VRF Lite

External connectivity from data centers is a prime requirement. External connectivity where workload is the part of data center fabric which can communicate with outside fabric over WAN/Backbone services. The VRF Lite feature is used between Data Center Border devices and the External fabric Edge router for connecting the fabric to an external Layer 3 domain, north-south traffic communication. In Virtual eXtensible Local Area Network (VXLAN) Ethernet Virtual Private Network (EVPN) fabric can be a border router or a Border Gateway router.

You can enable VRF Lite on below devices:

- Border
- Border Spine
- Border Gateway
- Boder Gateway Spine
- Border Super Spine
- Border Gateway Super Spine

VRF-Lite

External connectivity from data centers is a prime requirement where workloads that part of data center fabric can communicate with outside fabric over WAN/Backbone services. To enable Layer-3 for north-south traffic flow use BRF-Lite peering between data center border devices and the external fabric edge routers.

In a VXLAN (Virtual extensible Local Area Network) EVPN (Ethernet Virtual Private Network) fabric, this can be a border router or a Border Gateway router. You can enable VRF-Lite on the following devices:

- Border
- Border Spine
- Border Gateway
- Boder Gateway Spine
- Border Super Spine

Prerequisites and Guidelines

Prerequisites

- VRF Lite requires Cisco Nexus 9000 Series NX-OS (Nexus Operating System) Release 7.0(3)I6(2) or later.
- Familiarity with VXLAN BGP (Border Gateway Protocol) EVPN data center fabric architecture and VXLAN Overlays provisioning through the NDFC.
- Fully configured VXLAN BGP EVPN fabrics including underlay and overlay configurations on the various leaf and spine devices, external fabric configuration through NDFC, and relevant external fabric device configuration (edge routers, for example).
    - A VXLAN BGP EVPN fabric (and its connectivity to an external Layer 3 domain for north-south traffic flow) can be configured manually or using NDFC. This document explains the process to connect the fabric to an edge router (outside the fabric, towards the external fabric) through NDFC. So, you must know how to configure and deploy VXLAN BGP EVPN and external fabrics through NDFC.
    - VRF Lite can be enabled on physical Ethernet interface or Layer 3 port-channel. Sub-interface over physical interface or Layer 3 port-channel interface is created by NDFC at the VRF extension moment for each VRF lite link the VRF is extended over.
- To delete a VRF Lite IFC, remove all VRF extensions enabled on the IFC. Else, an error message is reported. Once the VRF Lite attachments are removed (detached) recalculate and deploy the fabric to remove any pending Layer-3 extension configurations. This removes the per-VRF sub-interface and per-VRF eBGP configuration on the devices.
- When you create a VXLAN VRF, ensure that the following 3 fields:
    - **Advertise Host Routes** – By default, over the VRF Lite peering session, only non-host (/32 or /128) prefixes are advertised. But if host routes (/32 or /128) need to be enabled and advertised from the border device to the edge/WAN router, then the **Advertise Host Routes** check box can be enabled. Route-map does outbound filtering. By default, this check box is disabled.
    - **Advertise Default Route** – This field controls whether a network statement 0/0 will be enabled under the vrf. This in turn will advertise a 0/0 route in BGP. By default, this field is enabled. When the check box is enabled, this will ensure that a 0/0 route is advertised inside the fabric over EVPN Route-type 5 to the Leafs there by providing a default route out of the Leafs toward the border devices.
    - **Config Static 0/0 Route** –This field controls whether a static 0/0 route to the edge/WAN router, must be configured under the VRF on the border device. By default, this field is enabled. If

WAN/edge routers are advertising a default route over the VRF Lite peering, to the border device in the fabric, then this field should be disabled. In addition, the **Advertise Default Route** field must be disabled. This is because the 0/0 route advertised over eBGP sends over EVPN to the Leafs without the need for any additional configuration. The clean iBGP EVPN separation inside the fabric with eBGP for external out-of-fabric peering provides for this desired behavior. By default, this check box is enabled.

Sample Scenarios

Scenarios explained in this document:

- VRF Lite between Cisco Nexus 9000 based Border and Cisco Nexus 9000 based Edge Router
- VRF Lite between Cisco Nexus 9000 based Border and Non-Cisco device
- VRF Lite between Cisco Nexus 9000 based Border and Non-Nexus device. A typical use-case of having Cisco ASR 9000 based Edge Router in Managed mode

Easy Fabric Settings

There are 4 modes for VRF Lite Deployments. By default, VRF Lite deployment is set to Manual. You can change the settings based on the requirement:

- **Manual** - Use the NDFC application to deploy the VRF Lite IFCs manually between Source and Destination device. This is default setting on NDFC.
- **ToExternal Only** - Configure a VRF Lite IFC on each physical interface of a border leaf device in the VXLAN fabric that is connected to a device with the **Edge Router** role in the external fabric.
- **Back-to-Back Only** - Configure VRF Lite IFCs between directly connected border leaf device interfaces of different VXLAN fabrics.
- **Back2Back&ToExternal** - Use this option to configure IFCs for the modes **To External Only** and **Back-to-Back Only**.

Note DCI (Data Center Interconnectivity) subnet is required, even if the VRF Lite mode is **Manual** for NDFC resource handling.

The **Manual** mode is the default mode in fabric settings. To change the default mode to other mode, click **Edit** fabric settings. On **Resoruce** tab, modify VRF Lite deployment field to above mentioned auto configuraiton modes. Here in the example, **ToExternal Only** check box is chosen.

Auto Deploy Both - This check box is applicable for the symmetric VRF Lite deployment. When you check this check box, the **Auto Deploy Flag** is set to true for auto created IFCs to turn on symmetric VRF Lite configuration. You can check or uncheck this check box when the **VRF Lite Deployment** field is not set to **Manual**. The value you choose takes priority. This flag only affects the new auto created IFC and it does not affect the existing IFCs.

VRF Lite Subnet IP Range: The IP address for VRF Lite IFC deployment is chosen from this range. The default value is 10.33.0.0/16. The best practice is to ensure that each fabric has its own unique range and is

distinct from any underlay range in order to avoid possible duplication. These addresses are reserved with the Resource Manager.

VRF Lite Subnet Mask: By default, it's set to /30 which is best practice for P2P links.

Automatic VRF Lite (IFC) Configuration

You can enable VRF Lite auto-configuration by changing the fabric settings of the **VRF Lite Deployment** field under the **Resources** tab from **Manual** to any of the auto-configuration settings.

Note On the fabric topology screen, you can view only the individual fabric and the external fabric connected.

- The topology displays VXLAN BGP EVPN fabrics **Easy60000** (at the left) and **Easy7200** (at the right) and external fabric **External65000** (at the top). The border leaf of one VXLAN fabric is connected to the border leaf of the other through the edge router **n7k1-Edge1** in the external fabric.
- The border leafs are special devices that allow clear control and data plane segregation from the fabric to the external Layer 3 domain while allowing for policy enforcement points for any inter-fabric traffic. Multiple border devices in the fabric ensure redundancy in the case of failures and effective load distribution. This document shows you how to enable Layer 3 north-south traffic between the VXLAN fabrics and the external fabric.
- Before VRF Lite configuration, end hosts associated with a specific VRF can send traffic to each other, but only within the fabric. After VRF Lite configuration, end hosts can send traffic across fabrics.
- Network configurations for the VXLAN fabric are provisioned through NDFC.

  The template used for VRF Lite IFC auto configuration is **ext_fabric_setup_11_1**. You can edit the **ext_fabric_setup_11_1** template or create a new one with custom configurations.

Automatic VRF Lite Creation Rules

- The Auto IFC is supported for the Cisco Nexus devices only.
- You can configure a Cisco ASR 1000 Series routers and Cisco Catalyst 9000 Series switches as edge routers, set up a VRF Lite IFC, and connect it as a border device with an easy fabric.
- If the device in the External fabric is non-Nexus, then IFC must be created manually.
- Ensure that no user policy is enabled on the interface that connects to the edge router. If a policy exists, then the interface will not be configured.
- Auto configuration is only provided for the **Border** or **Border Spine** role in the VXLAN fabric and **Edge Router** role in the connected external fabric device.

  Auto configuration is provided for the following cases:

  - **Border** role in the VXLAN fabric and **Edge Router** role in the connected external fabric device
  - **Border Gateway** role in the VXLAN fabric and **Edge Router** role in the connected external fabric device
  - **Border** role to another **Border** role directly

Note that auto configuration is not provided between two BGWs.

If you need a VRF Lite between any other roles, then you have to deploy it manually through the NDFC GUI.

- To deploy configurations in the external fabric, ensure that the **Fabric Monitor Mode** check box is cleared in the external fabric settings of the **External65000** fabric. When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches.

VRF Lite between Cisco Nexus 9000 based Border and Cisco Nexus 9000 based Edge Router

In below figure, a VXLAN EVPN Fabric with name DC-Vxlan is connected to WAN-Vxlan cloud. The Easy fabric has border leaf role and WAN-Vxlan cloud has a device with role edge router. NDFC shows physical and logical representation of the toplogy with CDP/LLDP Link discovery.

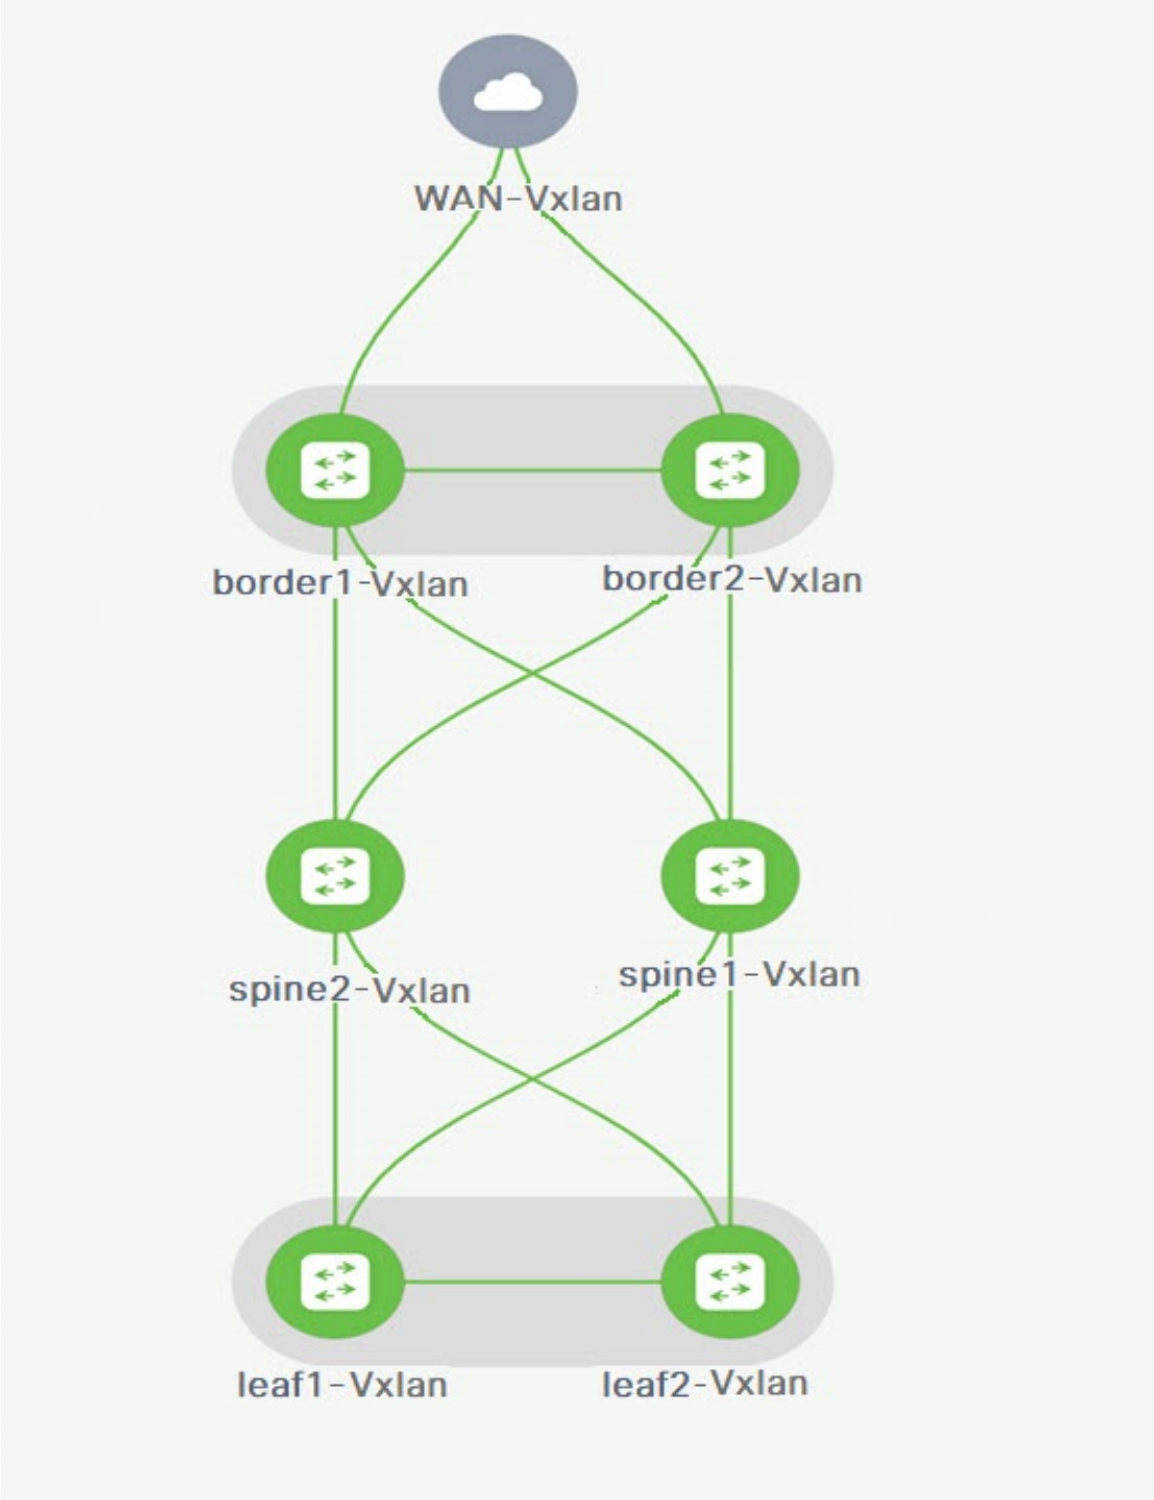

In this example, you can enable VRF Lite connections between DC-Vxlan Border leaf and WAN-Vxlan Edge router.

For VRF Lite configuration, you must enable eBGP peering between the fabric's Border interfaces and the edge router's interfaces, through point-to-point connections. The Border physical interfaces are:

- **eth1/1** on **border1-Vxlan**, towards **eth1/1** on **WAN1-Vxlan**.
- **eth1/2** on **border2-Vxlan**, towards **eth1/2** on **WAN1-Vxlan**.

1. Verify the links between the Border and the Edge router. Navigate **LAN** > **Fabrics**, double-click on **DC-Vxlan** fabric.

   On **Fabric Overview** window, click on **Links** tab. You can view the links detected by NDFC and appropriate policy **ext_fabric_setup** is assigned automatically.

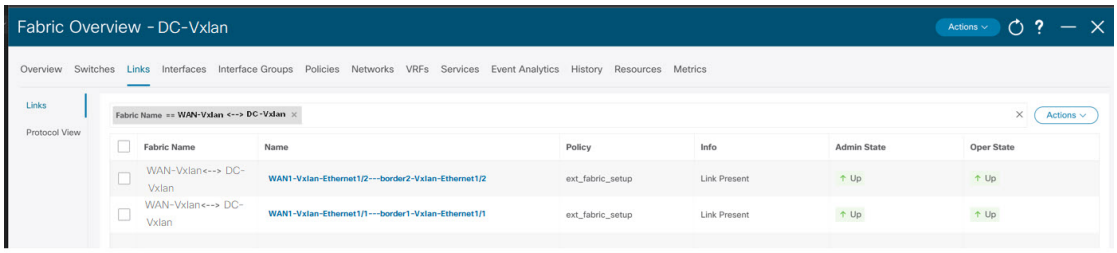

2. To verify the VRF Lite configurations, choose fabric name and choose **Actions** > **Edit**.

   Click on appropriate **Links**, choose **Actions** > **Edit**.

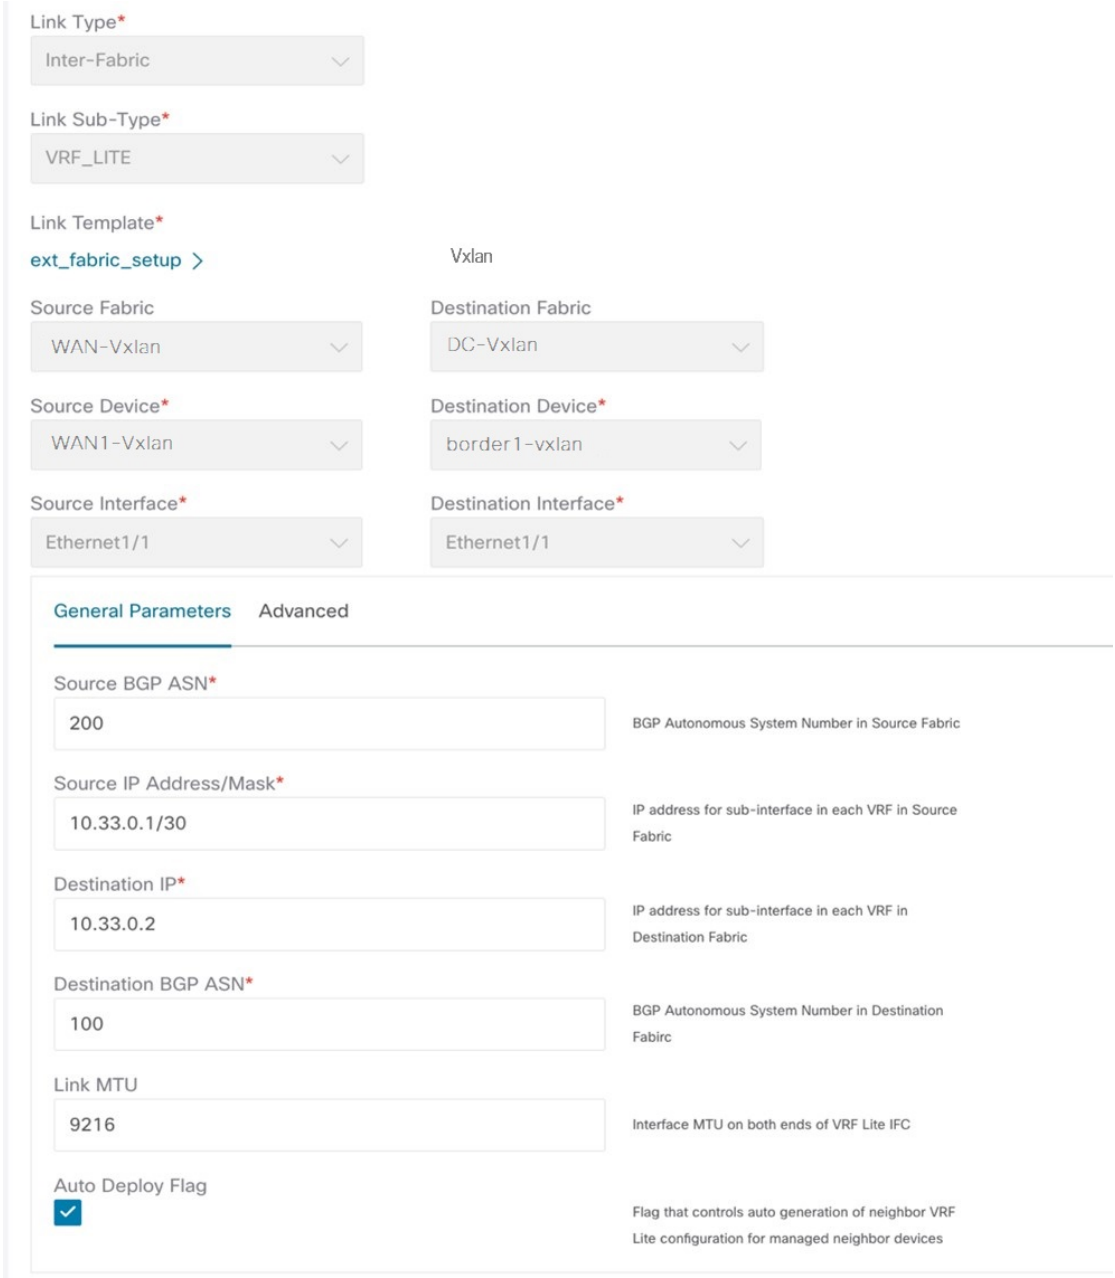

Link Type – Specifies the Inter-fabric link between two different fabrics within NDFC.

Link Sub-Type – Specifies the sub-type of link. By default, the **VRF_LITE** option is displayed.

Link Template – Specifies the template for the link. The default template for a VRF Lite IFC is **ext_fabric_setup** is displayed. The template enables the source and destination interfaces as Layer 3 interfaces, configures the **no shutdown** command, and sets their MTU to 9216.

The Source and Destination Fabric, Device, and Interfaces are auto detected and selected by NDFC based on CDP/LLDP discovery.

On the **General Parameters** tabs, the fields in this tab are:

Source BGP ASN – BGP ASN of selected Source fabric

Source IP Address/Mask - NDFC auto allocated IP pool from Resource manager pool of VRF Lite subnet pool for the **Ethernet1/1** sub interfaces, the source interface of the IFC. A sub-interface is created for each VRF extended over this IFC, and a unique 802.1Q ID is assigned to it. The IP address/Mask entered here, along with the BGP Neighbor IP field (explained below) will be used as the default values for the sub-interface created at VRF extension and can be overwritten.

For example, an 802.1Q ID of 2 is associated with subinterface Eth 1/1.2 for VRF CORP traffic, and 802.1Q ID of 3 is associated with Eth 1/1.3 and VRF ENG, and so on.

The IP prefix is reserved with the NDFC resource manager. Ensure that we use a unique IP address prefix for each IFC we create in the topology.

Destination IP - NDFC auto allocated IP pool from Resource manager pool of VRF Lite subnet pool. This is a BGP neighbor IP on the device.

Inter-fabric traffic from different VRFs for an IFC will have the same source IP address (10.33.0.1/30) and destination IP address (10.33.0.2) as an example.

Destination BGP ASN – BGP ASN of selected Destination fabric

Link MTU – Default 9216

Auto Deploy Flag – Default Auto selected based on fabric settings. This knob will auto configure the neighbor VRF on neighboring managed device. For example, it will automatically create VRF on the Edge router inside WAN-Vxlan External fabric.

The **Advanced** tab is added in the **Link Profile** section. The fields in this tab are:

- **Source Interface Description**
- **Destination Interface Description**
- **Source Interface Freeform Config**
- **Destination Interface Freeform Config**

Click **Save** to save the configuration.

3. To attach VRF and VRF Lite extensions on the Border devices.
   a. Click on **VRFs** > **VRF Attachments** tab.
   b. Choose **VRF Name**, click **Actions** > **Edit**

      The **Edit window** appears.
   c. Edit details in **Extension** field as mentioned below:

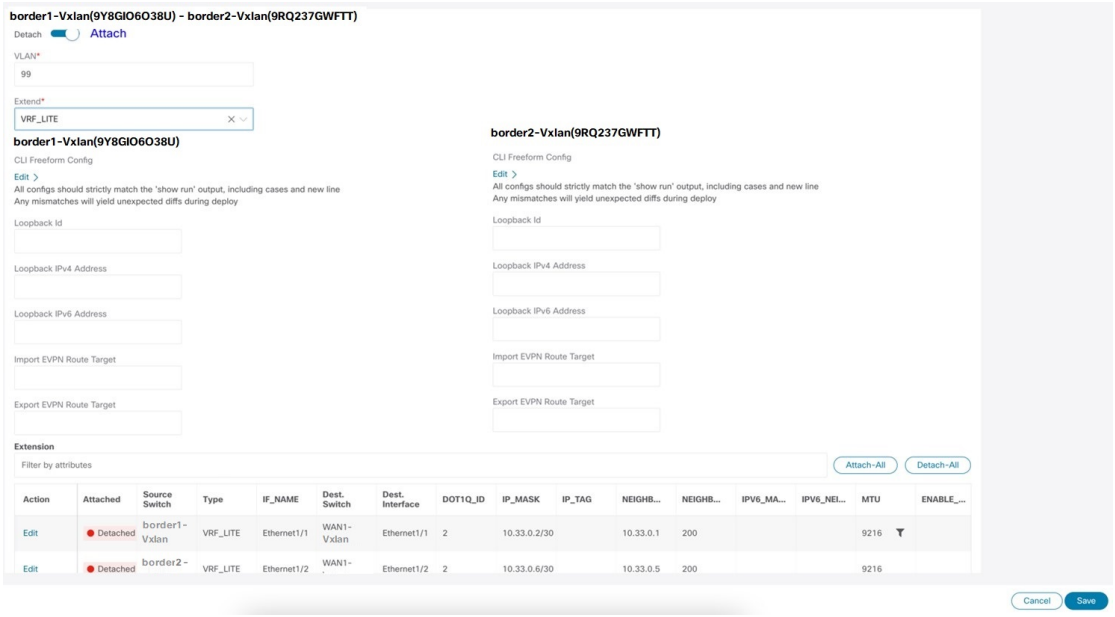

- Toggle the knob to **Attach**
- In **Extend** field, choose **VRF_LITE** from drop-down list.
- On **Extension** area, choose one after another switch and click **Edit**, enter details for **PEER_VRF_NAME**. This will auto deploy the VRF on the neighboring device.

When you extend VRF Lite consecutive scenario, the VRF must be in the peer fabric and VRF name must be same. If the VRF is not in the peer fabric and if you try to extend VRF Lite, an error message is generated displaying the issue.

When you extend VRF Lite between an easy fabric and an external fabric, the VRF name can be same as name of source fabric, or default name, or an other VRF name. Enter required VRF name in **PEER_VRF_NAME** field. The child PTIs for subinterface, VRF creation and BGP peering on external fabric have source values populated in it, hence the policies cannot be edited or deleted.

Follow above procedure for other link.

- On **Edit** window, click **Attach-All**, to attach the required VRF Extension on the border devices, and then click **Save**.

4. To Recalculate and deploy configurations on VXLAN EVPN Easy Fabric.

   On **Fabric** window double-click on appropriate fabric to navigate to **Fabric Overview** window. Click **Actions** > **Recalculate & Deploy**.

   OR

   Choose required **VRF Name** on **VRF attachments** tab, click **Actions**> **Deploy** to initiate VRF and VRF Lite configurations on the border devices.

5. Recalculate and Deploy on VXLAN EVPN Easy fabric.

   Click on the top Action button and then Recalculate and Deploy OR simply select the VRF attachments then edit and finally deploy button. This will trigger the VRF and VRF Lite configurations on the Border devices.

6. To recalculate and deploy configurations on External fabric, choose external fabric and follow the above procedure.

VRF Lite between Cisco Nexus 9000 based Border and Non-Cisco device

This example displays procedure to enable VRF-Lite connections between the DC-Vadodara Border leaf and a non-Cisco device in External fabric.

It is recommended to use meta definition of a device instead of importing devices in external fabric. This allows VRF-Lite configurations to extend Nexus 9000 managed border devices in easy fabric. NDFC will not manage destination non-cisco device. you must configure relevant VRF-Lite configurations on the destination device.

1. To create new IFC links between border and edge router.
   a. On **Fabrics** window, double click on the fabric.

      The **Fabric Overview** window appears.

   b. Click on **Links** tab. On **Links** tab, click **Actions** > **Create a new link**.

c.

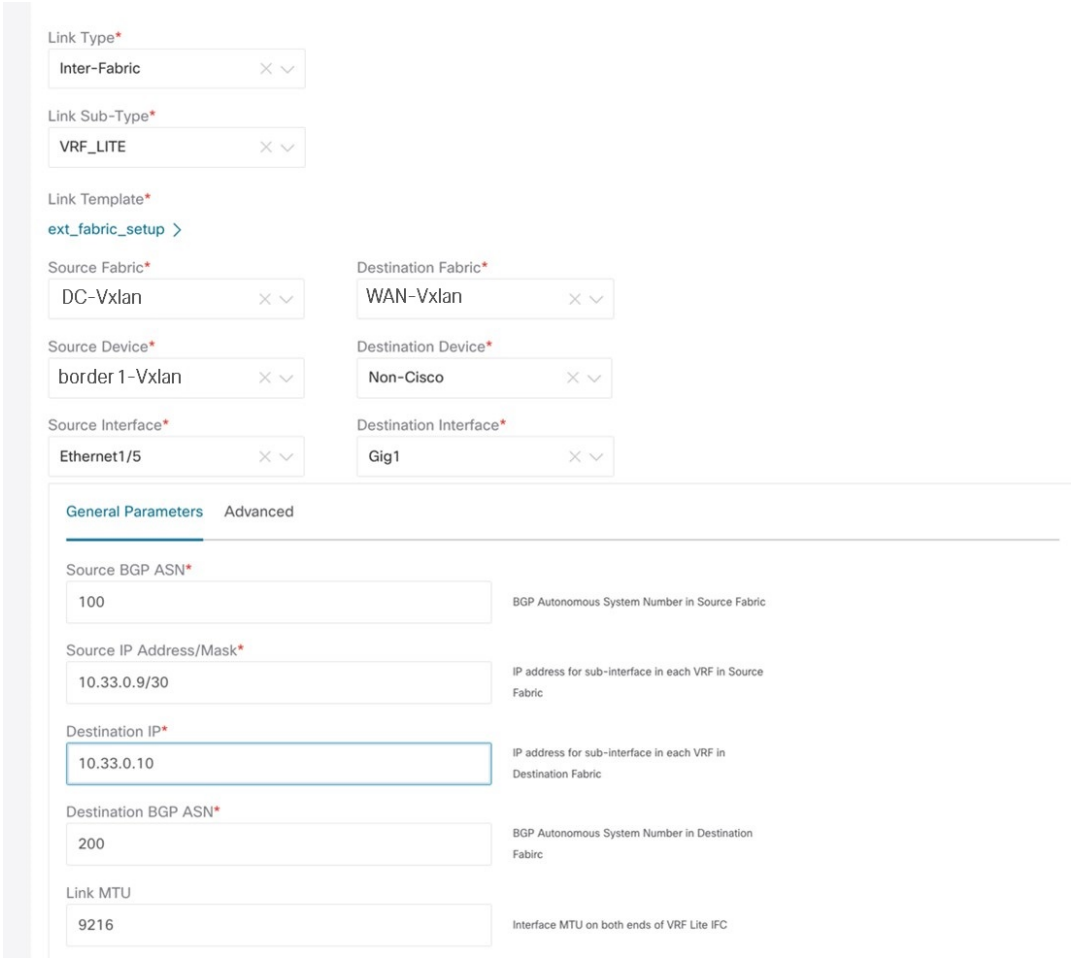

Enter the below required parameters in the window:

- **Link Type** – Select the Inter-fabric link. This is the IFC between two different fabrics within NDFC.
- **Link Sub-Type** - By default, the **VRF_LITE** option is displayed.
- **Link Template** – The default template for a VRF Lite IFC, **ext_fabric_setup**, is displayed. The template enables the source and destination interfaces as Layer 3 interfaces, configures the **no shutdown** command, and sets their MTU to 9216.
- **Source Fabric** – Select the Source Fabric. This will be the Easy fabric where Nexus 9000 based Border device resides.
- **Destination Fabric** – Select any External/Classic LAN fabric. It can be monitor mode as well.
- **Source Device** – Select the Source Device. This will be the Nexus 9000 based Border device.
- **Destination Device** – Now, we will be creating a “meta device definition”. Type any name and click on create. As an example, “non-cisco”.
- **Source Interface** – Select the interface on the border device where the non-cisco device will be connected.

- **Destination Interface** – Now, we will be creating a “meta device interface”. Type any interface name and click on create. Valid examples are “gig1, tengig1/10, eth1/1”

The **General Parameters** tab has the following:

- **Source BGP ASN** – BGP ASN of selected Source fabric
- **Source IP Address/Mask** - Provide IP address and mask for the **Ethernet1/5** sub interfaces, the source interface of the IFC. A sub-interface is created for each VRF extended over this IFC, and a unique 802.1Q ID is assigned to it. The IP address/Mask entered here, along with the BGP Neighbor IP field (explained below) will be used as the default values for the sub-interface created at VRF extension and can be overwritten.

  For example, an 802.1Q ID of 2 is associated with subinterface Eth 1/5.2 for VRF CORP traffic, and 802.1Q ID of 3 is associated with Eth 1/5.3 and VRF ENG, and so on.

  The IP prefix is reserved with the NDFC resource manager. Ensure that we use a unique IP address prefix for each IFC we create in the topology.
- **Destination IP** - NDFC auto allocated IP pool from Resource manager pool of VRF-LITE subnet pool. This is a BGP neighbor IP on the device.

  Inter-fabric traffic from different VRFs for an IFC will have the same source IP address (10.33.0.1/30) and destination IP address (10.33.0.2) as an example.
- **Destination BGP ASN** – BGP ASN of selected Destination fabric
- **Link MTU** – Default 9216
- **Auto Deploy Flag** – Not applicable as the destination device is Non-Nexus and Non-Cisco.

The **Advanced** tab has the following fields:

- **Source Interface Description**
- **Destination Interface Description**
- **Source Interface Freeform Config**
- **Destination Interface Freeform Config**

2. Click **Save** to creat new link with parameters mentioned.
3. To attach VRF and VRF-Lite extensions on the Border devices, double-click on **DC-Vadodara Easy fabric**. On **Fabric Overview** window, navigate to **VRFs** > **VRF Attachments** and edit the details as shown in image below:

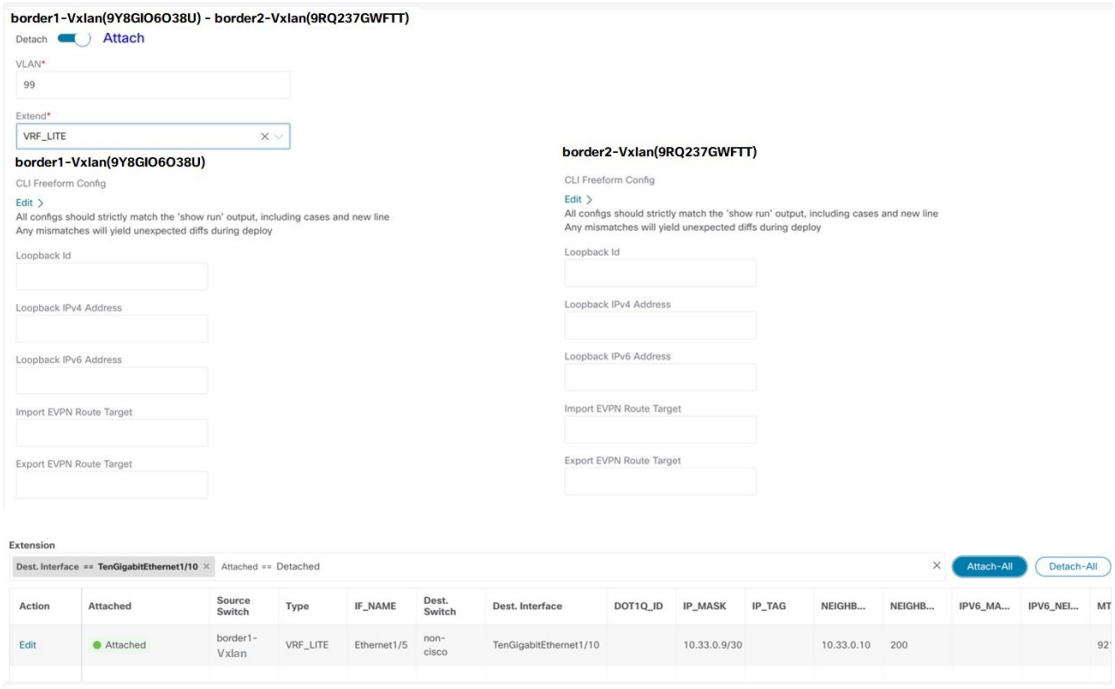

- Click on **Attach-all** to attch the required VRF Extension on the border devices and then click **Save**.

4. To recalculate and deploy configurations on VXLAN EVPN Easy fabric, click on appropriate fabric on **Fabric** window.
   a. On **Fabric Overview** window, click **Actions** > **Recalculate & Deploy**, or navigate to **VRFVRF attachments**, choose appropriate VRF attachments then edit and finally deploy button. This will initiate the VRF and VRF-LITE configurations on the Border devices.

VRF Lite between Cisco Nexus 9000 based Border and Non-Nexus device

In this example you can enable VRF Lite connections between DC-Vadodara border leaf and a non-Nexus device in external fabric.

Before Cisco NDFC Release 12.0.1a, ASR 9000 was supported for external fabric in monitor mode only. Starting with Release 12.0.1a, ASR 9000 is supported in managed mode with edge router role.

The following are the supported platforms:

- ASR 9000
- NCS 5500
- ASR 8000

Configuration compliance is enabled for IOS-XR switches in external fabric. NDFC

Config-compliance is enabled for IOS-XR switches, like the way Nexus switches are handled in External Fabric. NDFC will send commit in the end of deployment.

Note Ensure that the VXLAN BGP EVPN border device is active.

Procedure

Step 1 Naivgate to **LAN** > **Fabrics**, create an external fabric.

Step 2 On **Create Fabric** window, enter appropriate ASN number, uncheck **monitor mode** check box, and click **Save**.

Step 3 Navigate to ui**Switches** window, click **Actions** > **Add switches**.

Note Ensure that the IOS-XR device has the IP address reachability to NDFC with SNMP configurations for the discovery.

To add non-Nexus devices to external fabric, refer to Adding Non-Nexus Devices to External Fabrics section.

Step 4 On **Add Switches** window, choose **Discover** check box, and **IOS-XR** from drop-down list for **Device Type** field.

Step 5 After the router is discovered, you can view the switch name in the **Discovery Results** field.

Step 6 Choose the discovered router and add to the fabric and ensure that the **Discovery Status** displays **OK** in the status column. Edge router role is supported.

After successful discovery, you can view the links between the devices in the **Links** tab.

Step 7 To create VRF Lite IFC from easy fabric with Nexus 9000 border leaf, choose the link and click **Actions** > **Edit**.

Step 8 On **Edit Link** window, fill the required

Step 9 Enter the required details for IFC creation. Few fields are auto-populated.

Note For non-NXOS device auto deploy flag is not applicable.

Step 10 To extend VRF Lite configurations on VXLAN border device, navigate to **VRF** > **VRF Attachment** tab, choose the VRF name and click **Actions** > **Edit** and then extend it as VRF Lite

Step 11 Deploy the configuraiton on VXLAN border device.

Step 12 Navigate to external fabrics which has router, click **Apply** to VRF Lite BGP policies.

Step 13 Navigate to **Policies** tab, and add policies **ios_xr_base_bgp** and enter required details and save then and another policy **ios_xr_Ext_VRF_Lite_Jython**.

Step 14 Deploy the configurations on the XR router.

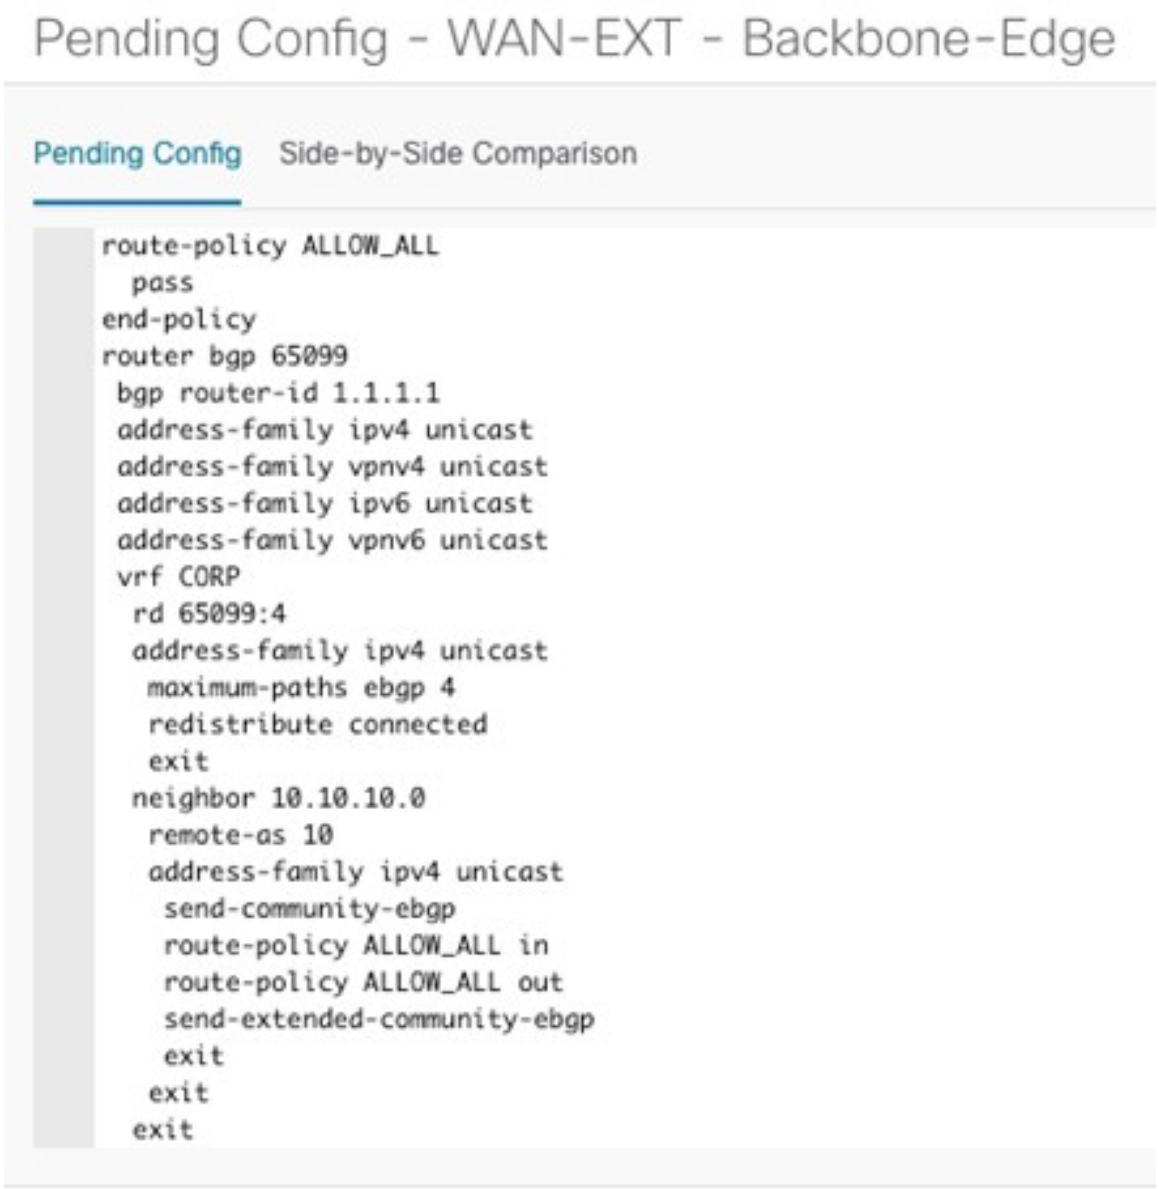

Appendix

Nexus 9000 Border device configurations

Border-Vxlan (base border configurations) generated by template ext_base_border_vrflite_11_1

switch(config)# refers to the global configuration mode. To access this mode, type the following on your switch: switch# **configure terminal**.

```
(config)#
ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
```

```
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map extcon-rmap-filter deny 10
    match ip address prefix-list default-route
route-map extcon-rmap-filter deny 20
    match ip address prefix-list host-route
route-map extcon-rmap-filter permit 1000
route-map extcon-rmap-filter-allow-host deny 10
    match ip address prefix-list default-route
route-map extcon-rmap-filter-allow-host permit 1000
ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
route-map extcon-rmap-filter-v6 deny 10
    match ipv6 address prefix-list default-route-v6
route-map extcon-rmap-filter-v6 deny 20
    match ip address prefix-list host-route-v6
route-map extcon-rmap-filter-v6 permit 1000
route-map extcon-rmap-filter-v6-allow-host deny 10
    match ipv6 address prefix-list default-route-v6
route-map extcon-rmap-filter-v6-allow-host permit 1000
```

Border-Vxlan VRF Lite Extension configuration

```
(config)#

vrf context CORP
  ip route 0.0.0.0/0 2.2.2.2
exit
router bgp 100
  vrf CORP
    address-family ipv4 unicast
      network 0.0.0.0/0
      exit
    neighbor 2.2.2.2
      remote-as 200
      address-family ipv4 unicast
        send-community both
        route-map extcon-rmap-filter out
configure terminal
interface ethernet1/1.2
  encapsulation dot1q 2
  mtu 9216
  vrf member CORP
  ip address 2.2.2.22/24
  no shutdown
configure terminal
```

WAN-Vxlan (External fabric Edge Router) VRF Lite Extension configuration

```
(config) #
vrf context CORP
  address-family ipv4 unicast
exit
router bgp 200
  vrf CORP
    address-family ipv4 unicast
    neighbor 10.33.0.2
      remote-as 100
      address-family ipv4 unicast
        send-community both
        exit
      exit
    neighbor 10.33.0.6
      remote-as 100
      address-family ipv4 unicast
        send-community both
```

```
configure terminal
interface ethernet1/1.2
  mtu 9216
  vrf member CORP
  encapsulation dot1q 2
  ip address 10.33.0.1/30
  no shutdown
interface ethernet1/2.2
  vrf member CORP
  mtu 9216
  encapsulation dot1q 2
  ip address 10.33.0.5/30
  no shutdown
configure terminal
```