXLAN fabric	Editing Campus VXLAN Fabric Settings
Classic fabric	Editing Classic Fabric Settings
Data Center VXLAN fabric	Editing Data Center VXLAN Fabric Settings
External fabric	Editing External Fabric Settings
IPFM fabric	Editing IP Fabric for Media (IPFM) Fabric Settings
Routed fabric	Editing Routed Fabric Settings

Onboard ACI fabrics

Choose this option to set up multi-cluster connectivity for your ACI by onboarding your APIC cluster.

1. Navigate to the **Fabrics** page.

Manage > Fabrics

2. Choose **Fabrics > Local**.
3. In the dropdown underneath the **Local** tab, choose **Online fabrics**.

See [Understanding LAN and ACI fabrics and fabric groups](#) for more information on the different types of local fabrics.

4. Click **Create Fabric**.

The **Select a category** step in the fabric creation process appears.

5. In the **Onboard ACI fabric** area, click **Connect APIC cluster** to set up multi-cluster connectivity for your ACI by onboarding your APIC cluster.

The **Connect Cluster** page appears. Refer to the [Connecting Clusters](#) article for the procedures on connecting the APIC cluster.



You can also navigate to the **Connect Cluster** page by navigating to:

Admin > System Settings > Multi-cluster connectivity

and clicking on **Connect Cluster**.

Onboard snapshot LAN fabrics

To onboard a local snapshot LAN fabric:

1. Enable the snapshot feature at the system level, if necessary.
 - a. Navigate to **Admin > System Settings**.
 - b. With the **General** tab selected, locate the **Advanced Settings** area.
 - c. Determine if the **Fabric snapshot creation** feature is enabled or not.
 - If you see **Enabled** under the **Fabric snapshot creation** field, this feature has already been enabled. Go to Step 2.
 - If you see **Disabled** under the **Fabric snapshot creation** field, continue with these procedures.
 - d. Click **Edit** in the Advanced Settings area.

The **Advanced settings** slide-in pane appears.

- e. Click the checkbox next to **Enable fabric snapshot creation** to enable this feature, then click **Save**.

You will now see **Enabled** under the **Fabric snapshot creation** field.

2. Navigate to the **Fabrics** page.

Click **Manage > Fabrics** to navigate to the **Fabrics** page. You can view, create, delete, and modify fabrics and fabric groups in this page.

3. Click the **Fabrics** tab, then click the **Local** subtab.
4. In the dropdown underneath the **Local** tab, choose **Snapshot fabrics**.

See [Understanding LAN and ACI fabrics and fabric groups](#) for more information on the different types of local fabrics.

5. Click **Create fabric**.

The **Create Fabric** page appears. Navigate through the **Create Fabric** wizard to create a local snapshot fabric.

- [Select a category](#)
- [Basic settings](#)
- [Advanced settings](#)
- [Fabric summary](#)
- [Fabric creation](#)

Select a category

1. Click **Onboard Snapshot fabric**.

This allows you to onboard a snapshot fabric, which will have no Internet connectivity on the

controllers or switches.

2. Click **Next**.

You advance to [Basic settings](#).

Basic settings

1. Determine if you have a fabric snapshot file.
 - o If you have a fabric snapshot file, choose that file or drag and drop to upload the file into the **Upload file** area.
 - o If you don't have a fabric snapshot file yet:
 - a. Click **Download Script** to download the [data-collectors.tar.gz](#) to your machine.
 - b. Extract the file you downloaded and run the data collection script. Follow the instructions provided in the readme.md file. After the script is completed successfully, the data is collected in a [<filename>.tar.gz](#) file.



The collection script requires that you have Python3 installed on your system.

- c. Choose the file or drag and drop to upload the file into the **Upload file** area.
2. Click **Next**
 3. Enter the fabric name to identify the fabric on Nexus Dashboard.
 4. Choose the fabric location from the map to identify the fabric on Nexus Dashboard.
 5. Click **Next**.
 6. Verify the configuration.
 7. Click **Submit**.

After your fabric is onboarded and fully prepared, Nexus Dashboard will start the analysis to collect data from your fabric and display the fabric information in the **Fabrics** page. For more information, see [Fabric summary](#). The Fabric Analysis banner displays the progress of the analysis. The time to run the analysis depends on the size of the fabric.

1. Click **Next**.

You advance to [Fabric summary](#).

Fabric summary

1. Verify all of the information that is shown in the **Fabric summary** page is correct.
2. If all of the information shown in the page looks correct, click **Submit**.

Create fabric groups

You can create groups of VXLAN fabrics to form a VXLAN fabric group or to support logical groups of LAN or IPFM fabrics for simplified management.

1. Navigate to the **Fabric groups** page.

Click **Manage > Fabric** to navigate to the **Fabric** page. You can view, create, delete, and modify fabrics and fabric groups on this page.

2. Click the **Fabric groups** tab.

You can see fabric groups that have already been created on the **Fabric groups** page.

3. Click **Create Fabric Group**.

The **Create Fabric Group** page appears. Navigate through the **Create Fabric Group** wizard to create a fabric group.

- [Settings](#)
- [Fabric group summary](#)
- [Fabric group creation](#)

Settings

1. Enter a name for the fabric group in the **Name** field.
2. Choose the type of fabric group that you want to create.



If you choose **VXLAN** as the fabric group type, an additional step in the **Create fabric group** workflow appears (**Advanced settings**). Advanced settings are available only for VXLAN fabric group types and not for any other fabric group types.

Type	Description
VXLAN or Multi-Site Domain (MSD)	A VXLAN fabric group can contain individual VXLAN, external, or enhanced classic LAN fabrics. This type of fabric group allows for shared deployments for VXLAN overlays (networks and VRFs) and fabric interconnectivity.
Classic	A classic fabric group can contain enhanced classic LAN, classic LAN, or external fabrics. This fabric group allows for a combined visualization at a topology level. No group level deployments are available in this fabric group.
IPFM	An IPFM fabric group can contain individual IPFM, IPFM classic, or classic LAN fabrics. This type of fabric group allows for shared host and flow definitions.

3. Click **Next** to advance to the next step in the fabric group creation process.
 - If you chose **VXLAN** in the **Type** field above, the next step in the **Create fabric group** workflow is [Advanced settings](#).

- If you chose **Classic** or **IPFM** in the **Type** field above, the next step in the **Create fabric group** workflow is [Fabric group summary](#).

Advanced settings

1. Locate the [Editing Fabric Settings for Fabric Groups](#) article, which provides information on each of the available fields for the configuration settings for the VXLAN fabric group.
2. Make the necessary advanced configuration changes for your fabric group using the information provided in "Editing Fabric Settings for Fabric Groups" article.
3. Return to these procedures after you have completed the advanced configurations for your fabric group, then click **Next**.

You advance to [Fabric group summary](#).

Fabric group summary

1. Verify all of the information that is shown on the **Fabric group summary** page is correct.
2. If all of the information shown in the page looks correct, click **Submit**.

You advance to [Fabric group creation](#).

Fabric group creation

1. Monitor the creation of the fabric group.

You can see the fabric group creation progress on this page. If you close the session, your fabric group will still be created.

2. When the fabric group creation is completed, determine your next step.
 - To bring up the **Overview** page for the fabric group that you just created, click **View fabric group details**.
 - To go back to the main **Fabric groups** page with all of the configured fabric groups in your cluster listed, click **View all fabric groups**.
 - If you want to create another fabric group at this time, click **Create another fabric group**, then repeat these steps.

Add child fabrics to the fabric group



- If the child fabric is a VXLAN EVPN fabric, make sure that the **Underlay Routing Loopback IP Range** and the **Underlay VTEP Loopback IP Range** pool values do not overlap with any existing child fabric within the fabric group, so there is no IP address conflict on the routing loopback interface or the VTEP loopback interface.
- You can also add a fabric group that contains child fabrics.

Follow these steps to add child fabrics to a fabric group.

1. Navigate to the **Overview** page for the fabric group.
 - a. Click **Manage > Fabrics > Fabric Groups**.
 - b. Click on the fabric group where you want to add child fabrics.
2. Click the **Inventory** tab.
3. Click the **Child Fabrics** subtab.
4. Click **Actions > Add child fabric**.

A list of available child fabrics appears.

5. On the **Add child fabric** page, click a child fabric that you want to add to the fabric group.

You can add only one child fabric at a time on this page.

6. Click **Select**.

The child fabric now appears under the **Child Fabrics** tab for this fabric group.

7. Repeat steps 4 – 6 to add additional child fabrics to this fabric group.

Delete fabric groups

Follow these steps to delete fabric groups.

1. Navigate to the **Fabric groups** page.
 - a. Click **Manage > Fabric** to navigate to the **Fabric** page.
 - b. Click **Fabric groups**.
2. Choose the fabric group that you want to delete, then choose **Actions > Delete Fabric Group**.

Create multi-cluster fabric groups

You can create groups of VXLAN fabrics to form a multi-cluster fabric group where VXLAN fabrics span across clusters for simplified management.

Guidelines and limitations for creating a multi-cluster fabric group

- If you are viewing a single cluster, you can create fabrics and fabric group within a single cluster.
- On the **All Clusters > Fabrics** page, you can view all the fabrics from all the member clusters. On the **Fabric groups** tab, you can view all the fabric groups from all the member clusters.
- Even though the button displays as **Create fabric group** on the **All Clusters > Multi-cluster fabric groups** page, Nexus Dashboard creates a multi-fabric group rather than a fabric group.
- On the **All Clusters** page, if you click on **Primary**, you see the member fabrics for the primary cluster.
- The health status for an NX-OS fabric will be shown if that fabric is owned by the cluster that you are viewing, but the health status will not be shown for an NX-OS fabric that is owned by another cluster in a multi-cluster fabric group.

Configure a multi-cluster fabric group

Follow these steps to configure a multi-cluster fabric group.

1. Ensure that you have configured multiple clusters for multi-cluster connectivity. For more information on multi-cluster connectivity in Nexus Dashboard, see the section "Connecting Nexus Dashboard clusters" in [Connecting Clusters](#).
2. Navigate to **All Clusters > Clusters** and click on **All Clusters**.
3. Click **Manage > Fabrics** to navigate to the **Fabrics** page. You can view, create, delete, and modify fabrics and fabric groups on this page.
4. Click the **Multi-cluster fabric groups** tab.

You can see the multi-cluster fabric groups that have already been created on the **Multi-cluster fabric groups** page.

5. Click **Create Fabric Group**.

The **Create multi-cluster fabric group** page appears.

6. Navigate through the **Create multi-cluster fabric group** wizard to create a multi-cluster fabric group.

VXLAN is the default domain for creating multiple VXLAN and External fabrics.

7. Click **Next**.
8. Follow the steps for creating a fabric group. For more information, see [Create fabric groups](#).

Add child fabrics to a multi-cluster fabric group

Follow these steps to add child fabrics to a multi-cluster fabric group.

1. Navigate to the **All Clusters** page for a multi-cluster fabric group.
2. Click **Manage > Fabrics > Multi-Cluster fabric groups**.
3. Click on the multi-cluster fabric group where you want to add a child fabric.
4. Click the **Inventory** tab.
5. Click the **Child Fabrics** subtab.
6. Click on the child fabric that you want to add to the multi-cluster fabric group.
7. Click **Actions > Add child fabric**.

A list of eligible child fabrics appears on the **Add child fabric** page.

8. On the **Add child fabric** page, click a child fabric that you want to add to a multi-cluster fabric group.

You can add only one child fabric at a time on this page.

9. Click **Select**.

The child fabric now appears under the **Child Fabrics** tab for this multi-cluster fabric group.

10. From the **Actions** drop-down list, click **Recalculate and deploy**.
11. Repeat steps 6 - 10 to add additional child fabrics to this multi-cluster fabric group.

Remove child fabrics from a multi-cluster fabric group

Follow these steps to remove child fabrics from a multi-cluster fabric group.

1. Navigate to the **All Clusters** page for a multi-cluster fabric group.
2. Click **Manage > Fabrics > Multi-Cluster fabric groups**.
3. Click on the multi-cluster fabric group where you want to remove a child fabric.
4. Click the **Inventory** tab.
5. Click the **Child Fabrics** subtab.
6. Click on the child fabric that you want to remove from the multi-cluster fabric group.
7. From the **Actions** drop-down list, click **Actions > Remove Child Fabric**.
8. Click **Ok**.

The child fabric no longer displays under the **Child Fabrics** tab for this multi-cluster fabric group.

9. From the **Actions** drop-down list, click **Recalculate and deploy**.
10. Repeat steps 5 - 9 to remove additional child fabrics from this multi-cluster fabric group.

Back up and restore multi-cluster fabric group configurations

- [\[Back Up multi-cluster fabric group configurations\]](#)
- [Restore multi-cluster fabric group configurations](#)

Back up multi-cluster fabric group configurations

Follow these steps to back up a multi-cluster fabric group configuration.

1. Navigate to the **All Clusters** page for a multi-cluster fabric group.
2. Click **Manage > Fabrics > Multi-Cluster fabric groups**.
3. Click on the appropriate multi-cluster fabric group to display the overview information for that fabric group.
4. Click **Actions > Maintenance > Backup Fabric Group**.

The **Create Fabric Backup** window appears.

5. In the **Backup Tag** area, enter a name for the backup, then click **Create Backup**.

Restore multi-cluster fabric group configurations

Follow these steps to restore a multi-cluster fabric group configuration.

1. Navigate to the **All Clusters** page for a multi-cluster fabric group.
2. Click **Manage > Fabrics > Multi-Cluster fabric groups**.
3. Click on the appropriate multi-cluster fabric group to display the overview information for that fabric group.
4. Click **Actions > Maintenance > Restore Fabric Group**.

The **Restore Fabric Group** window appears.

5. Review the backups shown on this page.

This table describes the columns that appear on the **Select Backup** tab.

Fields	Descriptions
Backup Date	Specifies the backup date.
Backup Version	Specifies the version of backup.
Backup Tag	Specifies the backup name.
Backup Type	Specifies the backup type (for example, a golden backup).

This table describes the fields that appear on the **Action** tab.

Actions	Descriptions
Mark as golden	To mark an existing backup as a golden backup, choose Mark as golden . Click Confirm in the confirmation window.
Remove as golden	To remove an existing backup from a golden backup, choose Remove as golden . Click Confirm in the confirmation window.

6. In the **Select Backup** step, click the radio button for the fabric backup that you want to restore, then click **Next**.
7. In the **Restore Preview** step, verify that the information is correct for the backup that you want to restore.

You can preview the details about the configuration in the backup file. You can also view the name and serial numbers for the switches in the Fabric backup. Click on **Delta Config** to view the configuration difference on the switches in the fabric.

8. Click **Restore Intent**.
9. In the **Restore Status** step, you can view the status of restoring the intent.
10. Click **Next** to view the preview configuration.
11. In the **Configuration Preview** step, you can resync the configurations on specific switches.

For the desired switch, check the **Switch Name** check box, and click **ReSync**.

12. Click **Deploy** to complete the **Restore Fabric Group** operation.

Migrate a Nexus Dashboard 3.2.x Orchestration-managed fabric to a multi- cluster fabric group

You can migrate a Nexus Dashboard 3.2.x Orchestration-managed fabric to a Nexus Dashboard 4.1.1 multi-cluster fabric group. For more information on a Nexus Dashboard 3.2.x multi-cluster fabric, see [Managing and Monitoring Multi-Cluster Fabrics Using One Manage, Release 12.2.2/12.2.3](#).

Prerequisites for migrating a Nexus Dashboard 3.2.x Orchestration-managed fabric to a multi-cluster fabric group

- Upgrade from Nexus Dashboard 3.2.x to Nexus Dashboard 4.1.1. For more information, see [Upgrade from Nexus Dashboard 3.2.x to Nexus Dashboard 4.1.1](#).
- Configure at least two clusters for multi-cluster connectivity. For more information, see the section "Create or onboard a local online LAN Fabric" in [Creating LAN and ACI Fabrics and Fabric Groups](#).
- Onboard all the clusters that were managed by Nexus Dashboard Orchestrator. For more information, see [Connecting Clusters, Release 4.1.1](#)
- You can add either a remote user or a local user accessing from primary cluster using the multi-cluster login domain. For more information, see the section "Add primary cluster as remote authentication domain" in [Configuring Users, Roles, and Security](#).

You need to add a local user for the VXLAN multi-cluster fabric group domain.

Guidelines and limitations for migrating a Nexus Dashboard 3.2.x Orchestration-managed fabric to a multi-cluster fabric group

- You can migrate any Nexus Dashboard 3.2.x Orchestration-managed fabric to a Nexus Dashboard 4.1.1 multi-cluster fabric group.
- After upgrading to Nexus Dashboard 4.1.1, we do not recommend adding VRFs or networks before migrating to a multi-cluster fabric group.

Upgrade from Nexus Dashboard 3.2.x to Nexus Dashboard 4.1.1

Refer to the "Upgrading an Existing Nexus Dashboard Cluster to This Release" section in the [Cisco Nexus Dashboard Deployment and Upgrade Guide, Release 4.1.x](#) for upgrade procedures.

How to migrate a Nexus Dashboard 3.2.x Orchestration-managed fabric to a multi-cluster fabric group

Follow these steps to migrate a Nexus Dashboard 3.2.x Orchestration-managed fabric to a Nexus Dashboard 4.1.1 multi-cluster fabric group.

1. Connect your Nexus Dashboard clusters. For more information, see the section "Connecting multiple ACI fabrics through the Orchestration page" in [Connecting Multiple ACI Fabrics and Working with Orchestration](#).

After migration, Nexus Dashboard converts the Multi-Site Orchestration (MSO) fabric type to a VXLAN fabric type.

You now have two connected clusters using the VXLAN fabric type.

2. Log in as a superadmin, admin, or fabric admin.



Nexus Dashboard requires a fabric admin in a multi-cluster fabric group.

3. Configure a multi-cluster fabric group as a regular user from the primary cluster. For more information, see [Create multi-cluster fabric groups](#).
4. Ensure that all the Orchestration-managed fabrics are in-sync before adding the Orchestration-managed fabrics to a multi-cluster fabric group.
5. Add the fabric groups managed by Nexus Dashboard Orchestration as a member to a multi-cluster fabric group. You should migrate Orchestration-managed fabrics from all the clusters managed by Nexus Dashboard Orchestration. For more information, see [Add child fabrics to a multi-cluster fabric group](#).
6. Perform a **Recalculate and deploy** operation.

This will migrate the multisite overlay links to the multi-cluster fabric group. You can manage the multisite underlay port configuration pushed by NDO in the future by editing the policies with the template name `multisite_dci_underlay_sitelocal_jython` from the individual fabrics. For more information, see [Working with Configuration Policies for Your Nexus Dashboard LAN or IPFM Fabrics](#).

+ After migrating an Orchestration-managed fabric from Nexus Dashboard 3.2.x, we do not expect there to be a difference in the configurations.

1. Ensure that there are no pending configurations after performing a **Recalculate and deploy** operation.
2. Create a network or add a VRF. For more information, see the sections "Working with VRFs" and "Working with networks" in [Working with Segmentation and Security for Your Nexus Dashboard VXLAN Fabric](#).

View the migrated multi-cluster fabric group on the Topology page

You can view the migrated multi-cluster fabric group on the **Topology** page after migrating your Nexus Dashboard 3.2.x Orchestration-managed fabric to a Nexus Dashboard 4.1.1 multi-cluster

fabric group.

Follow these steps to view the migrated multi-cluster fabric group details on the **Topology** page.

1. Navigate to **All Clusters > Clusters** and click on **All Clusters**.
2. Navigate to **Home > Topology**.

Nexus Dashboard displays the network topology from the connected clusters.

Additional settings

The following sections provide information for additional settings that might be necessary when creating LAN fabrics or fabric groups.

Understanding the Fabric Summary page

Click on a fabric to open the side kick panel. The following sections display the summary of the fabric:

- **Health** - Shows the health of the Fabric.
- **Alarms** - Displays the alarms based on the categories.
- **Fabric Info** - Provides basic about the Fabric.
- **Inventory** - Provides information about Switch Configuration and Switch Health.

Click the **Launch** icon to the right top corner to view the Fabric Overview.

Prerequisites to creating a fabric

- The ESXi host default setting on the vSphere Client for promiscuous mode is supported. For more information, see *ESXi Networking for Promiscuous Mode* section. The vNIC of the POD that has the Persistent IP shares the same MAC address of Nexus Dashboard bond0 or bond1 interface. Therefore, the POD sources the packets using the same MAC address of Nexus Dashboard bond0 or bond1 interfaces that are known by the VMware ESXi system.
- Configure the persistent IP addresses in Cisco Nexus Dashboard. For more information, see *Cluster Configuration* section in [Cisco Nexus Dashboard User Guide](#).

Change persistent IP address

You can change the persistent IP addresses that are assigned for mandatory pods, such as POAP-SCP and SNMP traps.

To change the persistent IP address, perform the following steps:

1. On the Nexus Dashboard Web UI, navigate to **Admin > System Settings > Fabric Management**.
2. Under **Advanced Settings**, click **Admin**.
3. In the **LAN Device Management Connectivity** field, change **Management** to **Data** or vice versa.

Changing the option results in a migration of SNMP and POAP-SCP pods to the persistent IP addresses associated with **External Service Pool** on Nexus Dashboard associated with the new **LAN Device Management Connectivity** option. After the completion of this process, the following message is displayed:

Some features have been updated. Reload the page to see latest changes.

Click **Reload the page**.

4. On the Nexus Dashboard Web UI, navigate to **Admin > System Settings > General**.
5. In the **External pools** card, click **Edit** to change the required IP addresses for **Persistent**

management IPs or **Persistent data IPs**.

6. Navigate back to **Admin > System Settings > Fabric Management > Advanced Settings > Admin**, then change the option in **LAN Device Management Connectivity** drop-down list to its initial selection.

Restoring this option to initial settings results in migration of the SNMP and POAP-SCP pods to use the updated persistent IP address from the appropriate external Service IP pool.

Configuring overlay mode

You can create a VRF or network in CLI or config-profile mode at the fabric level. The overlay mode of member fabrics in a VXLAN fabric group is set individually at the member-fabric level. Overlay mode can only be changed before deploying overlay configurations to the switches. After the overlay configuration is deployed, you cannot change the mode unless all the VRF and network attachments are removed.

If the switch has config-profile based overlays, you can import it in the **config-profile** overlay mode only. If you import it in the **cli** overlay mode, an error appears during brownfield import.

For brownfield import, if overlay is deployed as **config-profile** mode, it can be imported in **config-profile** mode only. However, if overlay is deployed as **cli**, it can be imported in either **config-profile** or **cli** modes.

To choose the overlay mode of VRFs or networks in a fabric, perform the following steps:

1. Navigate to the **Edit Fabric** page.
2. Go to the **Advanced** tab.
3. From the **Overlay Mode** drop-down list, choose **config-profile** or **cli**.

The default mode is **config-profile**.

Configuring Netflow support

Configuring Netflow at the fabric level allows you to collect, record, export, and monitor network flow and data to determine network traffic flow and volume for further analysis and troubleshooting. You can configure Netflow for VXLAN, Routed (BGP), External/inter-fabric connectivity, and Classic LAN fabric templates.

After Netflow is enabled for fabric, you can configure Netflow on a network, or an interface (VLAN, SVI, physical interface, sub-interface, or port-channel). Before enabling Netflow on the interface or network, ensure that the specified monitor name is defined in the fabric settings.

When Netflow is enabled at the fabric level, the configuration is generated for Netflow capable switches (FX/GX/EX) in the fabric except for spine/super-spine or switches with **no_netflow** policy. In a Multi-Site domain configuration, Netflow is configured per Easy Fabric and not for the entire Multi-Site domain.



Nexus Dashboard does not validate the **Netflow Monitor** name.

The following are the guidelines for Netflow configuration on other network elements:

- For VRF Lite IFC, the Netflow configuration is not inside the configuration profile, regardless of overlay mode.
- For networks, Netflow configurations are not inside the configuration profile, regardless of overlay mode.
- You can configure Netflow for Layer 2 Interface on trunk ports, access ports, dot1q tunnels, Layer2 port-channel, and VPC ports.
- You can configure Netflow for the Layer 3 interface on SVI, Routed host, L3 Port-Channel, and sub-interfaces.
- Netflow configuration for VLANs uses **vlan_netflow** Record Template. In Brownfield deployment, the Netflow configuration for VLANs is in switch freeform.
- You can enable Netflow under SVI (for routed traffic) or Vlan Configuration (for switched traffic).
- To configure IPv6 flow monitoring, use **switch_freeform** or **interface freeform**.
- Netflow configuration under the trunk or routed port is in **interface freeform**.
- For Host port resync, Netflow configuration is captured in interface freeform.
- There is no explicit support for Netflow in Intra-Fabric link or Multisite Underlay IFC. Note that you can use freeform configuration.

Netflow support for brownfield deployments

For brownfield deployments, global Netflow configuration for export, record, and monitor are not captured due to the telemetry use case. After brownfield import, to avoid global level Netflow command being removed, you can perform the following actions:

- Do not turn on strict CC.
- Include the Netflow global configuration in **switch freeform**.
- Enable Netflow in the fabric setting matching with the switch configuration.

Interface and VLAN-level Netflow configuration on the switch is captured in **freeform**.

- SVI Netflow config is captured in **switch_freeform** tied to the network.
- Netflow configuration for trunk or routed ports is in the **interface freeform**.
- Netflow configuration for VLANs is in the **switch_freeform**.
- The sub-interface configuration for VRF-Lite extensions is in **int_freeform**.

VXLAN OAM

In Nexus Dashboard, VXLAN OAM is supported on VXLAN, Routed (eBGP), External/interfabric-connectivity, and Classic LAN fabrics. You can track details such as reachability and actual path of the flows in a VXLAN EVPN based-fabric topology.

Guidelines

- OAM must be enabled on the switches before using the OAM trace.
- VXLAN OAM IPv6 is now supported.
- NX-API and NX-API on HTTP port must be enabled.

- vPC advertise-pip must be enabled.
- For switch-to-switch OAM, ensure that the VRFs are configured along with loopback interfaces with IPv4 and/or IPv6 addresses under those VRFs.
- For host-to-host OAM, ensure that the Networks are configured along with IPv4 and/or IPv6 gateway configuration.
- IPv6 underlay is supported with VXLAN OAM. To enable the VXLAN OAM support over IPv6 underlay, perform any one of the following steps:
 - On the **Topology** window:
 - Choose **Actions > Add Fabric**.
 - On the **General Parameters** tab, check the **Enable IPv6 Underlay** check box.
 - On the **Fabrics** window:
 - Choose **Actions > Create Fabric**.
 - On the **General Parameters** tab, check the **Enable IPv6 Underlay** check box.



Changing of IPv4 to IPv6 underlay is not supported for existing fabric settings.

To change the fabric settings from IPv4 to IPv6 underlay, delete the existing fabric and create new fabric with Underlay IPV6 enabled.

UI Navigation

- In the **Topology** window: Click **Actions**. Choose **VXLAN OAM** option from the drop-down list.
- From the **Fabrics** window: Choose **Manage > Fabrics**. Navigate to the fabric overview window of a fabric. Click **Actions**. Choose **VXLAN OAM** option from the drop-down list.

The VXLAN OAM window appears. The **Path Trace Settings** pane on the left displays the **Switch to Switch** and **Host to Host** tabs. Nexus Dashboard highlights the route on the topology between the source and destination switch for these two options.

The **Switch to Switch** option provides the VXLAN OAM ping and traceroute test results for the VTEP-to-VTEP use-case. Provide the following values to enable search by using the **Switch to Switch** option:

- In the **Source Switch** drop-down list, choose the source switch.
- In the **Destination Switch** drop-down list, choose the destination switch.
- From the **VRF** drop-down list, choose or enter the VRF details.
- Check the **All paths included** check box to include all the paths in the search results.

The **Host to Host** option provides the VXLAN OAM path trace results for the exact path that is taken by a given flow from the VTEP or switch that is connected to the source host to VTEP or switch that is connected to the destination host. For the **Host to Host** use-case, there are two options:

- VRF or SVI for a network is instantiated on the switches in the VXLAN EVPN fabric. In such a scenario, the IP address information of the end hosts is required.
- Layer 2 configuration for a given network is instantiated on the switches in the VXLAN EVPN fabric. In such a scenario, both the MAC and IP address information of the end hosts are required.

Provide the following values to enable search using the **Host to Host** option:

- From the **Source Host IP** field, enter the IPv4/IPv6 address of the source host.
- From the **Destination Host IP** field, enter the IPv4/IPv6 address of the destination host.
- In the **VRF** field, choose VRF from the drop-down list or enter the VRF name that is associated with the hosts.
- In the **Source Port** field, choose Layer 4 source port number from the drop-down list or enter its value.
- In the **Destination Port** field, choose destination port number or enter its value.
- In the **Protocol** field, choose the protocol value from the drop-down list or enter its value. This is the Layer 4 protocol, usually TCP or UDP.
- Check the **Layer 2 only** check box to search the VXLAN-EVPN fabric that is deployed in Layer 2 only mode for some networks, that is, Layer 2 VNIs. No SVIs or VRFs should be instantiated in the fabric for these networks when you use this search option. When you check this option, you have to enter details of the source MAC address, destination MAC address, and VNI too.

Click **Run Path Trace** to view the path trace from switch to switch or host to host.

You can view the forward path and reverse path as well in the topology. The summary of the path trace appears in the **Summary** tab. You can view the details of the forward and reverse paths as well under **Forward Path** or **Reverse Path** tabs. Filter the results by attributes, if needed.

AI QoS classification and queuing policies

These sections provide information about the AI QoS classification and queuing policies.

- [Understanding AI QoS classification and queuing policies](#)
- [Guidelines and limitations for AI QoS classification and queuing policies](#)
- [Configure AI QoS classification and queuing policies](#)
- [Create a policy using the custom QoS templates](#)

Understanding AI QoS classification and queuing policies

Support is available for configuring a low latency, high throughput, and lossless fabric configuration that can be used for artificial intelligence (AI) and machine learning (ML) traffic.

The AI QoS feature allows you to:

- Easily configure a network with homogeneous interface speeds, where most or all of the links run at 400Gb, 100Gb, or 25Gb speeds.
- Provide customizations to override the predominate queuing policy for a host interface.

When you apply the AI QoS policy, Nexus Dashboard will automatically pre-configure any inter-fabric links with QoS and system queuing policies, and will also enable Priority Flow Control (PFC). If you enable the AI QoS feature on a VXLAN EVPN fabric, then the Network Virtual (NVE) interface will have the attached AI QoS policies.

Use the following areas to enable this feature:

- When configuring a BGP fabric, new fields are available to enable the feature and to set the queuing policy parameters based on the interface speed.
- You can also use the following AI-specific switch templates to create custom device policies, which can be used on host interfaces:
 - **AI_Fabric_QoS_Classification_Custom**: An interface template that is available for applying a custom queuing policy to an interface.
 - **AI_Fabric_QoS_Queueing_Custom**: A switch template that is available for user-defined queuing policy configurations.

Policies defined with these custom Classification and Queueing templates can be used in various host interface policies. For more information, see [Create a policy using the custom QoS templates](#).

When enabling the AI feature, **priority-flow-control watchdog-interval on** is enabled on all of your configured devices, intra-fabric links, and all your host interfaces where Priority Flow Control (PFC) is also enabled. The PFC watchdog interval is for detecting whether packets in a no-drop queue are being drained within a specified time period. This release also adds the **Priority flow control watch-dog interval** field on the **Advanced** tab. When you create or edit a Data Center VXLAN EVPN fabric or other fabrics and AI is enabled, you can set the **Priority flow control watch-dog interval** field to a non-system default value (the default is 100 milliseconds). For more information on the PFC watchdog interval for Cisco NX-OS, see [Configuring a priority flow control watchdog Interval](#) in the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide*.

If you perform an upgrade from an earlier release, and then do a **Recalculate and deploy**, you may see additional **priority-flow-control watchdog-interval on** configurations.

Guidelines and limitations for AI QoS classification and queuing policies

Following are the guidelines and limitations for the AI QoS and queuing policy feature:

- On Cisco Nexus N9K-C9808 and N9K-C9804 series switches, the command **priority-flow-control watch-dog-interval** is not supported in either global or interface configuration modes and the command **hardware qos nodrop-queue-thresholds queue-green** is not supported in global configuration mode.
- Cisco Nexus N9K-C9808 and N9K-C9804 series switches only support AI fabric type from NX-OS version 10.5(1) and later.
- This feature does not automate any per-interface speed settings.
- This feature is supported only on Nexus devices with Cisco Cloud Scale technology, such as the Cisco Nexus 9300-FX2, 9300-FX3, 9300-GX, and 9300-GX2 series switches.
- This feature is not supported in fabrics with devices that are assigned with a ToR role.

Configure AI QoS classification and queuing policies

Follow these steps to configure AI QoS and queuing policies:

1. Enable AI QoS and queuing policies at the fabric level.
 - a. Create a fabric as you normally would.
 - b. In the **Advanced** tab in those instructions, make the necessary selections to configure AI QoS

and queuing policies at the fabric level.

- c. Configure any remaining fabric-level settings as necessary in the remaining tabs.
- d. When you have completed all the necessary fabric-level configurations, click **Save**, then click **Recalculate and deploy**.

At this point in the process, the network QoS and queuing policies are configured on each device, the classification policy is configured on NVE interfaces (if applicable), and priority flow control and classification policy is configured on all intra-fabric link interfaces.

2. For host interfaces, selectively enable priority flow control, QoS, and queuing by editing the policy associated with that host interface.

See [Working with Connectivity for LAN Fabrics](#) for more information.

- a. Within a fabric where you enabled AI QoS and queuing policies in the previous step, click the **Interfaces** tab.

The configured interfaces within this fabric are displayed.

- b. Locate the host interface where you want to enable AI QoS and queuing policies, then click the box next to that host interface to select it and click **Actions > Edit**.

The **Edit Interfaces** page is displayed.

- c. In the **Policy** field, verify that the policy that is associated with this interface contains the necessary fields that will allow you to enable AI QoS and queuing policies on this host interface.

For example, these policy templates contain the necessary AI QoS and queuing policies fields:

- int_access_host
- int_dot1q_tunnel_host
- int_pvlan_host
- int_routed_host
- int_trunk_host

- d. Locate the **Enable priority flow control** field and click the box next to this field to enable Priority Flow Control for this host interface.
- e. In the **Enable QoS Configuration** field, click the box next to this field to enable AI QoS for this host interface.

This enables the QoS classification on this interface if AI queuing is enabled at the fabric level.

- f. If you checked the box next to the **Enable QoS Configuration** field in the previous step and you created a custom QoS policy using the procedures provided in [Create a policy using the custom QoS templates](#), enter that custom QoS classification policy in the **Custom QoS Policy for this interface** field to associate that custom QoS policy with this host interface, if necessary.

If this field is left blank, then Nexus Dashboard will use the default QOS_CLASSIFICATION policy, if available.

- g. If you created a custom queuing policy using the procedures provided in [Create a policy using the custom QoS templates](#), enter that custom queuing policy in the **Custom Queuing Policy for this interface** field to associate that custom queuing policy with this host interface, if desired.
- h. Click **Save** when you have completed the AI QoS and queuing policy configurations for this host interface.

Create a policy using the custom QoS templates

Follow these procedures to use the custom QoS templates to create a policy, if desired. See [Managing Your Template Library](#) for general information on templates.

1. Within a fabric where you enabled AI QoS and queuing policies, click **Inventory > Switches**, then double-click the switch that has the host interface where you enabled AI QoS and queuing policies.

The **Switch Overview** page for that switch appears.

2. Choose **Configuration Policies > Policies**.
3. Click **Actions > Add policy**.

The **Create Policy** page appears.

4. Set the priority and enter a description for the new policy.

Note that the priority for this policy must be lower (must come before) the priority that was set for the host interface.

5. In the **Select Template** field, click the **No Policy Selected** text.

The **Select Policy Template** page appears.

6. Select the appropriate custom Classification or Queuing template from the list, then click **Select**.

The following templates are specific to the AI QoS and queuing policies feature. Use these templates to create policies that can be used on one or more host interfaces:

- **AI_Fabric_QOS_Classification_Custom**: An interface template that is available for applying a custom queuing policy to an interface.
- **AI_Fabric_QOS_Queueing_Custom**: A switch template that is available for user-defined queuing policy configurations.

7. Make the necessary QoS classification or queuing configurations in the template that you selected, then click **Save**.

Any custom QoS policy created using these procedures are now available to use when you configure QoS and queuing policies for the host interface.

Configuring downstream VNI

In a VXLAN fabric, the Layer 3 and Layer 2 virtual network identifier (VNIs) are centrally managed using a VXLAN fabric group for inter-fabric connectivity. All the VXLAN fabrics within the VXLAN

fabric group use the same VNI value for the Layer 3 or Layer 2 network. The problem is that before the fabrics are brought in for inter-fabric connectivity, the fabrics have been managed as standalone fabrics with existing VRFs and networks. The Layer 2 and Layer 3 overlays have been configured independently and have conflicting VNIs among fabrics.

There are two types of VNI conflicts:

- The same VNI is used by different VRFs or networks in different VXLAN fabrics.
- The same VRF or network uses different VNIs in different VXLAN fabrics.

This feature is supported when creating or editing fabric types:

- VXLAN
- Campus VXLAN

Prior to Nexus Dashboard release 4.1.1, you could not add a VXLAN fabric to a VXLAN fabric group if there was a VNI conflict. With Nexus Dashboard release 4.1.1, downstream VNI (DSVNI) allows VXLAN fabrics with a VNI conflict to communicate with each other. On the border gateway, Nexus Dashboard uses different VNIs for exchanging intra-fabric and inter-fabric traffic. The existing VNI continues to be used for intra-fabric traffic. Stitching of the VNI occurs at the border gateways for inter-fabric traffic between fabrics using different VNIs.

Nexus Dashboard added downstream VNI options in a VXLAN fabric group for configuring a global Layer 2 VNI and a Layer 3 VNI pool. The two ranges should not conflict with VNIs already in use in existing VXLAN fabrics. We suggest picking from the high end of the VNI range for the global **Layer 3 VXLAN VNI Global Range** and the **Layer 2 VXLAN VNI Global Range** in the **General Parameters** page for the fabric. Nexus Dashboard generates a new VRF and a network VNI allocation based on these two ranges.



Enabling downstream VNI in a VXLAN fabric group does not affect existing VRFs and networks.

When Nexus Dashboard allocates a VNI for a VRF or a network, and its intra-fabric VNI is different from the downstream VNI, Nexus Dashboard requires additional configuration on the border gateway.

VRF CLIs on border gateways for downstream VNI

```
ip extcommunity-list standard <vrf-name> seq 10 permit rt 2324:50005
route-map MS-FABRIC-TO-EXTERNAL-RMAP permit 100
  match extcommunity <vrf-name>
  set extcomm-list <vrf-name> delete

vrf context <vrf-name>
  address-family ipv4 unicast
    route-target both <local-asn>:<DSVNI>
    route-target both <local-asn>:<DSVNI> evpn
  address-family ipv6 unicast
    route-target both <local-asn>:<DSVNI>
```

```
route-target both <local-asn>:<DSVNI> evpn
```

Network CLIs on border gateways for downstream VNI

```
ip extcommunity-list standard <network-name> permit rt 27:30001
route-map MS-FABRIC-TO-EXTERNAL-RMAP permit 200
  match extcommunity <network-name>
  set extcomm-list <network-name> delete
route-map MS-FABRIC-TO-EXTERNAL-RMAP permit 65535
evpn
  vni <local-VNI> l2
  route-target both <local-asn>:<DSVNI>
```

Nexus Dashboard generates an additional route target containing the downstream VNI allowing border gateways in different fabrics to exchange routes.

extcommunity-list and **route-map** removes the local VNI from the BGP updates towards Data Center Interconnectivity (DCI). This prevents route leaking if the same VNI is used in a different fabric for a different VRF or network. **route-map MS-FABRIC-TO-EXTERNAL-RMAP** is applied to all multi-site overlay inter-fabric links if downstream VNI is enabled in the fabric group, regardless of whether the inter-fabric link is manually created or auto-generated by Nexus Dashboard.

Benefits of downstream VNI

- Resolves overlapping VNIs when using conflicted VNIs for a VRF or a network in different fabrics because of not changing the default VNI pool, which is the same for all VXLAN fabrics
- Supports route exchange using a route target
- Prevents route leaking

Use cases for downstream VNI

When you add a new VXLAN fabric to a VXLAN fabric group for inter-fabric connectivity, and if Nexus Dashboard detects a VNI conflict, Nexus Dashboard allocates a new downstream VNI range.

- If the same VNI is used by a VRF or a network in a VXLAN fabric group and the incoming VXLAN fabric (**vrf1** in the fabric group and **vrf2** in the incoming VXLAN fabric), Nexus Dashboard allocates a new VNI from the global VNI pool for both the VRF and the network. In this example, Nexus Dashboard allocates a new VNI for **vrf1** and **vrf2**.
- If a VRF or a network use a different VNI in a VXLAN fabric group and the incoming VXLAN fabric, Nexus Dashboard allocates a new VNI, if the VNI in the VXLAN fabric group is not already in the global VNI range.
- In these two use cases, if there is a VNI conflict between the VXLAN fabric group and the incoming VXLAN fabric, and the VNI in the incoming VXLAN fabric is already in the global VNI range, you cannot add the VXLAN fabric to the VXLAN fabric group. In this use case, you can edit the global **L2 VNI** and or the **L3 VNI** range so that the VNI range does not overlap with the VNIs already used in the incoming VXLAN fabric.

Supported platforms

Ensure that all border gateways have the correct platform and NX-OS version that supports downstream VNI. For the list of supported platforms and NX-OS versions, see the [Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide](#).

Guidelines and limitations for downstream VNI

- You must use unique names for VRFS and networks.

This means that **vrf foo** on fabric1 and **vrf foo** on fabric2 refer to the same VRF. The same applies for network names.

- You cannot disable downstream VNI in a VXLAN fabric group if there are any existing VRFs or networks with a fabric VNI that differs from the VNI of the VXLAN fabric group.

These features are not supported for downstream VNI:

- Using the same VRF or network but using a different VRF name or network name in a different VXLAN fabric
 - IPv6 underlay
 - Security groups
 - PVLAN
 - TRM and TRMv6
 - CloudSec VXLAN tunnel encryption
 - Brownfield deployment where downstream VNI is already configured on a switch
-

First Published: 2025-01-31
Last Modified: 2025-01-31