



Understanding LAN Fabrics, Release 12.2.2

Table of Contents

New and Changed Information	1
Understanding LAN Fabrics	2
Terms Used In This Document	2
Fields Shown In the Fabrics Window	2
Actions Available In the Fabrics Window	2
Understanding the Fabric Summary Window	3
Understanding Fabric Templates	3
Prerequisites to Creating a Fabric	4
Creating a Fabric	6
Locating Information on LAN Fabric Templates	7
Changing Persistent IP Address	8
Issues with Persistent IP Addresses After Disabling and Enabling NDFC	8
Enabling ESXi Networking for Promiscuous Mode	10
Overlay Mode	12
Netflow Support	13
Netflow Support for Brownfield deployments	13
VXLAN OAM	15
Copyright	17

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
There were no major changes from the previous release.		

Understanding LAN Fabrics

The following sections provide information that will help you understand LAN fabrics:

- [Terms Used In This Document](#)
- [Fields Shown In the Fabrics Window](#)
- [Actions Available In the Fabrics Window](#)
- [Understanding the Fabric Summary Window](#)
- [Understanding Fabric Templates](#)
- [Prerequisites to Creating a Fabric](#)

Terms Used In This Document

The following terms are used in this document:

- Greenfield Deployments: Applicable for provisioning new VXLAN EVPN fabrics and eBGP-based routed fabrics.
- Brownfield Deployments: Applicable for existing VXLAN EVPN fabrics:
 - Migrate CLI-configured VXLAN EVPN fabrics to Nexus Dashboard Fabric Controller using the Data Center VXLAN EVPN fabric template.
 - NFM migration to Cisco Nexus Dashboard Fabric Controller using the Data Center VXLAN EVPN fabric template.

Note that in this document the terms *switch* and *device* are used interchangeably.

For information about upgrades, refer to the *Cisco Installation and Upgrade Guide for LAN Controller Deployment*.

Fields Shown In the Fabrics Window

The following table describes the fields that appear on **Manage > Fabrics**.

Field	Description
Fabric Name	Displays the name of the fabric.
Fabric Technology	Displays the fabric technology based on the fabric template.
Fabric Type	Displays the type of the fabric—Switch Fabric, LAN Monitor, or External
ASN	Displays the ASN for the fabric.
Fabric Health	Displays the health of the fabric.

Actions Available In the Fabrics Window

The following table describes the action items in the Actions menu drop-down list, that appear on

Manage > Fabrics.

Action Item	Description
Create Fabric	From the Actions drop-down list, select Create Fabric . For more instructions, see Creating a Fabric .
Edit Fabric	Select a fabric to edit. From the Actions drop-down list, select Edit Fabric . Make the necessary changes and click Save . Click Close to discard the changes.
Delete Fabric	Select a fabric to delete. From the drop-down list, select Delete Fabric . Click Confirm to delete the fabric.

Understanding the Fabric Summary Window

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.

Understanding Fabric Templates

The following table provides information about the available fabric templates:

Type of Fabric	Description	REST Template Name	API	Detailed Procedures
Data Center VXLAN EVPN	Fabric for a VXLAN EVPN deployment with Nexus 9000 and 3000 switches.	Easy_Fabric		Data Center VXLAN EVPN
Enhanced Classic LAN	Fabric for a fully automated 3-tier Classic LAN deployment with Nexus 9000 and 7000 switches.	Easy_Fabric_Classic		Enhanced Classic LAN
Campus VXLAN EVPN	Fabric for a VXLAN EVPN Campus deployment with Catalyst 9000 switches.	Easy_Fabric_IOS_XE		Campus VXLAN EVPN

BGP Fabric	Fabric for an eBGP based deployment with Nexus 9000 and 3000 switches. Optionally VXLAN EVPN can be enabled on top of the eBGP underlay.	Easy_Fabric_eBGP	BGP Fabric
Custom Network	Fabric for flexible deployments with a mix of Nexus and Non-Nexus devices.	External_Fabric	Custom Network
Fabric Group	Domain that can contain Enhanced Classic LAN, Classic LAN, and External Connectivity Network fabrics.	Fabric_Group	Fabric Group and LAN Monitor
Classic LAN	Fabric to manage a legacy Classic LAN deployment with Nexus switches.	LAN_Classic	Classic LAN
LAN Monitor	Fabric for monitoring Nexus switches for basic discovery and inventory management.	LAN_Monitor	Fabric Group and LAN Monitor
VXLAN EVPN Multi-Site	Domain that can contain multiple VXLAN EVPN Fabrics (with Layer-2/Layer-3 Overlay Extensions) and other Fabric Types.	MSD_Fabric	VXLAN EVPN Multi-Site
Classic IPFM	Fabric to manage or monitor existing Nexus 9000 switches in an IP Fabric for Media Deployment.	IPFM_Classic	IPFM and Classic IPFM
IPFM	Fabric for a fully automated deployment of IP Fabric for Media Network with Nexus 9000 switches.	Easy_Fabric_IPFM	IPFM and Classic IPFM
Multi-Site External Network	Fabric to interconnect VXLAN EVPN for Multi-Site deployments with a mix of Nexus and Non-Nexus devices	External_Fabric	Multi-Site External Network
External Connectivity Network	Fabric for core and edge router deployments with a mix of Nexus and Non-Nexus devices.	External_Fabric	External Connectivity Network

Prerequisites to Creating a Fabric

- From Cisco NDFC Release 12.1.2e, the ESXi host default setting on the vSphere Client for promiscuous mode is supported. For more information, see *ESXi Networking for Promiscuous Mode* section. From Nexus Dashboard release 2.3.1c, 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](#).

Creating a Fabric

To create a fabric using the Cisco Nexus Dashboard Fabric Controller Web UI:

1. Choose **Manage > Fabrics**.
2. From the **Actions** drop-down list, select **Create Fabric**.
3. Enter the fabric name and click **Choose Fabric**.
4. Specify the values for the fabric settings and click **Save**.

Locating Information on LAN Fabric Templates

The following table provides pointers to documents that give information and instructions specifically for each type of LAN fabric template.

Type of Fabric	Detailed Procedures
BGP Fabric	BGP Fabric
Campus VXLAN EVPN	Campus VXLAN EVPN
Classic IPFM	IPFM and Classic IPFM
Classic LAN	Classic LAN
Custom Network	Custom Network
Data Center VXLAN EVPN	Data Center VXLAN EVPN
Enhanced Classic LAN	Enhanced Classic LAN
External Connectivity Network	External Connectivity Network
Fabric Group	Fabric Group and LAN Monitor
IPFM	IPFM and Classic IPFM
LAN Monitor	Fabric Group and LAN Monitor
Multi-Site External Network	Multi-Site External Network
VXLAN EVPN Multi-Site	VXLAN EVPN Multi-Site

Changing Persistent IP Address

You can change 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 Cisco NDFC Web UI, navigate to **Settings > Server Settings > Admin** under **LAN Device Management Connectivity** drop-down list change **Management** to **Data** or conversely.

Changing option results in migration of SNMP and POAP-SCP pods to the persistent IP addresses associated with **External Service Pool** on Nexus Dashboard connected 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**.

2. On Cisco Nexus Dashboard Web UI, navigate to **Infrastructure > Cluster Configuration > General**, in **External Service Pools** card, change the required IP addresses for **Management Service IP Usage** or **Data Service IP Usage**.
3. Navigate to NDFC Web UI **Server Settings** page, 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.

Issues with Persistent IP Addresses After Disabling and Enabling NDFC

A known condition exists where you might have persistent IP addresses configured and assigned through NDFC, but then you disable the NDFC app through Nexus Dashboard, which puts those persistent IP addresses in an unassigned state. This results in the following situations when you enable the NDFC app again at a later date through Nexus Dashboard:

- For persistent IP addresses that were deleted, the deleted persistent IP addresses will reappear in an assigned state when you enable the NDFC app again at a later date through Nexus Dashboard.
- For persistent IP addresses that were changed, the original persistent IP addresses will reappear rather than the new, updated persistent IP addresses when you enable the NDFC app again at a later date through Nexus Dashboard.

The following workarounds are available to resolve this issue:

- You can force NDFC to reassign the persistent IP addresses by changing the discovery to management or data, and then changing it back. See the section "Persistent IP Addresses" in [Nexus Dashboard Infrastructure Management](#) for more information.
- If you are running on Nexus Dashboard release 3.1 or later, you can also perform the following as a workaround:

1. Take a backup of the cluster.
 2. Restore from the backup with the **Ignore External Service IP Configuration** option checked after re-initializing the cluster and adding new IP addresses in the correct external IP pool.
- See [Backing Up and Restoring LAN Operational Mode Setups](#) for more information.

Enabling ESXi Networking for Promiscuous Mode

From Cisco NDFC Release 12.1.2e, you can run NDFC on top of virtual Nexus Dashboard (vND) instance with promiscuous mode that is disabled on port groups that are associated with Nexus Dashboard interfaces where External Service IP addresses are specified. vND comprises Nexus Dashboard management interface and data interface. By default, for fabric controller persona, two external service IP addresses are required for the Nexus Dashboard management interface subnet.

Before the NDFC Release 12.1.2e, if Inband management or Endpoint Locator or POAP feature was enabled on NDFC, you must also enable promiscuous mode for the Nexus Dashboard data or fabric interface port-group. This setting was mandatory for traffic flow that is associated for these features.

Enabling promiscuous mode raise risk of security issues in NDFC, it is recommended to set default setting for promiscuous mode.



- Disabling promiscuous mode is supported from Cisco Nexus Dashboard Release 2.3.1c.
- You can disable promiscuous mode when Nexus Dashboard nodes are layer-3 adjacent on the Data network, BGP is configured, and fabric switches are reachable through the data interface.
- You can disable promiscuous mode when Nexus Dashboard interfaces are layer-2 adjacent to switch mgmt0 interface.

If Inband management or EPL is enabled, you must specify External Service IP addresses in the Nexus Dashboard data interface subnet. You can disable promiscuous mode for the Nexus Dashboard data or fabric interface port-group. For more information, refer to the [Cisco Nexus Dashboard Deployment Guide](#).



Default option for promiscuous mode is **Reject**.

1. Log into your **vSphere** Client.
2. Navigate to the ESXi host.
3. Right-click the host and choose **Settings**.

A sub-menu appears.

4. Choose **Networking > Virtual Switches**.

All the virtual switches appear as blocks.

5. Click **Edit Settings** of the VM Network.
6. Navigate to the **Security** tab.
7. Update the **Promiscuous mode** settings as follows:
 - Check the **Override** check box.
 - Choose **Accept** from the drop-down list.

8. Click **OK**.

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 of an MSD fabric 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 you upgrade from Cisco DCNM Release 11.5(x), the existing config-profile mode functions the same.

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** window.
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**.

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. From Cisco NDFC Release 12.0.2, you can configure Netflow for Easy Fabrics, Easy Fabric eBGP, External Fabric, and LAN Classic 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.



NDFC does not validate the **Netflow Monitor** name.

The following are the guidelines for Netflow configuration on other networks 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 will be 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 Fabric Controller, VXLAN OAM is supported on VXLAN Fabric, eBGP VXLAN Fabric, External, and Lan Classic fabric technologies. 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.
- From Cisco NDFC Release 12.1.1e, 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 **LAN 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 **LAN 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 Fabric Controller 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.

Copyright

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

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 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/go/trademarks>. 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. (1110R)

© 2017–2024 Cisco Systems, Inc. All rights reserved.