



Campus VXLAN EVPN, Release 12.1.3

Table of Contents

New and Changed Information	1
About Campus VXLAN EVPN Fabric	2
Creating Campus VXLAN EVPN Fabric	3
General Parameters	3
Replication	5
Protocols	5
Advanced	5
Resources	7
Bootstrap	9
Configuration Backup	11
Border Gateway	12
Adding Cisco Catalyst 9000 Series Switches and Nexus 9000 Series Switches to a Campus VXLAN EVPN Fabric	15
Recalculating and Deploying Configurations	18
Creating VRFs in Campus VXLAN EVPN Fabric	19
Attaching VRFs to Switches in Campus VXLAN EVPN Fabrics	21
Creating and Deploying Networks in Campus VXLAN EVPN Fabrics	22
Creating Networks for Campus VXLAN EVPN Fabrics	22
Attaching Networks in Campus VXLAN EVPN Fabric	24
Deploying Networks in Campus VXLAN EVPN Fabrics	24
Creating DCI Links for Switches in Campus VXLAN EVPN Fabric	25
Copyright	27

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 12.1.3	Reorganized content	Content within this document was originally provided in the <i>Cisco NDFC-Fabric Controller Configuration Guide</i> or the <i>Cisco NDFC-SAN Controller Configuration Guide</i> . Beginning with release 12.1.3, this content is now provided solely in this document and is no longer provided in those documents.

About Campus VXLAN EVPN Fabric

NDFC supports Campus VXLAN EVPN fabric type to automate and manage Enterprise Campus VXLAN BGP EVPN networks based on Catalyst 9000 series switches. Optionally, admins can integrate a Nexus 9000 switch with Border Gateway functionality to interconnect with remote Data Centers and Campus for VXLAN EVPN Multi-Site Layer 2 and Layer 3 extensions.

This document describes how to create a Campus VXLAN EVPN fabric with Cisco Catalyst 9000 Series Switches and Nexus 9000 Series Switches using the **Campus VXLAN EVPN** fabric template. This fabric supports OSPF as the underlay protocol and BGP EVPN as the overlay protocol. Using this fabric template, Cisco Nexus Dashboard Fabric Controller manages all the configurations of a VXLAN EVPN fabric consisting of Cisco Catalyst 9000 IOS XE and Nexus 9000 NX-OS switches. Backing up and restoring this fabric is similar to **Data Center VXLAN EVPN** backup and restore.

Starting with Release 12.1.3, Cisco Nexus Dashboard Fabric Controller provides support for the following features with Campus VXLAN EVPN fabrics:

- Tenant Routed Multicast (TRM)
- Zero-Touch Provisioning (ZTP) on Cisco Catalyst switches through Plug and Play (PnP)
- Deploying Campus VXLAN EVPN fabric as a child fabric in VXLAN EVPN Multi-Site.

Guidelines for Configuring Campus VXLAN EVPN Fabrics

- Provides support for EVPN VXLAN Distributed Anycast Gateway when each SVI is configured with the same Anycast Gateway MAC.
- Provides support for Cisco Catalyst switches with Stackwise or Stackwise Virtual.
- Provides support for spine, leaf, and border roles on Cisco Catalyst switches. Whereas, Cisco Nexus 9000 series switches support border gateway, border gateway spine and border gateway super spine roles.
- Does not support Brownfield deployments.
- Does not support IPv6 Underlay and Anycast RP.
- Does not support ISIS, ingress replication, unnumbered intra-fabric link, and 4 bytes BGP ASN.



For information about configuration compliance, see [Configuration Compliance in External Fabrics](#).

Creating Campus VXLAN EVPN Fabric

Perform the following steps to create the **Campus VXLAN EVPN** fabric for Cisco Catalyst 9000 Series Switches and Nexus 9000 Series Switches:

1. In Cisco Nexus Dashboard Fabric Controller, navigate to **LAN > Fabrics**.
2. Choose **Actions > Create Fabric**.

The **Create Fabric** window appears.

3. Enter a unique name for the fabric in the **Fabric Name** field and click **Choose Fabric**.

A list of all available fabric templates is listed.

4. From the available list of fabric templates, choose the **Campus VXLAN EVPN fabric** template, then click **Select**.
5. Enter the necessary field values to create a fabric.

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](#)
- [Replication](#)
- [Protocols](#)
- [Advanced](#)
- [Resources](#)
- [Bootstrap](#)
- [Configuration Backup](#)
- [Border Gateway](#)

6. When you have completed the necessary configurations, click **Save**.
 - Click on the fabric to display a summary in the slide-out pane.
 - Double-click on the fabric name to open the **Fabric Overview** page.

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	Specifies the unique Autonomous System (AS) number used to identify the fabric. The range is 1-65535.
Underlay Subnet IP Mask	Specifies the subnet mask for the fabric interface IP addresses.
Link-State Routing Protocol	Specifies the supported routing protocol which is OSPF.

Field	Description
Route-Reflectors	<p>The number of spine switches that are used as route reflectors for transporting BGP traffic. Choose option 2 or 4 from the drop-down list. The default value is 2.</p> <p>To deploy spine devices as RRs, Cisco Nexus Dashboard Fabric Controller sorts the spine devices based on their serial numbers and designates two or four spine devices as RRs. If you add more spine devices, existing RR configuration won't change.</p> <p>Increasing the count - You can increase the route reflectors from two to four at any point in time. Configurations are automatically generated on the other two spine devices designated as RRs.</p> <p>Decreasing the count - When you reduce four route reflectors to two, remove the required route reflector devices from the fabric.</p> <p>Follow these steps to reduce the count from 4 to 2.</p> <ol style="list-style-type: none"> 1. Change the value in the drop-down list to 2. 2. Identify the spine switches designated as route reflectors. An instance of the rr_state policy is applied on the spine switch if it's a route reflector. 3. Delete the spine devices that are not required from the fabric. <p>If you delete existing RR devices, the next available spine switch is selected as the replacement RR.</p> <ol style="list-style-type: none"> 4. Click Deploy Config in the fabric topology window. <p>You can preselect RRs and RPs before performing the first Save & Deploy operation.</p> <p>For more information, see [Preselecting Switches as Route-Reflectors and Rendezvous-Points].</p>
Anycast Gateway MAC	Specifies the shared MAC address for the leaf switches.
Enable Performance Monitoring	<p>Enables performance monitoring on the switches.</p> <p>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.</p>

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

Replication

The fields in the **Replication** tab are described in the following table. All the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Replication Mode	Specifies the mode of replication used in the fabric for BUM (Broadcast, Unknown Unicast, Multicast) traffic. Multicast is selected by default.
Multicast Subnet Group	<p>Specifies the IP address prefix used for multicast communication. A unique IP address is allocated from this group for each overlay network.</p> <p>The replication mode change is not allowed if a policy template instance is created for the current mode. For example, if a multicast related policy is created and deployed, you cannot change the mode to Ingress.</p>
Enable Tenant Routed Multicast (TRM)	Enables Tenant Routed Multicast (TRM) that allows overlay multicast traffic support over EVPN/MVPN in the VXLAN EVPN fabric.
Rendezvous-Points	Specifies the number of spine switches acting as rendezvous points.
Underlay RP Loopback Id	Specifies the loopback ID used for the RP, for multicast protocol peering purposes in the fabric underlay. The default is 254.

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

Protocols

The fields in the **Protocols** tab are described in the following table. All the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Underlay Routing Loopback Id	Specifies the loopback interface ID. By default, value 0 is populated as loopback0 that is normally used for fabric underlay IGP peering purposes.
Underlay VTEP Loopback Id	Specifies the loopback interface ID. By default, value 1 is populated as loopback1 used as the VTEP address.
OSPF Process Id	Specifies the OSPF process tag.
OSPF Area Id	Specifies the OSPF unique 32-bit area ID denoted in dotted decimal format.

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

Advanced

The fields in the **Advanced** tab are described in the following table. All the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields

if needed.

Field	Description
VRF Template	Specifies the VRF template for creating VRFs. By default, the system uses the pre-defined Default_VRF_Universal template for overlay configuration for leaf switches.
Network Template	Specifies the network template for creating networks. By default, the system uses the pre-defined Default_Network_Universal template for leaf switches.
VRF Extension Template	Specifies the VRF extension template for enabling VRF extension to other fabrics. By default, the system uses the pre-defined Default_VRF_Extension_Universal template for border switches.
Network Extension Template	Specifies the network extension template for extending the network to other fabrics. By default, the system uses the pre-defined Default_Network_Extension_Universal template for border switches.
Intra Fabric Interface MTU	<p>Specifies the MTU for the intra fabric interface. The value must be an even number.</p> <p>The valid values range from 576 to 9216. This is a mandatory field.</p>
Layer 2 Host Interface MTU	<p>Specifies the MTU for the layer 2 host interface. This value must be an even number. The valid values range from 1500 to 9216.</p> <p>The default MTU for both fabric interface and host interface is 9198. The default MTU for IOS XE is 1500.</p>
IOS XE System MTU	Specifies the MTU for an IOS XE device. This value must be an even number. The valid values range from 1500 to 9198.
Enable Tenant DHCP	<p>Check the check box to enable DHCP and associated configurations globally on all the switches in the fabric. This is a prerequisite for supporting DHCP for overlay networks that are part of the tenant VRFs.</p> <div>  <p>Ensure that Enable Tenant DHCP is enabled before enabling DHCP-related parameters in the overlay profiles.</p> </div>
VTP Mode	By default, the VTP mode is Off . Transparent mode allows you to relay all VTP protocol packets that it receives on a trunk port to all other trunk ports.
Enable NDFC as Trap Host	Allows you to configure Cisco Nexus Dashboard Fabric Controller as an SNMP trap destination.


Field	Description
Enable Overlay Template Conversion	<p>Allows you to convert all the existing VRFs and networks to use the default templates. In existing deployments using the IOS_XE_VRF and IOS_XE_Network templates, enabling this field converts the templates to use the Default_VRF_Universal, Default_Network_Universal, Default_VRF_Extension_Universal and Default_Network_Extension_Universal templates after performing a Recalculate and Deploy.</p> <p>When adding child fabrics to a VXLAN EVPN Multi-Site, Campus VXLAN EVPN fabric does not allow you to add fabrics which are configured with IOS_XE_VRF and IOS_XE_Network as the templates for existing VRFs and networks. Enable the Enable Overlay Template Conversion field to convert the existing VRFs and networks using IOS_XE_VRF and IOS_XE_Network templates to the Default_VRF and Default_Network templates.</p>
Leaf Freeform Config	Configures additional CLIs for all the Cisco Catalyst leaf switches in the fabric.
Spine Freeform Config	Configures additional CLIs for all the Cisco Catalyst spine switches in the fabric.
Intra-fabric Links Additional Config	Configures additional CLIs for all the intra fabric links.

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

Resources

The fields in the **Resources** tab are described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Underlay Routing Loopback IP Range	Specifies the loopback IPv4 addresses for protocol peering.
Underlay VTEP Loopback IP Range	Specifies the loopback IP address range for VTEPs.
Underlay RP Loopback IP Range	Specifies anycast or phantom RP IP address range.
Underlay Subnet IP Range	Specifies the IP addresses for underlay P2P routing traffic between interfaces.
Layer 2 VXLAN VNI Range	Specify the VXLAN VNI IDs for the fabric.
Layer 3 VXLAN VNI Range	Specify the VXLAN VNI IDs for the fabric.


Field	Description
Network VLAN Range	VLAN range for the per switch overlay network (min:2, max:4094).
VRF VLAN Range	VLAN range for the per switch overlay Layer 3 VRF (min:2, max:4094).
Subinterface Range Dot1q	Specifies the subinterface range when L3 sub-interfaces are used.
VRF Lite Deployment	<p>Specifies the VRF Lite method for extending inter fabric connections.</p> <p>The VRF Lite Subnet IP Range field specifies resources reserved for IP address used for VRF Lite when VRF Lite IFCs are auto-created. If you select Back2Back&ToExternal, then VRF Lite IFCs are auto-created.</p>
Auto Deploy for Peer	<p>This check box is applicable for VRF Lite deployment. When you select this checkbox, auto-created VRF Lite IFCs will have the Auto Generate Configuration for Peer field in the VRF Lite tab set, if the peer is a Cisco device.</p> <p>To access VRF Lite IFC configuration, navigate to the Links tab, select the link, and then choose Actions > Edit.</p> <p>You can check or uncheck the check box when the VRF Lite Deployment field is not set to Manual. This configuration only affects the new auto-created IFCs and does not affect the existing IFCs. You can edit an auto-created IFC and check or uncheck the Auto Generate Configuration for Peer field. This setting takes priority always.</p>
Auto Deploy Default VRF	<p>When you select this check box, the Auto Generate Configuration on default VRF field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the VRF Lite Deployment field is not set to Manual. The Auto Generate Configuration on default VRF field when set, automatically configures the physical interface for the border device, and establishes an eBGP connection between the border device and the edge device or another border device in a different VXLAN EVPN fabric.</p>
Auto Deploy Default VRF for Peer	<p>When you select this check box, the Auto Generate Configuration for NX-OS Peer on default VRF field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the VRF Lite Deployment field is not set to Manual. The Auto Generate Configuration for NX-OS Peer on default VRF field when set, automatically configures the physical interface and the eBGP commands for the peer NX-OS and IOS XE switches.</p> <div>  <p>To access the Auto Generate Configuration on default VRF and Auto Generate Configuration for NX-OS Peer on default VRF fields for an IFC link, navigate to the Links tab, select the link and choose Actions > Edit.</p> </div>
Redistribute BGP Route-map Name	Route Map used to redistribute BGP routes to IGP in default VRF for auto-created VRF Lite IFC links.

Field	Description
VRF Lite Subnet IP Range	<p>These fields are prefilled with the DCI subnet details. Update the fields as needed.</p> <p>The values shown on the page are automatically generated. If you want to update the IP address ranges, VXLAN Layer 2/Layer 3 network ID ranges or the VRF/network VLAN ranges, ensure that each fabric has its own unique range and is distinct from any underlay range to avoid possible duplication. You should only update one range of values at a time.</p>
VRF Lite Subnet Mask	<p>If you want to update more than one range of values, do it in separate instances. For example, if you want to update Layer 2 and Layer 3 ranges, you should do the following.</p> <ol style="list-style-type: none"> 1. Update the Layer 2 range and click Save. 2. Click the Edit Fabric option again, update the Layer 3 range, and click Save.
Auto Allocation of Unique IP on VRF Extension over VRF Lite IFC	When enabled, IP prefix allocated to the VRF Lite IFC is not reused on VRF extension over VRF Lite IFC. Instead, unique IP Subnet is allocated for each VRF extension over VRF Lite IFC.
Per VRF Loopback Per VTEP Auto-Provisioning	Enables you to auto provision a loopback on a VTEP on VRF attachment.
Per VRF Loopback Pool for VTEP IP	Indicates the prefix pool to assign IP addresses to loopbacks on VTEPs on a per VRF basis.

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

Bootstrap

The fields in the **Bootstrap** tab are described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Enable Bootstrap	<p>Select this check box to enable the bootstrap feature. Bootstrap allows easy day-0 import and bring-up of new devices into an existing fabric. Bootstrap leverages POAP for NX-OS and PnP for IOS XE.</p> <div>  <p>For NX-OS switches, POAP will not work if you have set Bootstrap Script Download Protocol (in the Server Settings for LAN) as https.</p> </div>

Field	Description
Enable Local DHCP Server	<p>Select this check box to initiate enabling of automatic IP address assignment through a local DHCP server. When you select this check box, the DHCP Scope Start Address and DHCP Scope End Address fields become editable.</p> <p>If you want to configure a remote or external DHCP server for automatic IP address assignment, enter details about the external DHCP server in the Switch Mgmt Default Gateway and Switch Mgmt IP Subnet Prefix fields.</p>
DHCP Version	<p>Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the Switch Mgmt IPv6 Subnet Prefix field is disabled. If you select DHCPv6, the Switch Mgmt IP Subnet Prefix is disabled.</p> <div>  <p>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 L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.</p> </div>
Domain name	Specifies the domain name of the DHCP server.
DHCP Scope Start Address and DHCP Scope End Address	Specifies the first and the 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	<p>Specifies the prefix for the management interface on the switch. The prefix should be between 8 and 30.</p> <p><i>DHCP scope and management default gateway IP address specification:</i> 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.</p>
Switch Mgmt IPv6 Subnet Prefix	Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be between 64 and 126. This field is editable if you enable IPv6 for DHCP.

Field	Description
Bootstrap Freeform Config (IOS-XE)	<p>(Optional) Enter additional commands for IOS XE switches, as needed. For example, if you require some additional configurations to be pushed to the device and be available post device bootstrap, they can be captured in this field, to save the desired intent. After the devices boot up, they will contain the configuration defined in the Bootstrap Freeform Config field.</p> <p>Copy-paste the running-config to a freeform config field with correct indentation, as seen in the running configuration on the IOS XE switches. The freeform config must match the running config. For more information, see Enabling Freeform Configurations on Fabric Switches.</p>
Bootstrap Freeform Config (NXOS)	<p>(Optional) Enter additional commands for NX-OS switches, as needed. For example, if you require some additional configurations to be pushed to the device and be available post device bootstrap, they can be captured in this field, to save the desired intent. After the devices boot up, they will contain the configuration defined in the Bootstrap Freeform Config field.</p> <p>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.</p>
DHCPv4 Multi Subnet Scope	<p>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:</p> <p>DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix</p> <p>For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24</p>

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

Configuration Backup

The fields in the **Configuration Backup** tab is described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields, if needed.

Field	Description
Hourly Fabric Backup	Select the check box to enable an hourly backup of fabric configurations and the intent. The hourly backups are triggered during the first 10 minutes of the hour.

Field	Description
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	<p>Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the Scheduled Fabric Backup check box.</p> <p>Select both the check boxes to enable both back up processes. The backup process is initiated after you click Save.</p> <p>The scheduled backups are triggered exactly at the time you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status.</p> <p>The number of fabric backups that will be retained on NDFC is decided by the Settings > Server Settings > LAN Fabric > Maximum Backups per Fabric.</p> <p>The number of archived files that can be retained is set in the # Number of archived files per device to be retained: field in the Server Properties window.</p> <p>Note: To trigger an immediate backup, do the following:</p> <ol style="list-style-type: none"> 1. Choose LAN > Topology. 2. Click within the specific fabric box. The fabric topology screen comes up. 3. Right-click on a switch within the fabric, then select Preview Config. 4. In the Preview Config window for this fabric, click Re-Sync All. <p>You can also initiate the fabric backup in the fabric topology window. Click Backup Now in the Actions pane.</p>

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

Border Gateway

The **Border Gateway** tab is applicable only to Cisco Nexus 9000 switches. The fields in the tab are described in the following table. Most of the fields are automatically populated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Site Id	Specifies the ID for this fabric when you are moving this fabric within a VXLAN EVPN Multi-Site. The site ID is mandatory for a member fabric to be a part of a VXLAN EVPN Multi-Site. Each member fabric of a VXLAN EVPN Multi-Site has a unique site ID for identification.

Field	Description
Anycast Border Gateway advertise-pip	Advertises Anycast Border Gateway PIP as VTEP.
vPC Peer Link VLAN	Specifies the VLAN used for the vPC peer link SVI. The vPC fields become active only if the switch role is border gateway.
Make vPC Peer Link VLAN as Native VLAN	Enables vPC peer link VLAN as Native VLAN.
vPC Peer Keep Alive option	<p>Allows you to configure routed links between vPC peers using management or loopback interfaces.</p> <p>To use IP addresses assigned to the management port and the management VRF, choose management. To use IP addresses assigned to loopback interfaces and a non-management VRF, choose underlay routing loopback with IPv6 address for PKA. Both the options are supported for IPv6 underlay.</p>
vPC Auto Recovery Time (In Seconds)	Specifies the vPC auto recovery time-out period in seconds.
vPC Delay Restore Time (In Seconds)	Specifies the vPC delay restore period in seconds.
vPC Peer Link Port Channel ID	Specifies the Port Channel ID for a vPC Peer Link. By default, the value in this field is 500.
vPC IPv6 ND Synchronize	Enables IPv6 Neighbor Discovery synchronization between vPC switches. The check box is enabled by default.
vPC advertise-pip	Select the check box to enable the Advertise PIP feature. You can enable the advertise PIP feature also on a specific vPC.
vPC Domain Id Range	Specifies the vPC Domain Id range to use for new pairings.
Enable NX-API	Enables NX-API on HTTPS. This check box is checked by default.
NX-API HTTPS Port Number	Specifies the port on which NX-API is enabled. By default, NX-API is enabled on HTTPS port 443.
Enable HTTP NX-API	Enables NX-API to use HTTP connections. This option is enabled by default. However, it is recommended to use HTTPS for secure communication.
NX-API HTTP Port Number	Specifies the port on which NX-API is enabled. By default, NX-API is enabled on HTTP port 80.
Enable TCAM Allocation	Automatically generates TCAM commands for VXLAN and vPC Fabric Peering, when enabled.
Nexus Border Gateway Freeform Config	Allows you to configure additional CLIs for all the border gateway switches.
Nexus Intra-fabric Links Additional Config	Allows you to configure additional CLIs for all the intra fabric links.

Field		Description
Greenfield Option	Cleanup	Enables cleaning up the switches imported into Nexus Dashboard Fabric Controller with Preserve-Config=No, without a switch reload. This option is typically recommended only for the fabric environments with Cisco Nexus 9000v Switches to improve the switch clean up time.

What's next: Complete the configurations in another tabs if necessary or click **Save** when you have completed the necessary configurations for this fabric.

Adding Cisco Catalyst 9000 Series Switches and Nexus 9000 Series Switches to a Campus VXLAN EVPN Fabric

Cisco Catalyst 9000 series switches and Nexus 9000 Series Switches are discovered using SSH. Before adding the switches to the fabric, it is necessary that you configure the switches for SSH discovery as described in the [\[Before You Begin\]](#) section.

Choose one of the following navigation paths to add switches to the fabric.

- Choose **LAN > Fabrics**. Select the required Campus VXLAN EVPN fabric from the list and choose **Actions > Add Switches**.
- Choose **LAN > Fabrics**. Select the required Campus VXLAN EVPN fabric from the list. Go to the **Switches** tab and then choose **Actions > Add Switches**.
- Choose **LAN > Switches** and then choose **Actions > Add Switches**. Click **Choose Fabric**, select the required Campus VXLAN EVPN fabric, and then click **Select**.

Before you begin

- Set the default credentials for the device in the **LAN Credentials Management** window if not already set. To navigate to the **LAN Credentials Management** window, from Cisco Nexus Dashboard Fabric Controller choose **Settings > LAN Credentials Management**.
- For StackWise and StackWise Virtual switches, configure the StackWise/StackWise Virtual settings before adding them to the fabric.
- Run the following SSH commands on the Cisco Catalyst 9000 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 15 secret <password>
switch (config)# aaa new-model
switch (config)# aaa authentication login default local
switch (config)# aaa authorization exec default local none
```

Enter values for the following fields:

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

Seed IP	Enter the IP address of the switch in one of the following formats - "2.2.2.20" or "10.10.10.40-60" or "2.2.2.20, 2.2.2.21". You can import more than one switch by providing the IP address range. The switches must be properly cabled and reachable from Cisco Nexus Dashboard Fabric Controller.
Authentication Protocol	Choose the authentication protocol from the drop-down list.
Device Type	Choose IOS XE or NX-OS from the drop-down list. If you select IOS XE , the CAT9K radio button appears which is selected by default.
Username	Enter the username for the switch.
Password	Enter the password for the switch.
Set as individual device write credential	Check the checkbox to set the discovery/read credentials as LAN/Write credentials for individual devices.



You can change the Discover and LAN credentials only after discovering the switch.

Perform the following steps to add switches to Campus VXLAN EVPN fabric:

1. Click **Discover Switches**.

The switch details are populated.

Cisco Nexus Dashboard Fabric Controller supports the import of Cisco Catalyst 9500 Switches running in StackWise Virtual. The StackWise Virtual configuration to form a pair of Cisco Catalyst 9500 Switches into a virtual switch has to be in place before the import.

For more information on how to configure StackWise Virtual, see the [Configuring Cisco StackWise Virtual](#) chapter in the *High Availability Configuration Guide (Catalyst 9500 Switches)* for the required release.

2. Check the check boxes next to the switches you want to import.

You can import only switches with the **manageable** status.



Note that the existing configuration on the switches will be erased after adding the switches to the Campus VXLAN EVPN fabric.

3. Click **Add Switches**.

The switch discovery process is initiated, and the discovery status is updated under the **Discovery Status** column in the **Switches** tab.

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

After the discovery of the device, the discovery status changes to **ok** in green.

What to do next:

1. Set the appropriate role. The supported roles are:

Cisco Catalyst 9000 Series Switches	Nexus 9000 Series Switches
Leaf	Border gateway
Spine	Border gateway spine
Border	Border gateway super spine

To set the role, select the switch and choose **Actions > Set role**. Select a role and click **Select**.

After discovering the switches, Nexus Dashboard Fabric Controller usually assigns **Leaf** as the default role for Cisco Catalyst 9000 Series switches and Border gateway for Nexus 9000 Series switches.

Optionally, you can form a VPC pair if there are Cisco Nexus 9000 switches with border gateway role in the fabric.



After setting the switch role, if the switch credentials and the default credentials do not match, the **Discovery Status** or the **Mode** column for the switches displays an error. In such case, set the LAN credentials for the switch. For more information, see section "LAN Credentials Management" in [Overview and Initial Setup of Cisco NDFC LAN](#).

2. Recalculate the configurations and deploy the configurations to the switches. Proceed to the next section for the steps to perform recalculate and deploy.

Recalculating and Deploying Configurations

To recalculate and deploy the configurations to the switch(es) in the Campus VXLAN EVPN fabric, perform the following steps to recalculate configurations:

Before you begin:

Set the role of the switch(es) in the fabric.

1. In Cisco Nexus Dashboard Fabric Controller, navigate to **LAN > Fabrics**.
2. Double-click the fabric name to open the **Fabric Overview** window.
3. Choose **Actions > Recalculate and Deploy**.

Recalculation of configurations starts on the switch(es).

Creating VRFs in Campus VXLAN EVPN Fabric

Perform the following steps to create tenant VRFs in Campus VXLAN EVPN fabric:

1. In Cisco Nexus Dashboard Fabric Controller, choose **LAN > Fabrics**.
2. From the list of available fabrics, double-click the Campus VXLAN EVPN fabric that you have created in the previous step.

The **Fabric Overview** page appears.

3. Navigate to the **VRFs** tab and choose **Actions > Create**.

The **Create VRF** window appears.

4. Enter the required details in the mandatory fields. Some of the fields are autopopulated with default values. You can make changes, as required.

- o **VRF Name** - Accept the default value or enter a name for VRF. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).
- o **VRF ID** - Accept the default or enter an ID for the VRF.
- o **VLAN ID** - Specifies the corresponding tenant VLAN ID for the network. Enter an ID for the VLAN. If you want to propose a new VLAN for the network, click **Propose Vlan**.
- o **VRF Template** - Accept the default autopopulated template or choose another template from the list.

The default template is the template name specified in the **VRF Template** field in the **Advanced** tab on the **Create Fabric** window.

- o **VRF Extension Template** - Accept the default autopopulated template or choose another template from the list.

The default template is the template name specified in the **VRF Extension Template** field in the **Advanced** tab on the **Create Fabric** window.

5. Configure the following fields under **General Parameters**, as needed.

- o **VRF VLAN Name** - Enter the VLAN name for the VRF.
- o **VRF Description** - Enter a description for the VRF.
- o **VRF Interface Description** - Enter a description for the VRF interface.

6. Click the **Advanced** tab to optionally specify the advanced profile settings.

Some of the fields in the **Advanced** tab are mentioned here. For more information about configuring the fields in the **Advanced** tab, see "Creating VRF" section in the [About Fabric Overview for LAN Operational Mode Setups](#).

- o **VRF Interface MTU** - Specifies VRF interface MTU.
- o **Loopback Routing Tag** - If a VLAN is associated with multiple subnets, then this tag is associated with the IP prefix of each subnet. Note that this routing tag is associated with overlay network creation also.

- **Redistribute Direct Route Map** - Specifies the redistribute direct route map name.
- **Max BGP Paths** - Specifies the maximum BGP paths. The valid value range is 1-64 for NX-OS and 1-32 for IOS XE.
- **Max iBGP Paths** - Specifies the maximum iBGP paths. The valid value range is 1-64 for NX-OS and 1-32 for IOS XE.
- **Advertise Host Routes** - Check the check box to control advertisement of /32 and /128 routes to Edge routers.
- **Advertise Default Route** - Check the check box to control advertisement of default route internally.
- **Config Static 0/0 Route** - Check the check box to control configuration of static default route.

7. Click **Create** to create the VRF or click **Cancel** to discard the VRF.

A message appears indicating that the VRF is created.

The new VRF appears on the **VRFs** tab. The status displays **NA** as the VRF is created but not yet deployed. Now that the VRF is created, you can create and deploy networks on the devices in the fabric.

What to do next

1. Attach the VRF.
2. Create a loopback interface and select the **VRF_LITE** extension.

Attaching VRFs to Switches in Campus VXLAN EVPN Fabrics

To attach the VRFs and VRF Lite extensions to the switches in the Campus VXLAN EVPN fabric:

1. On the **VRFs** tab in the **Fabric Overview** window, double-click the VRF that you created in the previous section.

The **VRF Overview** page opens.

2. Go to the **VRF Attachments** tab and choose the VRF corresponding to the switch by checking the check box next to it.
3. Choose **Actions > Edit**.

The **Edit VRF Attachment** page opens.

4. Toggle the knob to **Attach** and click **Save**.

Similarly, you can create a loopback interface, and select the **VRF_LITE** extension.

For more information about attaching and detaching VRFs, see the section "VRF Attachments" in [About Fabric Overview for LAN Operational Mode Setups](#).

What to do next

Deploy the configurations as follows:

1. Click **Actions** in **Fabric Overview**.
2. Choose **Deploy config to switches**.
3. Click **Deploy** after the configuration preview is complete.
4. Click **Close** after the deployment is complete.

Creating and Deploying Networks in Campus VXLAN EVPN Fabrics

The next step is to create and deploy networks in Campus VXLAN EVPN Fabrics.

Creating Networks for Campus VXLAN EVPN Fabrics

To create network for Campus VXLAN EVPN fabric from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. In Cisco Nexus Dashboard Fabric Controller, choose **LAN > Fabrics**.
2. From the list of available fabrics, double-click the Campus VXLAN EVPN fabric that you have created in the previous step.

The **Fabric Overview** page appears.

3. Navigate to the **Networks** tab and choose **Actions > Create**.

The **Create Network** window appears.

4. Enter the required details in the mandatory fields. Some of the fields are autopopulated with default values. You can make changes, as required.

The fields in the **Create Network** window are:

- **Network Name** - Specifies the name of the network. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-).
- **Layer 2 Only** - Enables you to create a Layer 2 only network.
- **VRF Name** - Allows you to select the VRF that you have created for the fabric.

When no VRF is created, this field appears as blank. If you want to create a new VRF, click **Create VRF**. The VRF name should not contain any white spaces or special characters except underscore (_), hyphen (-), and colon (:).

- **VLAN ID** - Specifies the corresponding tenant VLAN ID for the network. If you want to propose a new VLAN for the network, click **Propose VLAN**.
- **Network Template** - Auto-populates the universal template. This is only applicable for leaf switches.
- **Network Extension Template** - Auto-populates the universal extension template. This allows you to extend this network to another fabric. The VRF Lite extension is supported. The template is applicable for border leaf switches.
- **Generate Multicast IP** - If you want to generate a new multicast group address and override the default value, click **Generate Multicast IP**.

5. Configure the following fields in the **General Parameters** tab:



If the network is a non Layer 2 network, then it is mandatory to provide the gateway IP address.

- o **IPv4 Gateway/NetMask** - Specifies the IPv4 address with subnet.

Specify the anycast gateway IP address for transporting the L3 traffic from a server belonging to MyNetwork_30000 and a server from another virtual network. The anycast gateway IP address is the same for MyNetwork_30000 on all switches of the fabric that have a network.



If the same IP address is configured in the **IPv4 Gateway** and **IPv4 Secondary GW1** or **IPv4 Secondary GW2** fields of the network template, Nexus Dashboard Fabric Controller does not show an error, and you will be able to save this configuration. However, after the network configuration is pushed to the switch, it would result in a failure as the configuration is not allowed by the switch.

- o **IPv6 Gateway/Prefix List** - Specifies one or more IPv6 addresses with subnets.
- o **Vlan Name** - Enter a name for the VLAN.
- o **Interface Description** - Enter a description for the interface. This interface is a switch virtual interface (SVI).
- o **IPv4 Secondary GW1** - Enter the gateway IP address for the additional subnet.
- o **IPv4 Secondary GW2** - Enter the gateway IP address for the additional subnet.
- o **IPv4 Secondary GW3** - Enter the gateway IP address for the additional subnet.
- o **IPv4 Secondary GW4** - Enter the gateway IP address for the additional subnet.

6. Configure the following fields in the **Advanced** tab.

- o **Multicast Group Address** - The multicast IP address for the network is auto-populated. Multicast group address is a per fabric instance variable and remains the same for all networks by default.
- o Configure DHCP relay server fields as follows:
 - a. Under **DHCP Relay Server Information**, choose **Actions > Add**.
 - b. Enter the DHCP relay IP address of the first DHCP server in the **Server 1 V4 Address** field.
 - c. Enter the DHCP server VRF ID in the **Server VRF** field.
 - d. Click **Save**.

Note that you can configure up to 16 DHCP servers.

- o **Loopback ID for DHCP Relay interface (Min:0, Max:1023)** - Enter the loopback ID for DHCP relay interface.
- o **TRM Enable** - Check the checkbox to enable TRM.
- o **L2 VNI Route-Target Both Enable** - Check the check box to enable automatic importing and exporting of route targets for all L2 virtual networks. This is applicable only for Cisco Nexus 9000 switches.
- o **Interface Vlan Netflow Monitor** - Specifies the Netflow monitor specified for Layer 3 record for the VLAN interface. This is applicable only if Layer 2 Record is not enabled in the Netflow Record for the fabric. This is applicable only for Cisco Nexus 9000 switches.
- o **Vlan Netflow Monitor** - Enter the Netflow monitor name defined in the fabric setting for Layer 3 Netflow Record. This is applicable only for Cisco Nexus 9000 switches.

- **Enable L3 Gateway on Border** – Check the check box to enable a Layer 3 gateway on the border switches.

7. Click **Create**.

A message appears indicating that the network is created. The new network appears on the **Networks** page that comes up.

The Status appears as **NA** since the network is created but not yet deployed on the switches. Now that the network is created, you can create more networks if needed and deploy the networks on the devices in the fabric.

Attaching Networks in Campus VXLAN EVPN Fabric

To attach a network in Campus VXLAN EVPN fabric from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. In Cisco Nexus Dashboard Fabric Controller, choose **LAN > Fabrics**.
2. From the list of available fabrics, double-click the Campus VXLAN EVPN fabric.

The **Fabric Overview** page appears.

3. Navigate to the **Networks** tab and double-click the network that you created in the previous section to open the **Network Overview** page.
4. On the **Network Attachments** tab, select the required network and choose **Actions > Edit**.

The **Edit Network Attachment** page opens.

5. Use the toggle switch to enable **Attach** and then click **Save**.

Deploying Networks in Campus VXLAN EVPN Fabrics

To attach networks in Campus VXLAN EVPN fabric from the Cisco Nexus Dashboard Fabric Controller Web UI, perform the following steps:

1. In Cisco Nexus Dashboard Fabric Controller, choose **LAN > Fabrics**.
2. From the list of available fabrics, double-click the Campus VXLAN EVPN fabric.

The **Fabric Overview** page appears.

3. Navigate to the **Switches** tab, select the switches and choose **Actions > Deploy**.
4. Click **Deploy All** after the configuration preview is complete.
5. Click **Close** after the deployment is complete.

Creating DCI Links for Switches in Campus VXLAN EVPN Fabric

You can create VRF-Lite IFC between a Cisco Catalyst 9000 Series Switch or a Cisco Nexus 9000 switch with border role in Campus VXLAN EVPN fabric, and another switch in a different fabric. The other switch can be a Cisco Nexus 9000 switch in External Fabric, LAN Classic fabric, or Campus VXLAN EVPN fabric. It can also be a Catalyst 9000 switch in External Fabric or Campus VXLAN EVPN fabric. The link can be created only from Campus VXLAN EVPN fabric. The other switch can be in External, Classic LAN, Campus VXLAN EVPN or Data Center VXLAN EVPN fabrics. The link can be created only from the Campus VXLAN EVPN fabric.

For more information, see the section "Links" in [About Fabric Overview for LAN Operational Mode Setups](#) and [Templates](#).



When creating DCI links for Campus VXLAN EVPN fabric, auto-deploy is supported only if the destination device is a Cisco Nexus 9000 switch.

To create links for Campus VXLAN EVPN fabric, perform the following procedure:

1. Navigate to the **Links** tab in the fabric overview.

The list of previously created links is displayed. The list contains intra-fabric links, which are between switches within a fabric, and inter-fabric links, which are between border switches in this fabric and switches in other fabrics.

2. Choose **Actions > Create**.

The **Create Link** window appears. By default, the **Intra-Fabric** option is chosen as the link type.

3. From the **Link Type** drop-down box, choose **Inter-Fabric**.
4. From the **Link Sub-Type** drop-down list, choose **VRF_LITE**.
5. In the **Link Template** field, ensure **ext_fabric_setup** template is auto populated for VRF_LITE IFC.

The templates are autopopulated with corresponding pre-packaged default templates that are based on your selection. The template to use for **VRF_LITE** IFC is **ext_fabric_setup**.

6. From the **Source Fabric** drop-down list, choose the campus VXLAN EVPN fabric.
7. From the **Destination Fabric** drop-down list, choose a destination fabric.
8. Choose the **Source Device** and Ethernet interface that connects to the destination device.
9. Choose the **Destination Device** and Ethernet interface that connects to the source device.
10. Enter values in other fields, as required. For more information about configuring the fields in the **Default VRF** tab, see [VRF Lite](#)
11. Click **Save**.

Instead of the **Create** option, you can also use **Edit** to create VRF-Lite IFC(s) using the existing inter-fabric link(s). Choose the **VRF_Lite** link subtype. By default, if you select **Edit**, then the data for the fields **Link-Type**, **Source Fabric**, **Destination Fabric**, **Source Device**, **Destination Device**, **Source Interface** and **Destination Interface** are auto-populated in the **Edit Link** window.

Choose **VRF_LITE** as the link sub-type and **ext_fabric_setup** template for VRF_LITE IFC. To complete the procedure, repeat step 4 to step 10 mentioned above.

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.