



External Connectivity Network, Release 12.2.2/12.2.3

Table of Contents

New and Changed Information	1
Fabric Templates Associated with External_Fabric	2
External Fabrics	3
Guidelines and Limitations	3
Move an External Fabric Under an MSD Fabric	5
External Fabric Depiction in an MSD Fabric Topology	5
Creating an External Fabric	6
General Parameters	6
Advanced	7
Resources	10
Configuration Backup	10
Bootstrap	11
Flow Monitor	14
Adding Switches to the External Fabric	17
Switch Settings for External Fabrics	19
Discovering New Switches	21
Adding Non-Nexus Devices to External Fabrics	24
Configuration Compliance in External Fabrics	24
Special Configuration CLIs Ignored for Configuration Compliance	26
Managing Cisco IOS-XR Devices using NDFC	26
Configuring IOS-XR as an Edge Router	27
Configuring Non-Nexus Devices for Discovery	27
Configuring IOS-XE Devices for Discovery	27
Configuring Arista Devices for Discovery	28
Configuring and Verifying Cisco IOS-XR Devices for Discovery	30
Discovering Non-Nexus Devices in an External Fabric	31
Managing Non-Nexus Devices to External Fabrics	33
Creating a vPC Setup	35
Undeploying a vPC Setup	37
Precision Time Protocol for External Fabrics	38
Copyright	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
NDFC release 12.2.2	Support for Cisco Plug and Play Connect (PnP) with out-of-band (OOB) management for Cisco Catalyst 9000 series switches in an External Connectivity Network or Custom Network fabric	With this feature, you can enable automatic Cisco Plug n Play (PnP) IP assignment for Cisco Catalyst 9000 series switches in an External Connectivity Network or a Custom Network fabric using the Create Fabric or Edit Fabric > Bootstrap tab. For more information, see External Fabrics and Creating an External Fabric .

Fabric Templates Associated with External_Fabric

Any reference to External_Fabric in this document refers to one of the following 3 fabric templates:

- Multi-Site Interconnect Network
- External Connectivity Network
- Flexible Network

The type of fabric will be seen as **External_Fabric** aka the fabric template name, in the following cases:

1. Upgrade and Restore from DCNM 11.5(4) .
2. Upgrade from NDFC 12.0.2f/12.1.1e

All existing functionalists will continue to work similarly to the previous release. You can optionally edit the fabric and choose one of the three options **Flexible Network**, **External Connectivity Network**, **Multi-Site Interconnect Network**. If you edit the fabric settings without choosing one of these options, then the default option **Flexible Network** will be picked. You can toggle between these three options as desired without any loss of functionality. The type of fabric is stored in nvPairs in a variable called **EXT_FABRIC_TYPE**. This can be optionally provided in the payload during fabric creation. If not provided, the default option of **Flexible Network** is picked.

External Fabrics

You can add switches to an external fabric.

Guidelines and Limitations

- NDFC will not generate "no router bgp". If you want to change it, go to the switch and do a "no feature bgp" followed by a re-sync, if you don't have anything and want to update the ASN.
- The external fabric is a monitor-only or managed-mode fabric.
- Beginning with Cisco Nexus Dashboard Fabric Controller release 12.2.2, you can enable Cisco PnP for automatic IP assignment for Cisco Catalyst 9000 Series switches for an External Connectivity Network or a Custom Network fabric. For more information, see the **Bootstrap** tab in [Creating an External Fabric](#).
- From Cisco Nexus Dashboard Fabric Controller release 12.0.1, Cisco IOS-XR family devices Cisco ASR 9000 Series Aggregation Services routers and Cisco Network Convergence System (NCS) 5500 Series switches are supported in an external fabric in managed mode and monitor mode. NDFC generates and pushes configurations to these switches, and configuration compliance is also enabled for these switches.
- From Cisco Nexus Dashboard Fabric Controller release 12.1.1e, you can also add Cisco 8000 Series routers to external fabrics both in managed mode and monitored mode, and configuration compliance is also supported.
- You can import, remove, and delete switches for an external fabric.
- For Inter-Fabric Connection (IFC) cases, you can choose Cisco 9000, 7000 and 5600 Series switches as destination switches in an external fabric.
- You can use non-existing switches as destination switches.
- The template that supports an external fabric is the External Connectivity Network fabric.
- If an external fabric is an MSD fabric member, then the MSD topology screen displays the external fabric with its devices, along with the member fabrics and their devices.

When viewed from an external fabric topology screen, any connections to non-Nexus Dashboard Fabric Controller managed switches are represented by a cloud icon labeled as **Undiscovered**.

- You can set up a Multi-Site or a VRF-Lite IFC by manually configuring the links for the border devices in the VXLAN fabric or by using an automatic Deploy Border Gateway Method or VRF Lite IFC Deploy Method. If you are configuring the links manually for the border devices, we recommend using the Core Router role to set up a Multi-Site eBGP underlay from a Border Gateway device to a Core Router and the Edge Router role to set up a VRF-Lite IFC from a border device to an edge device.
- If you are using the Cisco Nexus 7000 Series switch with Cisco NX-OS release 6.2(24a) on the LAN Classic or an External Connectivity Network fabric, make sure to enable **AAA IP Authorization** in the fabric settings.
- In NDFC release 12.2.1, Nexus 9800 switches are supported only with roles of Spine and Super-Spine in a VXLAN EVPN fabric with VXLAN OAM disabled. If a Nexus 9800 switch is positioned as a route server (if it has been assigned with a Core Router role) in an External Connectivity Network fabric, that Nexus 9800 switch must be running on NX-OS release 10.4.2 or later.

- You can discover the following non-Nexus devices in an External Connectivity Network fabric:
 - IOS-XE family devices: Cisco CSR 1000v, Cisco ASR 1000 Series routers, and Cisco Catalyst 9000 Series switches
 - IOS-XR family devices: ASR 9000 Series routers, IOS XR release 6.5.2, and Cisco NCS 5500 Series routers, IOS XR release 6.5.3
 - Arista 4.2 (any model)
- Configure all the non-Nexus devices, except Cisco CSR 1000v, before adding them to an external fabric.
- You can configure non-Nexus devices as borders. You can create an IFC between a non-Nexus device in an external fabric and a Cisco Nexus device in an easy fabric. The interfaces supported for these devices are:
 - Routed
 - Subinterface
 - Loopback
- You can configure a Cisco ASR 1000 Series router and a Cisco Catalyst 9000 Series switch as an edge router, set up a VRF-Lite IFC, and connect it as a border device with an easy fabric.
- Before a VDC reload, discover Admin VDC in the fabric. Otherwise, the reload operation does not occur.
- You can connect a Cisco data center to a public cloud using Cisco CSR 1000v. For more information, see the "Setting Up the Infra Configuration for Hybrid Cloud and Multi-Cloud Connectivity Deployment" section in the [Hybrid Cloud Connectivity Deployment for Cisco NX-OS Guide](#) for the use case.
- In an external fabric, when you add the **switch_user** policy and provide the username and password, the password must be an encrypted string that is displayed in the **show run** command.

For example:

```
username admin password 5
$5$!4sapkBh$S7B7UcPH/iVTihLKH5sgldBeS3O2X1StQsvv3cmbYd1 role network-admin
```

In this case, the entered password should be **\$5\$!4sapkBh\$S7B7UcPH/iVTihLKH5sgldBeS3O2X1StQsvv3cmbYd1**.

- For the Cisco Network Insights for Resources (NIR) release 2.1 and later, and flow telemetry, the **feature lldp** command is one of the required configurations.

Cisco Nexus Dashboard Fabric Controller pushes **feature lldp** on the switches only for the easy fabric deployments, that is, for the eBGP-routed fabric or a VXLAN EVPN fabric.

Therefore, NIR users need to enable **feature lld** on all the switches in the following scenarios:

- External fabric in monitored or managed mode
- LAN Classic fabric in monitored or managed Mode
- Backup/restore is only supported for Nexus devices on external fabrics.



Before you do a fabric or switch restore, ensure that the target device is supported. If the target device is not supported, then per-switch restore will be blocked, and the same will be shown as not supported during a fabric-wide restore.

Move an External Fabric Under an MSD Fabric

You should go to the MSD fabric page to associate an external fabric as its member.

1. On the **Overview > Topology** page, click within the MSD-Parent-Fabric. From the **Actions** drop-down list, select **Move Fabrics**.

The **Move Fabric** page displays. It contains a list of fabrics. The external fabric is displayed as a standalone fabric.

2. Select the radio button next to the external fabric and click **Add**.

Now, in the **Scope** drop-down list at the top right, you can see that the external fabric appears under the MSD fabric.

External Fabric Depiction in an MSD Fabric Topology

The MSD topology page displays MSD member fabrics and external fabrics together. The external fabric External65000 is displayed as part of the MSD topology.



When you deploy networks or VRFs for the VXLAN fabric, the deployment page (MSD topology view) shows the VXLAN and external fabrics that are connected to each other.

Creating an External Fabric

To create an external fabric using the Cisco Fabric Controller Web UI, perform the following steps:

1. Choose **Manage > Fabrics**.
2. From the **Actions** drop-down list, select **Create Fabric**.
3. Enter a unique name for the fabric and click **Choose Fabric**.
4. From the drop-down list, choose the **External Connectivity Network** template.

The tabs and their fields in the screen are explained in the following sections. The fabric-level parameters are included in these tabs.

- [General Parameters](#)
- [Advanced](#)
- [Resources](#)
- [Configuration Backup](#)
- [Bootstrap](#)
- [Flow Monitor](#)

5. When you have completed the necessary configurations, click **Save**.
 - Click on the fabric to display a summary in the slide-in pane.
 - Click on the **Launch** icon to display the **Fabric Overview**.

General Parameters

The **General Parameters** tab is displayed by default. The fields in this tab are described in the following table.

Field	Description
BGP ASN	This field becomes editable if you selected ebgp in the Routing Protocol field. Enter the BGP AS number the fabric is associated with. This must be same as existing fabric.

Field	Description
Fabric Monitor Mode	Clear the check box if you want Nexus Dashboard Fabric Controller to manage the fabric. Keep the check box selected to enable a monitor-only external fabric. From the Cisco Nexus Dashboard Fabric Controller Release 12.1.1e, you can also add Cisco 8000 Series Routers to external fabrics both in managed mode and monitored mode. When you create an Inter-Fabric Connection (IFC) from a VXLAN fabric to this external fabric, the BGP AS number is referenced as the external or neighbor fabric AS number. When an external fabric is set to Fabric Monitor Mode Only, you cannot deploy configurations on its switches. If you click Deploy Config, it displays an error message. The configurations must be pushed for non-Nexus devices before you discover them in the fabric. You cannot push configurations in the monitor mode.
Enable Performance Monitoring (For NX-OS and IOS XE Switches Only)	Check this check box to enable performance monitoring on NX-OS switches only. Ensure that you do not clear interface counters from the command-line interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both clear counters and clear counters snmp commands (not all switches have the clear counters snmp command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the clear counters interface ethernet slot/port command followed by the clear counters interface ethernet slot/port snmp command. This can lead to a one time spike.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations.

Advanced

The fields in the **Advanced** tab are described in the following table.

Field	Description
Power Supply Mode	Choose the appropriate power supply mode.
Enable MPLS Handoff	Select the check box to enable the MPLS Handoff feature. For more information, see MPLS SR and LDP Handoff .
Underlay MPLS Loopback Id	Specifies the underlay MPLS loopback ID. The default value is 101.
Enable AAA IP Authorization	Enables AAA IP authorization, after IP authorization is enabled on the AAA server.

Field	Description
Enable NDFC as Trap Host	Check this check box to enable the Nexus Dashboard Fabric Controller as a trap host.
Enable CDP for Bootstrapped Switch	Check the check box to enable CDP for the bootstrapped switch.
Enable NX-API	Specifies enabling of NX-API on HTTPS. This check box is unchecked by default.
NX-API HTTPS Port Number	Specifies the NX-API HTTPS port number.
Enable HTTP NX-API	Specifies enabling of NX-API on HTTP. This check box is unchecked by default. Enable this check box and the Enable NX-API check box to use HTTP. If you uncheck this check box, the applications that use NX-API and supported by Cisco Nexus Dashboard Fabric Controller, such as Endpoint Locator (EPL), Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using the HTTPS instead of HTTP.  If you check the Enable NX-API check box and the Enable NX-API on HTTP check box, applications use HTTP.
NX-API HTTP Port Number	Specifies the NX-API HTTP port number.

Field	Description
Inband Mgmt	For External and Classic LAN fabrics, this knob enables Nexus Dashboard Fabric Controller to import and manage switches with inband connectivity (reachable over switch loopback, or a routed interface, or SVI interfaces), in addition to management of switches with out-of-band connectivity (reachable over the switch mgmt0 interface). The only requirement is that for inband-managed switches, there should be IP reachability from Nexus Dashboard Fabric Controller to the switches over the Nexus Dashboard data interface, also known as an inband interface. For this purpose, static routes may be needed on the Nexus Dashboard Fabric Controller, that in turn can be configured from Admin > System Settings > Routes. After enabling inband management, during discovery, provide the IPs of all the switches to be imported using inband management and set maximum hops to 0. Nexus Dashboard Fabric Controller has a precheck that validates that the inband-managed switch IPs are reachable over the Nexus Dashboard data interface. After completing the precheck, Nexus Dashboard Fabric Controller discovers and learns about the interface on that switch that has the specified discovery IP in addition to the VRF that the interface belongs to. As part of the process of switch import/discovery, this information is captured in the baseline intent that is populated on the Nexus Dashboard Fabric Controller. For more information, see the section "Inband Management in External Fabrics and LAN Classic Fabrics" in Configuring Inband Management, Inband POAP Management, and Secure POAP.  Bootstrap or POAP is only supported for switches that are reachable over out-of-band connectivity, that is, over switch mgmt0. The various POAP services on the Nexus Dashboard Fabric Controller are typically bound to the eth1 or out-of-band interface. In scenarios, where Nexus Dashboard Fabric Controller eth0/eth1 interfaces reside in the same IP subnet, the POAP services are bound to both interfaces.
Enable Precision Time Protocol (PTP)	Enables PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. You can also edit the PTP Source Loopback Id and PTP Domain Id fields. For more information, see the section Precision Time Protocol for External Fabrics .
PTP Source Loopback Id	Specifies the loopback interface ID loopback that is used as the source IP address for all Precision Time Protocol (PTP) packets. The valid values range from 0-1023. The PTP loopback ID cannot be the same as RP, Phantom RP, NVE, or the MPLS loopback ID. Otherwise, an error is generated. The PTP loopback ID can be the same as the BGP loopback or user-defined loopback that is created from Nexus Dashboard Fabric Controller. If the PTP loopback ID is not found during a Save & Deploy , the following error is generated: Loopback interface to use for the PTP source IP is not found. Please create a PTP loopback interface on all the devices to enable the PTP feature.

Field	Description
PTP Domain Id	Specifies the PTP domain ID on a single network. The valid values range from 0-127.
Fabric Freeform	You can apply configurations globally across all the devices that are discovered in the external fabric using this freeform field. The devices in the fabric should belong to the same device type and the fabric should not be in monitor mode. The different device types are: ▪ NX-OS ▪ IOS-XE ▪ IOS-XR ▪ Others Depending on the device type, enter the configurations accordingly. If some of the devices in the fabric do not support these global configurations, they go out-of-sync or fail during the deployment. Hence, ensure that the configurations you apply are supported on all the devices in the fabric or remove the devices that do not support these configurations.
AAA Freeform Config	You can apply AAA configurations globally across all devices that are discovered in the external fabric using this freeform field.

Resources

The fields in the **Resources** tab are described in the following table.

Field	Description
Subinterface Range Dot1q	The subinterface 802.1Q range and the underlay routing loopback IP address range are autopopulated.
Underlay Loopback IP Range MPLS	Specifies the underlay MPLS SR or LDP loopback IP address range. The IP range should be unique, that is, it should not overlap with IP ranges of the other fabrics.

Configuration Backup

The fields in the **Configuration Backup** tab are described in the following table.

Field	Description
Hourly Fabric Backup	Check the check box to enable an hourly backup of the fabric configurations and the intent. You can enable an hourly backup for fresh fabric configurations and the intent as well. If there is a configuration push in the previous hour, Nexus Dashboard Fabric Controller takes a backup. In case of the external fabric, the entire configuration on the switch is not converted to intent on Nexus Dashboard Fabric Controller as compared to the VXLAN fabric. Therefore, for the external fabric, both the intent and the running configuration are backed up. Intent refers to configurations that are saved in Nexus Dashboard Fabric Controller, but yet to be provisioned on the switches. The hourly backups are triggered during the first 10 minutes of the hour.
Scheduled Backup	Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.
Scheduled Time	Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the Scheduled Fabric Backup check box. Select both the check boxes to enable both backup processes. The backup process is initiated after you click Save. The scheduled backups are triggered exactly at the time that you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status. You can also initiate the fabric backup on the Overview > Topology page. Click Backup Fabric in the Actions drop-down list. The backups contain the running configuration and the intent that is pushed by Nexus Dashboard Fabric Controller. Configuration compliance forces the running configuration to be the same as the Nexus Dashboard Fabric Controller configuration. Note that for the external fabric, only some configurations are part of the intent and the remaining configurations are not tracked by Nexus Dashboard Fabric Controller. Therefore, as part of the backup, both Nexus Dashboard Fabric Controller intent and the running configuration from the switch are captured.

Bootstrap

You can configure automatic Cisco Plug n Play (PnP) IP assignment of Cisco Catalyst 9000 Series switches in an External Connectivity Network or a Custom Network fabric by configuring the following.

1. Navigate to the **Admin > Certificate Management > Bootstrap Certificates** tab, and click the **Upload Certificate** button.

The **Upload Certificate - Bootstrap Server** dialog displays.

2. Drag and drop your bootstrap certificate file to the dialog box or browse to the location of your file.

The following are the accepted file types: .pem, .cer, .key, or .crt.

3. Enter your password for your bootstrap server and click **Upload**.
4. Ensure that you configure the **Bootstrap Script Download Protocol** field in **Admin > System Settings > LAN-Fabric** as **http**.
5. On the **Create Fabric > Bootstrap** or the **Edit Fabric > Bootstrap** tab, configure the following:
 - a. Check the **Enable Bootstrap (For NX-OS and IOS-XE(Cat9K) Switches Only)** check box.
 - b. Check the **Enable Local DHCP Server** check box.
 - c. Check the **Enable Plug n Play for Cat9K** check box.

The fields in the **Bootstrap** tab are described in the following table.

Field	Description
Enable Bootstrap (For NX-OS and IOS-XE(Cat9K) Switches Only)	Check this check box to enable the bootstrap feature for NX-OS and IOS-XE Cisco Catalyst 9000 Series switches. After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment.
Enable Inband POAP	Check this check box to enable inband POAP.  You must enable Inband Mgmt on the Advanced tab to enable this option.
Enable Local DHCP Server	Enable the Local DHCP Server check box and enter details for the remaining mandatory fields. From Cisco NDFC Release 12.1.1e, you can choose inband POAP or out-of-band POAP for external fabrics.
Enable Plug n Play for Cat9K	Check this check box to enable PnP for automatic IP assignment for Cisco Catalyst 9000 Series switches.  Cisco PnP is supported for Cisco Catalyst 9000 Series switches only.

Field	Description
DHCP Version	Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, Switch Mgmt IPv6 Subnet Prefix field is disabled. If you select DHCPv6, the Switch Mgmt IP Subnet Prefix is disabled.  Cisco Nexus Dashboard Fabric Controller IPv6 POAP is not supported with Cisco Nexus 7000 Series switches. Cisco Nexus 9000 and 3000 Series switches support IPv6 POAP only when switches are either Layer 2 adjacent (eth1 or out-of-band subnet must be a /64) or they are Layer 3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.  NDFC supports Cisco PnP for IPv4 only. There is no support for IPv6. If you do not select this check box, Nexus Dashboard Fabric Controller uses the remote or external DHCP server for automatic IP address assignment.
Domain name	Specify the domain name for the DHCP server PnP block.
DHCP Scope Start Address and DHCP Scope End Address	Specifies the first and last IP addresses of the IP address range to be used for the switch out-of-band POAP.
Switch Mgmt Default Gateway	Specifies the default gateway for the management VRF on the switch.
Switch Mgmt IP Subnet Prefix	Specifies the prefix for the mgmt0 interface on the switch. The prefix range is 8-30.
DHCP scope and management default gateway IP address specification	If you specify the management default gateway IP address 10.0.1.1 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.2 and 10.0.1.254.
Switch Mgmt IPv6 Subnet Prefix	Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be from 112 through 126. This field is editable if you enable IPv6 for DHCP.
Enable AAA Config	Select this check box to include AAA configs from Advanced tab during device bootup.
Bootstrap Freeform Config (Optional)	Enter other commands as needed. For example, if you are using AAA or remote authentication-related configurations, add these configurations in this field to save the intent. After the devices boot up, they contain the intent that is defined in the Bootstrap Freeform Config field. Copy-paste the running-config to a freeform config field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see Enabling Freeform Configurations on Fabric Switches.

Field	Description
DHCPv4 Multi Subnet Scope	Specifies the field to enter one subnet scope per line. This field is editable after you check the Enable Local DHCP Server check box. The format of the scope should be defined as: DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix Example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24

Flow Monitor

The fields in the **Flow Monitor** tab are described in the following table.

Field	Description
Enable NetFlow	Check this check box to enable NetFlow on VTEPs for this fabric. By default, NetFlow is disabled. On Enable, NetFlow configuration will be applied to all VTEPS that support NetFlow.  When NetFlow is enabled on the fabric, you can choose not to have NetFlow on a particular switch by having a dummy no_netflow PTI. If NetFlow is not enabled at the fabric level, an error message is generated when you enable NetFlow at the interface, network, or VRF level. For information about NetFlow support for Cisco NDFC, see the section "Netflow Support" in Understanding LAN Fabrics.
NetFlow Exporter	Click Actions > Add to add one or more NetFlow exporters. This exporter is the receiver of the NetFlow data. The fields on this area are: • Exporter Name - Specifies the name of the exporter. • IP - Specifies the IP address of the exporter. • VRF - Specifies the VRF over which the exporter is routed. • Source Interface - Enter the source interface name. • UDP Port - Specifies the UDP port over which the NetFlow data is exported. Click Save to configure the exporter. Click Cancel to discard. You can also choose an existing Netflow exporter and select Actions > Edit or Actions > Delete to perform the relevant actions.

Field	Description
NetFlow Record	Click Actions > Add to add one or more NetFlow records. The fields on this area are: • Record Name - Specifies the name of the record. • Record Template - Specifies the template for the record. Enter one of the record templates names. In Release 12.0.2, the following two record templates are available for use. You can create custom NetFlow record templates. Custom record templates that are saved in the template library are available for use here. • netflow_ipv4_record - to use the IPv4 record template. • netflow_l2_record - to use the Layer 2 record template.
Is Layer 2 Record	Check this check box if the record is for a Layer 2 NetFlow. Click Save to configure the report. Click Cancel to discard. You can also choose an existing record and select Actions > Edit or Actions > Delete to perform the relevant actions.
NetFlow Monitor	Click Actions > Add to add one or more NetFlow monitors. The fields on this area are: • Monitor Name - Specifies the name of the monitor. • Record Name - Specifies the name of the record for the monitor. • Exporter1 Name - Specifies the name of the exporter for the NetFlow monitor. • Exporter2 Name - (optional) Specifies the name of the secondary exporter for the NetFlow monitor. The record name and exporters referred to in each NetFlow monitor must be defined in the Netflow Record and Netflow Exporter areas. Click Save to configure the monitor. Click Cancel to discard. You can also choose an existing Netflow monitor and select Actions > Edit or Actions > Delete to perform the relevant actions.

Field	Description
NetFlow Sampler	Click Actions > Add to add one or more NetFlow samplers. The fields on this area are: ▪ Sampler Name - Specifies the name of the Netflow sampler. ▪ Number of Samples - Specifies the number of Netflow samples. ▪ Number of Packets in Each Sampling - Specifies the number of packets in each Netflow sampling.  The Netflow Sampler is applicable to Cisco Nexus 7000 Series switches only. Click Save to configure the Netflow sampler. Click Cancel to discard. You can also choose an existing Netflow sampler and select Actions > Edit or Actions > Delete to perform the relevant actions.

Click **Save**.

After the external fabric is created, the external fabric topology page comes up.

After creating the external fabric, add switches to it.

Adding Switches to the External Fabric

Switches in each fabric are unique, and hence, each switch can only be added to one fabric. To add switches to the external fabric, perform the following steps:

1. Choose **Manage > Inventory > Switches**. From the Actions drop-down list, select **Add Switches**

You can also add switches to a Fabric from **Manage > Fabrics**. Select a fabric and view the **Summary**. On the **Switches** tab, from the **Actions** drop-down list, select **Add switches** to add switches to the selected Fabric.

From Topology, right click on the Fabric and select **Add Switches**.

2. Select **Discover** to discover new switches. Select **Move Neighbor Switches** to add existing switches to the Fabric.
3. If you select **Discover** option, perform the following steps:
 - a. Enter the IP address (Seed IP) of the switch.
 - b. In the **Authentication Protocol** field, from the drop-down list, select the appropriate protocol to add switches to the Fabric.
 - c. Choose the device type from the **Device Type** drop-down list.

The options are **NX-OS**, **IOS XE**, **IOS XR**, and **Other**.

- Select **NX-OS** to discover a Cisco Nexus switch.
- Select **IOS XE** to discover a CSR device.
- Select **IOS XR** to discover an ASR device.
- Select **Other** to discover non-Cisco devices.

Refer the *Adding non-Nexus Devices to External Fabrics* section for more information on adding other non-Nexus devices.

Config compliance is disabled for all non-Nexus devices except for Cisco CSR 1000v.

- a. Enter the administrator username and password of the switch.
- b. Click **Discovery Switches** at the bottom part of the screen.

The Scan Details section comes up shortly. Since the Max Hops field was populated with 2, the switch with the specified IP address and switches two hops from it are populated.

Select the check boxes next to the concerned switches and click **Add Switches** into fabric.

You can discover multiple switches at the same time. The switches must be properly cabled and connected to the Nexus Dashboard Fabric Controller server and the switch status must be manageable.

The switch discovery process is initiated. The **Progress** column displays the progress. After Nexus Dashboard Fabric Controller discovers the switch, click **Close** to revert to the previous screen.

4. If you select **Move Neighbor Switches** option, select the switch and click **Move Switch**.

The selected switch is moved to the External Fabric.

Switch Settings for External Fabrics

External Fabric Switch Settings vary from the VXLAN fabric switch settings. Double-click on the switch to view the Switch Overview screen to edit/modify options.

The options are:

Set Role - By default, no role is assigned to an external fabric switch. You can assign desired role to the fabric. Assign the Core Router role for a Multi-Site Inter-Fabric Connection (IFC) and the Edge Router role for a VRF Lite IFC between the external fabric and VXLAN fabric border devices.



Changing of switch role is allowed only before executing **Deploy Config. vPC Pairing** - Select a switch for vPC and then select its peer.

Change Modes - Allows you to modify the mode of switch from Active to Operational.

Manage Interfaces - Deploy configurations on the switch interfaces.

Straight-through FEX, Active/Active FEX, and breakout of interfaces are not supported for external fabric switch interfaces.

View/edit Policies - Add, update, and delete policies on the switch. The policies you add to a switch are template instances of the templates available in the template library. After creating policies, deploy them on the switch using the Deploy option available in the View/edit Policies screen.

History - View per switch deployment history.

Recalculate Config - View the pending configuration and the side-by-side comparison of the running and expected configuration.

Deploy Config - Deploy per switch configurations.

Discovery - You can use this option to update the credentials of the switch, reload the switch, rediscover the switch, and remove the switch from the fabric.

Click **Deploy** from the Actions drop-down list. The template and interface configurations form the configuration provisioning on the switches.

When you click **Deploy**, the **Deploy Configuration** screen comes up.

Click **Config** at the bottom part of the screen to initiate pending configuration onto the switch. The **Deploy Progress** screen displays the progress and the status of configuration deployment.

Click **Close** after the deployment is complete.



If a switch in an external fabric does not accept default credentials, you should perform one of the following actions:

- * Remove the switch in the external fabric from inventory, and then rediscover.
- * LAN discovery uses both SNMP and SSH, so both passwords need to be the same. You need to change the SSH password to match the SNMP password on the switch. If SNMP authentication fails, discovery is stopped with authentication error. If SNMP authentication passes but SSH authentication fails, Nexus Dashboard Fabric

Controller discovery continues, but the switch status shows a warning for the SSH error.

Discovering New Switches

To discover new switches, perform the following steps:

1. Power on the new switch in the external fabric after ensuring that it is cabled to the Nexus Dashboard Fabric Controller server.

Boot the Cisco NX-OS and setup switch credentials.

2. Execute the **write**, **erase**, and **reload** commands on the switch.

Choose **Yes** to both the CLI commands that prompt you to choose Yes or No.

3. On the Nexus Dashboard Fabric Controller UI, select the External Fabric. Choose **Edit Fabric** from the **Actions** drop-down list.

The **Edit Fabric** screen is displayed.

4. Click the **Bootstrap** tab and update the DHCP information.
5. Click **Save** at the bottom right part of the **Edit Fabric** screen to save the settings.
6. Double click on the Fabric to view the **Fabric Overview**.
7. On **Switches** tab, from the **Actions** drop-down list, select **Add Switches**.
8. Click the **POAP** tab.

In an earlier step, the reload command was executed on the switch. When the switch restarts to reboot, Nexus Dashboard Fabric Controller retrieves the serial number, model number, and version from the switch and displays them on the Inventory Management along screen. Also, an option to add the management IP address, hostname, and password are made available. If the switch information is not retrieved, refresh the screen using the Refresh icon at the top right part of the screen.



At the top left part of the screen, export and import options are provided to export and import the .csv file that contains the switch information. You can pre-provision a device using the import option too.

Select the checkbox next to the switch and add switch credentials: IP address and host name.

Based on the IP address of your device, you can either add the IPv4 or IPv6 address in the **IP Address** field.

You can provision devices in advance.

9. In the **Admin Password** and **Confirm Admin Password** fields, enter and confirm the admin password.

This admin password is applicable for all the switches displayed in the POAP window.



If you do not want to use admin credentials to discover switches, you can instead use the AAA authentication, that is, RADIUS or TACACS credentials for discovery only.

10. (Optional) Use discovery credentials for discovering switches.

- a. Click the **Add Discovery Credentials** icon to enter the discovery credentials for switches.
- b. In the **Discovery Credentials** window, enter the discovery credentials such as discovery username and password.

Click **OK** to save the discovery credentials.

If the discovery credentials are not provided, Nexus Dashboard Fabric Controller uses the admin user and password to discover switches.



The discovery credentials that can be used are AAA authentication based credentials, that is, RADIUS or TACACS.

The discovery credential is not converted as commands in the device configuration. This credential is mainly used to specify the remote user (or other than the admin user) to discover the switches. If you want to add the commands as part of the device configuration, add them in the **Bootstrap Freeform Config** field under the **Bootstrap** tab in the fabric settings. Also, you can add the respective policy from **View/Edit Policies** window.

11. Click **Bootstrap** at the top right part of the screen.

Nexus Dashboard Fabric Controller provisions the management IP address and other credentials to the switch. In this simplified POAP process, all ports are opened up.

After the added switch completes POAP, the fabric builder topology screen displays the added switch with some physical connections.

12. Monitor and check the switch for POAP completion.

13. Click **Deploy Config** from the **Actions** drop-down list on the **Fabric Overview** screen to deploy pending configurations (such as template and interface configurations) onto the switches.



If there is a sync issue between the switch and Nexus Dashboard Fabric Controller, the switch icon is displayed in red color, indicating that the fabric is Out-Of-Sync. For any changes on the fabric that results in the out-of-sync, you must deploy the changes. The process is the same as explained in the Discovering Existing Switches section.

The discovery credential is not converted as commands in the device configuration. This credential is mainly used to specify the remote user (or other than the admin user) to discover the switches. If you want to add the commands as part of the device configuration, add them in the **Bootstrap Freeform Config** field under the **Bootstrap** tab in the fabric settings. Also, you can add the respective policy from **View/Edit Policies** window.

During fabric creation, if you have entered AAA server information (in the **Manageability** tab), you must update the AAA server password on each switch. Else, switch discovery fails.

14. After the pending configurations are deployed, the **Progress** column displays 100% for all switches.

15. On the Topology screen, click **Refresh Topology** icon to view the update.

All switches must be in green color indicating that they are functional.

The switch and the link are discovered in Nexus Dashboard Fabric Controller. Configurations are built based on various policies (such as fabric, topology, and switch generated policies). The switch image (and other required) configurations are enabled on the switch.

16. Right-click and select History to view the deployed configurations.

Click the **Success** link in the **Status** column for more details. An example:

17. On the Nexus Dashboard Fabric Controller UI, the discovered switches can be seen in the fabric topology.

Up to this step, the POAP is completed with basic settings. All the interfaces are set to trunk ports. You must setup interfaces through the **Manage > Inventory > Interfaces** option for any additional configurations, but not limited to the following:

- o vPC pairing.
- o Breakout interfaces

Support for breakout interfaces is available for 9000 Series switches.

- o Port channels, and adding members to ports.



After discovering a switch (new or existing), at any point in time you can provision configurations on it again through the POAP process. The process removes existing configurations and provision new configurations. You can also deploy configurations incrementally without invoking POAP.

Adding Non-Nexus Devices to External Fabrics

From Cisco Nexus Dashboard Fabric Controller Release 12.0.1a, you can add Cisco IOS-XR devices to external fabrics in managed mode as well. You can manage the following Cisco IOS-XR devices in external fabrics:

- Cisco ASR 9000 Series Routers
- Cisco NCS 5500 Series Routers, IOS XR Release 6.5.3

From Cisco Nexus Dashboard Fabric Controller Release 12.1.1e, you can also add Cisco 8000 Series Routers to external fabrics both in managed mode and monitored mode.

You can discover non-Nexus devices in an external fabric and perform the configuration compliance of these devices as well. For more information, see the [Configuration Compliance in External Fabrics](#) section.

Refer the *Cisco Compatibility Matrix* to see the non-Nexus devices supported by Cisco Nexus Dashboard Fabric Controller.

Only Cisco Nexus switches support SNMP discovery by default. Hence, configure all the non-Nexus devices before adding it to the external fabric. Configuring the non-Nexus devices includes configuring SNMP views, groups, and users. See the [Configuring Non-Nexus Devices for Discovery](#) section for more information.

Cisco CSR 1000v is discovered using SSH. Cisco CSR 1000v does not need SNMP support because it can be installed in clouds where SNMP is blocked for security reasons. See the *Connecting Cisco Data Center and a Public Cloud* chapter to see a use case to add Cisco CSR 1000v to an external fabric.

However, Cisco Nexus Dashboard Fabric Controller can only access the basic device information like system name, serial number, model, version, interfaces, up time, and so on. Cisco Nexus Dashboard Fabric Controller does not discover non-Nexus devices if the hosts are part of CDP or LLDP.

The settings that are not applicable for non-Nexus devices appear blank, even if you get many options when you right-click a non-Nexus device in the fabric topology window. You cannot add or edit interfaces for ASR 9000 Series Routers and Arista switches.

You can add IOS-XE devices like Cisco Catalyst 9000 Series switches and Cisco ASR 1000 Series Routers as well to external fabrics.

Configuration Compliance in External Fabrics

With external fabrics, any Nexus switches, Cisco IOS-XE devices, Cisco IOS XR devices, and Arista can be imported into the fabric, and there is no restriction on the type of deployment. It can be LAN Classic, VXLAN, FabricPath, vPC, HSRP, etc. When switches are imported into an external fabric, the configuration on the switches is retained so that it is non-disruptive. Only basic policies such as the switch username and mgmt0 interface are created after a switch import.

In the external fabric, for any intent that is defined in the Nexus Dashboard Fabric Controller,

configuration compliance (CC) ensures that this intent is present on the corresponding switch. If this intent is not present on the switch, CC reports an Out-of-Sync status. Additionally, there will be a Pending Config generated to push this intent to the switch to change the status to In-Sync. Any additional configuration that is on the switch but not in intent defined in Nexus Dashboard Fabric Controller, will be ignored by CC, as long as there is no conflict with anything in the intent.

When there is user-defined intent added on Nexus Dashboard Fabric Controller and the switch has additional configuration under the same top-level command, as mentioned earlier, CC will only ensure that the intent defined in Nexus Dashboard Fabric Controller is present on the switch. When this user defined intent on Nexus Dashboard Fabric Controller is deleted as a whole with the intention of removing it from the switch and the corresponding configuration exists on the switch, CC will report an Out-of-Sync status for the switch and will generate **Pending Config** to remove the config from the switch. This **Pending Config** includes the removal of the top-level command. This action leads to removal of the other out-of-band configurations made on the switch under this top-level command as well. If you choose to override this behavior, the recommendation is that, you create a freeform policy and add the relevant top-level command to the freeform policy.

Let us see this behavior with an example.

1. A **switch_freeform** policy defined by the user in Nexus Dashboard Fabric Controller and deployed to the switch.
2. Additional configuration exists under **router bgp** in **Running config** that does not exist in user-defined Nexus Dashboard Fabric Controller intent **Expected config**. Note that there is no **Pending Config** to remove the additional config that exists on the switch without a user defined intent on Nexus Dashboard Fabric Controller.
3. The **Pending Config** and the **Side-by-side Comparison** when the intent that was pushed earlier via Nexus Dashboard Fabric Controller is deleted from Nexus Dashboard Fabric Controller by deleting the switch_freeform policy that was created in the Step 1.
4. A **switch_freeform** policy with the top-level **router bgp** command needs to be created. This enables CC to generate the configuration needed to remove only the desired sub-config which was pushed from Nexus Dashboard Fabric Controller earlier.
5. The removed configuration is only the subset of the configuration that was pushed earlier from Nexus Dashboard Fabric Controller.

For interfaces on the switch in the external fabric, Nexus Dashboard Fabric Controller either manages the entire interface or does not manage it at all. CC checks interfaces in the following ways:

- For any interface, if there is a policy defined and associated with it, then this interface is considered as managed. All configurations associated with this interface must be defined in the associated interface policy. This is applicable for both logical and physical interfaces. Otherwise, CC removes any out-of-band updates made to the interface to change the status to **In-Sync**.
- Interfaces created out-of-band (applies for logical interfaces such as port-channels, sub interfaces, SVIs, loopbacks, etc.), will be discovered by Nexus Dashboard Fabric Controller as part of the regular discovery process. However, since there is no intent for these interfaces, CC will not report an **Out-of-Sync** status for these interfaces.
- For any interface, there can always be a monitor policy associated with it in Nexus Dashboard Fabric Controller. In this case, CC will ignore the interface's configuration when it reports the

In-Sync or **Out-of-Sync** config compliance status.

Special Configuration CLIs Ignored for Configuration Compliance

The following configuration CLIs are ignored during configuration compliance checks:

- Any CLI having 'username' along with 'password'
- Any CLI that starts with 'snmp-server user'

Any CLIs that match the above will not show up in pending diffs and clicking Save & Deploy in the Fabric Builder window will not push such configurations to the switch. These CLIs will not show up in the Side-by-side Comparison window also.

To deploy such configuration CLIs, perform the following procedure:

1. Select **Manage > Fabrics**.

Double click on the fabric name to view **Fabric Overview** screen.

2. On the Switches tab, double click on the switch name to view **Switch Overview** screen.

On the Policies tab, all the policies applied on the switch within the chosen fabric are listed.

3. On the Policies tab, from the **Actions** drop-down list, select **Add Policy**.
4. Add a Policy Template Instances (PTIs) with the required configuration CLIs using the **switch_freeform** template and click **Save**.
5. Select the created policy and select **Push Config** from the **Actions** drop-down list to deploy the configuration to the switch(es).

Managing Cisco IOS-XR Devices using NDFC

In general, workload requires communication with services outside of the data center domain in a data center fabric. This includes users accessing an application and services from the internet and WAN. VXLAN EVPN fabrics with border devices are considered as a handoff for north-south connectivity. These border devices are in peer with IOS-XR routers, which are the backbone routers for WAN and internet connectivity.

In DCNM Release 11.5(x), users with an admin role can control VXLAN EVPN fabrics with capabilities such as monitoring, automation, and compliance. You can only monitor the IOS-XR routers in monitored mode. Therefore, there is a requirement for a single fabric controller to manage, and automate configurations between these devices to balance and check configuration compliance for communicating between different services.

From NDFC Release 12.0.1a, users with an admin role can manage IOS-XR routers that is limited to automation and compliance checking. New templates and policies are introduced to automate and manage eBGP VRF Lite handoff between border switches and IOS-XR routers. NDFC allows you to check configuration compliance for IOS-XR devices similar to Cisco Nexus switches in external fabrics.



For all non-Nexus devices, only the message-digest algorithm (MD5) protocol is supported for Simple Network Management Protocol version 3 (SNMPv3) authentication.



Beginning with NDFC 12.2.1, you do not need to configure the Simple Network Management Protocol (SNMP) for IOS-XR discovery of switches. NDFC uses Secure Shell (SSH) for IOS-XR device discovery.

Configuring IOS-XR as an Edge Router

To extend VRF Lite from a Cisco Nexus 9000 fabric with border devices for IOS-XR as the edge router, see the *VRF Lite Between Cisco Nexus 9000 Based Border and Non-Nexus Device* section for more information.

For more information, see the video at [Managing and Configuring ASR 9000 using NDFC](#).

Configuring Non-Nexus Devices for Discovery

Before discovering any non-Nexus device in Cisco Nexus Dashboard Fabric Controller, configure it on the switch console.

Configuring IOS-XE Devices for Discovery



In case of failure or issues configuring devices contact Cisco Technical Assistance Center (TAC). Before you discover the Cisco IOS-XE devices in Nexus Dashboard Fabric Controller, perform the following steps:

1. Run the following SSH commands on the switch console.

```
switch (config)# hostname <hostname>
switch (config)# ip domain name <domain_name>
switch (config)# crypto key generate rsa
switch (config)# ip ssh time-out 90
switch (config)# ip ssh version 2
switch (config)# line vty 1 4
switch (config-line)# transport input ssh
switch (config)# username admin privilege secret <password>
switch (config)# aaa new-model
switch (config)# aaa authentication login default local
switch (config)# aaa authorization exec default local none
```

2. Before you run SNMP command on the switch, ensure that the IP addresses, username and SNMP related configurations are defined on the switch. Run the following SNMP command on the switch console.

```
aaa new-model
```

```

aaa session-id common
ip domain name cisco
username admin privilege 15 secret 0 xxxxx
snmp-server group group1 v3 auth read view1 write view1
snmp-server view view1 mib-2 included
snmp-server view view1 cisco included
snmp-server user admin group1 v3 auth md5 xxxxx priv des xxxxx
line vty 0 4
privilege level 15
transport input all
line vty 5 15
privilege level 15
transport input all
line vty 16 31
transport input ssh

```

Configuring Arista Devices for Discovery

Enable Privilege Exec mode using the following command:

```

switch> enable
switch#

switch# show running configuration | grep aaa      /* to view the authorization*/
aaa authorization exec default local

```

Run the following commands in the switch console to configure Arista devices:

```

switch# configure terminal
switch (config)# username ndfc privilege 15 role network-admin secret cisco123
snmp-server view _view_name_ SNMPv2 included
snmp-server view _view_name_ SNMPv3 included
snmp-server view _view_name_ default included
snmp-server view _view_name_ entity included
snmp-server view _view_name_ if included
snmp-server view _view_name_ iso included
snmp-server view _view_name_ lldp included
snmp-server view _view_name_ system included
snmp-server view sys-view default included
snmp-server view sys-view ifmib included
snmp-server view sys-view system included
snmp-server community private ro
snmp-server community public ro
snmp-server group _group_name_ v3 auth read _view_name_

```

```
snmp-server user username _group_name_ v3 auth md5 _password_ priv aes _password_
```



SNMP password should be same as the password for username.

You can verify the configuration by running the **show run** command, and view the SNMP view output by running the **show snmp view** command.

Show Run Command

```
switch (config)# snmp-server engineID local f5717f444ca824448b00
snmp-server view _view_name_ SNMPv2 included
snmp-server view _view_name_ SNMPv3 included
snmp-server view _view_name_ default included
snmp-server view _view_name_ entity included
snmp-server view _view_name_ if included
snmp-server view _view_name_ iso included
snmp-server view _view_name_ lldp included
snmp-server view _view_name_ system included
snmp-server view sys-view default included
snmp-server view sys-view ifmib included
snmp-server view sys-view system included
snmp-server community private ro
snmp-server community public ro
snmp-server group _group_name_ v3 auth read _view_name_
snmp-server user _user_name__group_name_ v3 localized f5717f444ca824448b00 auth
md5 be2eca3fc858b62b2128a963a2b49373 priv aes
be2eca3fc858b62b2128a963a2b49373
!
spanning-tree mode mstp
!
service unsupported-transceiver labs f5047577
!
aaa authorization exec default local
!
no aaa root
!
username admin role network-admin secret sha512
$6$5ZKs/7.k2UxrWDg0$FOkdVQsBTnOquW/9AYx36YUBSPNLFdeuPlse9XgyHSdEOYXtPyT/
0sMUYYdkMffuljgn/d9rx/Do71XSbygSn/
username cvpadmin role network-admin secret sha512
$6$fLGFj/PUcuJT436i$Sj5G5c4y9cYjl/BZswjjmZW0J4npGrGqlyG3ZFk/ULza47Kz.d31q13jXA
7iHM677gwqQbFSH2/3oQEaHRq08.
username ndfc privilege 15 role network-admin secret sha512
$6$M48PNrCdG2EITEdG$iiB880nvFQQlrWoZwOMzdt5EfkuClraNqtEMRS0TJUHNKCCQnJN.VD
```

Show SNMP View Command

```
configure terminal# show snmp view
view_name SNMPv2 - included
view_name SNMPv3 - included
view_name default - included
view_name entity - included
view_name if - included
view_name iso - included
view_name lldp - included
view_name system - included
sys-view default - included
sys-view ifmib - included
sys-view system - included
leaf3-7050sx#show snmp user
```

```
User name : _user_name_
Security model : v3
Engine ID : f5717f444ca824448b00
Authentication : MD5
Privacy : AES-128
Group : _group_name_
```

Configuring and Verifying Cisco IOS-XR Devices for Discovery

To configure IOS-XR devices, run the following commands on the switch console:

```
switch# configure terminal
switch (config)# snmp-server viewview_namecisco included
snmp-server view _view_name_ mib-2 included
snmp-server group _group_name_ v3 auth read _view_name_ write _view_name_
snmp-server user _user_name__group_name_ v3 auth md5 password priv des56 password
SystemOwner
```

Below shown example of configuring IOS-XR device on a switch.

```
RP/0/RSP0/CPU0:ios(config)#snmp-server view view_name cisco included
RP/0/RSP0/CPU0:ios(config)#snmp-server view view_namemib-2 included
RP/0/RSP0/CPU0:ios(config)#snmp-server group group_name v3 auth read view_name
write view_name
RP/0/RSP0/CPU0:ios(config)#snmp-server user user_name_group_name_ v3 auth md5
```

```
password priv des56 passwordSystemOwner
RP/0/RSP0/CPU0:ios(config)#commit Day MMM DD HH:MM:SS Timezone
```

To verify IOS-XR devices, run the following command:

```
RP/0/RSP0/CPU0:ios(config)#
RP/0/RSP0/CPU0:ios(config)#show run snmp-server Day MMM DD HH:MM:SS Timezone
snmp-server user user_name group1 v3 auth md5 encrypted 10400B0F3A4640585851 priv
des56 encrypted 000A11103B0A59555B74 SystemOwner
snmp-server view _view_name_cisco included
snmp-server view _view_name_mib-2 included
snmp-server group group_name v3 auth read view_name write view_namev3 auth read
_view_name_ write _view_name_
```

Discovering Non-Nexus Devices in an External Fabric

To add non-Nexus devices to an external fabric on the fabric topology page, perform the following steps:

Before you begin:

Ensure that the configurations are pushed for non-Nexus devices before adding them to an external fabric. You cannot push configurations in a fabric in the monitor mode.

1. Navigate to **Manage > Fabrics**.
2. Double-click on an external fabric.

The **Fabric Overview** page appears.
3. Click the **Switches** tab.
4. Choose **Add switches** in the **Actions** drop-down list.

The **Add Switches - Fabric** page appears.

The **Discover** radio button is selected in the **Switch Addition Mechanism** area.

5. Enter values for the following fields in the **Seed Switch Details** area.

Field	Description
-------	-------------

Seed IP	Enter the IP address of the switch. You can discover more than one switch by providing the IP address range. For example: 10.10.10.40-60. The switches must be properly cabled and connected to the Nexus Dashboard Fabric Controller server and the switch status must be manageable.
Authentication Protocol	Choose the authentication protocol from the drop-down list. MD5 is the default.
Device Type	- Choose IOS XE from the drop-down list for adding Cisco CSR 1000v, Cisco ASR 1000 series routers, Cisco Catalyst 8000 series switches, and Cisco Catalyst 9000 series switches. - Choose IOS XR from the drop-down list for adding ASR 9000 series routers, Cisco NCS 5500 and NCS 5001 series routers, IOS XR Release 6.5.3, or Cisco 8000 series routers.  To add Cisco IOS XR devices in managed mode, navigate to the General Parameters tab in the fabric settings and uncheck the Fabric Monitor Mode check box. - Choose Other from the drop-down list for adding non-Cisco devices, like Arista switches.
Username	Enter the username.
Password	Enter the password.
Set as individual device write credential	Specify if you want to set discovery/read credentials on individual devices.

6. Click **Discover Switches**.

The **Discovery Results** section appears with the switch details populated.



An error message appears if you try to discover a device that is already discovered.

Set the password of the device on the **LAN Credentials** page if the password is not set. To navigate to the **LAN Credentials** page from the Cisco Nexus Dashboard Fabric Controller Web UI, choose **Admin > Switch Credentials > LAN**

Credentials Management.

7. Check the check boxes next to the switches you want to add.
8. Click **Add Switches**.

The switch discovery process is initiated. The **Progress** column displays the progress.

Discovering devices takes some time. A pop-up message appears at the bottom-right about the device discovery after the discovery progress is **100%**, or **done**. For example: **<ip-address> added for discovery**.



If you see the following error message after attempting to add the switch to the fabric:

Error while creating the (Seed interface) intent for basic switch configurations. Please retry using config Save/Deploy.

This might be because the permissions were not set properly for the switch before you tried to add it to the fabric. Set the permissions for the switch using the procedures in [Configuring Non-Nexus Devices for Discovery](#), then try adding the switch to the fabric again.

9. Click **Close**.

The fabric topology page appears with the switches.

10. (Optional) Click **Refresh** on the **Topology** page to view the latest topology view.
11. (Optional) Click **Fabric Overview**.

The switches and links page appears, where you can view the scan details. The discovery status is **discovering** in red with a warning icon next to it if the discovery is in progress.

12. (Optional) View the details of the device.

After the discovery of the device:

- o The discovery status changes to **Ok** in green with a check box checked next to it.
- o The value of the device under the **Config Status** column changes to **In-Sync**.



When a switch is in **Unreachable** discovery status, the last available information of the switch is retained in other columns. For example, if the switch was in **RUNNING** tracker status before it becomes unreachable, the value under the **Config Status** column for this switch will still be **RUNNING** despite the switch being in **Unreachable** discovery status.

What to do next:

Set the appropriate role. Right-click the device, choose **Set role**.

If you added these devices under managed mode, you can add policies too.

Managing Non-Nexus Devices to External Fabrics

From Nexus Dashboard Fabric Controller 12.0.1a, IOS-XR is supported in managed mode.



Configuration compliance is enabled for IOS-XE and IOS-XR switches, similar to the way the Nexus switches are handled in External Fabric. For more information, see [Configuration Compliance in External Fabrics](#).

Nexus Dashboard Fabric Controller sends commit at the end of deployment for IOS-XR devices.

Nexus Dashboard Fabric Controller provides a few templates for IOS-XR devices. Use the **ios_xr_Ext_VRF_Lite_Jython.template** for IOS-XR switch to be an edge router to establish eBGP peering with border. This will create config for vrf, eBGP peering for the vrf and the sub-interface. Similarly, **ios_xe_Ext_VRF_Lite_Jython** can be used for IOS-XE switch to be an edge router to establish eBGP peering with border.

Creating a vPC Setup

You can create a vPC setup for a pair of switches in an external fabric. Ensure that the switches are of the same role and connected to each other.

1. Right-click one of the two designated **vPC switches** and choose **vPC Pairing**.

The **Select vPC peer** dialog box comes up. It contains a list of potential peer switches. Ensure that the **Recommended** column for the vPC peer switch is updated as **true**.



Alternatively, you can also navigate to the **Tabular view** from the **Actions** pane. Choose a switch in the **Switches** tab and click **vPC Pairing** to create, edit, or unpair a vPC pair. However, you can use this option only when you choose a Cisco Nexus switch.

2. Click the radio button next to the vPC peer switch and choose **vpc_pair** from the **vPC Pair Template** drop-down list. Only templates with the **VPC_PAIR** template sub type are listed here.

The **vPC Domain** and **vPC Peerlink** tabs appear. You must populate the fields in the tabs to create the vPC setup. The description for each field is displayed at the extreme right.

Field	Description
vPC Domain tab	Enter the vPC domain details.
vPC+	If the switch is part of a FabricPath vPC + setup, enable this check box and enter the FabricPath switch ID field.
Configure VTEPs	Check this check box to enter the source loopback IP addresses for the two vPC peer VTEPs and the loopback interface secondary IP address for NVE configuration.
NVE interface	Enter the NVE interface. vPC pairing will configure only the source loopback interface. Use the freeform interface manager for additional configuration.
NVE loopback configuration	NVE loopback configuration: Enter the IP address with the mask. vPC pairing will only configure primary and secondary IP address for loopback interface. Use the freeform interface manager for additional configuration.
vPC Peerlink tab	Enter the vPC peer-link details.
Switch Port Mode	Choose trunk or access or fabricpath . If you select trunk , then corresponding fields (Trunk Allowed VLANs and Native VLAN) are enabled. If you select access , then the Access VLAN field is enabled. If you select fabricpath , then the trunk and access port related fields are disabled.

3. Click **Save**.

The **vPC setup** is created.

To update vPC setup details, do the following:

- a. Right-click a vPC switch and choose **vPC Pairing**.

The **vPC peer** dialog box comes up.

- b. Update the field(s) as needed.

When you update a field, the **Unpair** icon changes to **Save**.

- c. Click **Save** to complete the update.

After creating a vPC pair, you can view vPC details on the **vPC Overview** page.

Undeploying a vPC Setup

1. Right-click a **vPC** switch and choose **vPC Pairing**.

The vPC peer screen comes up.

2. Click **Unpair** at the bottom right part of the screen.

The vPC pair is deleted and the fabric topology window appears.

3. Click **Deploy Config**.

4. (Optional) Click the value under the **Recalculate Config** column.

View the pending configuration in the **Config Preview** dialog box. The following configuration details are deleted on the switch when you unpair: vPC feature, vPC domain, vPC peerlink, vPC peerlink member ports, loopback secondary IPs, and host vPCs. However, the host vPCs and port channels are not removed. Delete these port channels from the **Interfaces** window if required.



Resync the fabric if it is out of sync.

When you unpair, only PTIs are deleted for following features, but the configuration is not cleared on the switch during **Deploy Config**: NVE configuration, LACP feature, fabricpath feature, nv overlay feature, loopback primary ID. In case of host vPCs, port channels and their member ports are not cleared. You can delete these port channels from the **Interfaces** window if required. You can continue using these features on the switch even after unpairing.

If you are migrating from fabricpath to VXLAN, you need to clear the configuration on the device before deploying the VXLAN configuration.

Precision Time Protocol for External Fabrics

In the Fabric settings for the **External Fabric** template, select the **Enable Precision Time Protocol (PTP)** check box to enable PTP across a fabric. When you select this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the **PTP Loopback Id** and **PTP Domain Id** fields are editable.

The PTP feature is supported with Cisco Nexus 9000 Series cloud-scale switches, with NX-OS version 7.0(3)I7(1) or later. Warnings are displayed if there are non-cloud scale devices in the fabric, and PTP is not enabled. Examples of the cloud-scale devices are Cisco Nexus 93180YC-EX, Cisco Nexus 93180YC-FX, Cisco Nexus 93240YC-FX2, and Cisco Nexus 93360YC-FX2 switches. For more information, refer to <https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html>.



PTP global configuration is supported with Cisco Nexus 3000 Series switches; however, PTP and TTAG configurations are not supported.

For more information, see the *Configuring PTP* chapter in *Cisco Nexus 9000 Series NX-OS System Management Configuration Guide* and *Cisco Nexus Insights for Cisco User Guide*.

For External fabric deployments, you have to enable PTP globally, and also enable PTP on core-facing interfaces. The interfaces could be configured to the external PTP server like a VM or Linux-based machine. Therefore, the interface should be edited to have a connection with the grandmaster clock. For PTP and TTAG configurations to be operational on External fabrics, you must sync up of Switch Configs to Nexus Dashboard Fabric Controller using the **host_port_resync** policy. For more information, see the section "Out-of-Band Switch Interface Configurations" in [Add Interfaces for LAN Operational Mode](#).

It is recommended that the grandmaster clock should be configured outside of Data Center VXLAN EVPN and it is IP reachable. The interfaces toward the grandmaster clock need to be enabled with PTP via the interface freeform config.

All core-facing interfaces are auto-enabled with the PTP configuration after you click **Deploy Config**. This action ensures that all devices are PTP synced to the grandmaster clock. Additionally, for any interfaces that are not core-facing, such as interfaces on the border devices and leafs that are connected to hosts, firewalls, service-nodes, or other routers, the TTAG related CLI must be added. The TTAG is added for all traffic entering the VXLAN EVPN fabric and the TTAG must be stripped when traffic is exiting this fabric.

Here is the sample PTP configuration:

```
feature ptp
```

```
ptp source 100.100.100.10 -> IP address of the loopback interface (loopback0)  
that is already created, or user-created loopback interface in the fabric settings
```

```
ptp domain 1 -> PTP domain ID specified in fabric settings
```

```
interface Ethernet1/59 -> Core facing interface
```

ptp

interface Ethernet1/50 -> Host facing interface

ttag

ttag-strip

The following guidelines are applicable for PTP:

- The PTP feature can be enabled in a fabric when all the switches in the fabric have Cisco NX-OS Release 7.0(3)I7(1) or a higher version. Otherwise, the following error message is displayed:

PTP feature can be enabled in the fabric, when all the switches have NX-OS Release 7.0(3)I7(1) or higher version. Please upgrade switches to NX-OS Release 7.0(3)I7(1) or higher version to enable PTP in this fabric.

- For hardware telemetry support in NIR, the PTP configuration is a prerequisite.
- If you are adding a non-cloud scale device to an existing fabric which contains PTP configuration, the following warning is displayed:

TTAG is enabled fabric wide, when all devices are cloud-scale switches so it cannot be enabled for newly added non cloud-scale device(s).

- If a fabric contains both cloud-scale and non-cloud scale devices, the following warning is displayed when you try to enable PTP:

TTAG is enabled fabric wide when all devices are cloud-scale switches and is not enabled due to non cloud-scale device(s).

- TTAG configuration is generated for all the devices if host configuration sync up is performed on all the devices. TTAG configuration will not be generated for any newly added devices if host configuration sync up is not performed on all newly added devices.

If the configuration is not synced, the following warning is displayed:

TTAG on interfaces with PTP feature can only be configured for cloud-scale devices. It will not be enabled on any newly added switches due to the presence of non cloud-scale devices.

- PTP and TTAG configurations are deployed on host interfaces.
- PTP and TTAG Configurations are supported between switches in the same fabric (intra-fabric links). PTP is created for inter-fabric links, and TTAG is created for the inter-fabric link if the other fabric (Switch) is not managed by Nexus Dashboard Fabric Controller. Inter-fabric links do not support PTP or TTAG configurations if both fabrics are managed by Nexus Dashboard Fabric Controller.

- TTAG configuration is configured by default after the breakout. After the links are discovered and connected post breakout, perform **Deploy Config** to generate the correct configuration based on the type of port (host, intra-fabric link, or inter fabric link).

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)

© 2017–2025 Cisco Systems, Inc. All rights reserved.