



# Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site

---

This chapter explains LAN Fabric border provisioning using EVPN Multi-Site feature.

- [Overview, page 1](#)
- [Prerequisites, page 1](#)
- [Limitations, page 2](#)
- [Sample Scenario, page 2](#)
- [EVPN Multi-Site Configuration, page 4](#)
- [Deploying Networks and VRF Instances, page 20](#)
- [Additional References, page 26](#)
- [Appendix, page 27](#)

## Overview

This document explains how to connect two Virtual eXtensible Local Area Network (VXLAN) Border Gateway Protocol (BGP) Ethernet VPN (EVPN) fabrics through DCNM using EVPN Multi-Site feature. The EVPN Multi-Site configurations are applied on the Border Gateways (BGWs) of the two fabrics. Apart from VXLAN BGP EVPN fabrics, EVPN Multi-Site also allows you to extend Layer 2 and Layer 3 connectivity to data center networks built with older (legacy) technologies (Spanning Tree Protocol, virtual Port Channel [vPC], Cisco FabricPath, etc).

## Prerequisites

- The EVPN Multi-Site feature requires Cisco Nexus 9000 Series NX-OS Release 7.0(3)I7(1) or later.
- Familiarity with VXLAN BGP EVPN data center fabric architecture and configuration through DCNM.
- Fully configured VXLAN BGP EVPN fabrics and connected device (route servers, for example) configurations that are ready to be connected using the EVPN Multi-Site feature.

- VXLAN BGP EVPN fabrics (and their interconnection) can be configured manually or using Cisco® Data Center Network Manager (DCNM). This document explains the process to connect the fabrics through DCNM. So, you should know how to configure and deploy a VXLAN BGP EVPN fabric through DCNM. For more details, see the LAN Fabric Provisioning section under Configure chapter in [Cisco DCNM Web Client Online Help, 10.4\(2\) Release](#).



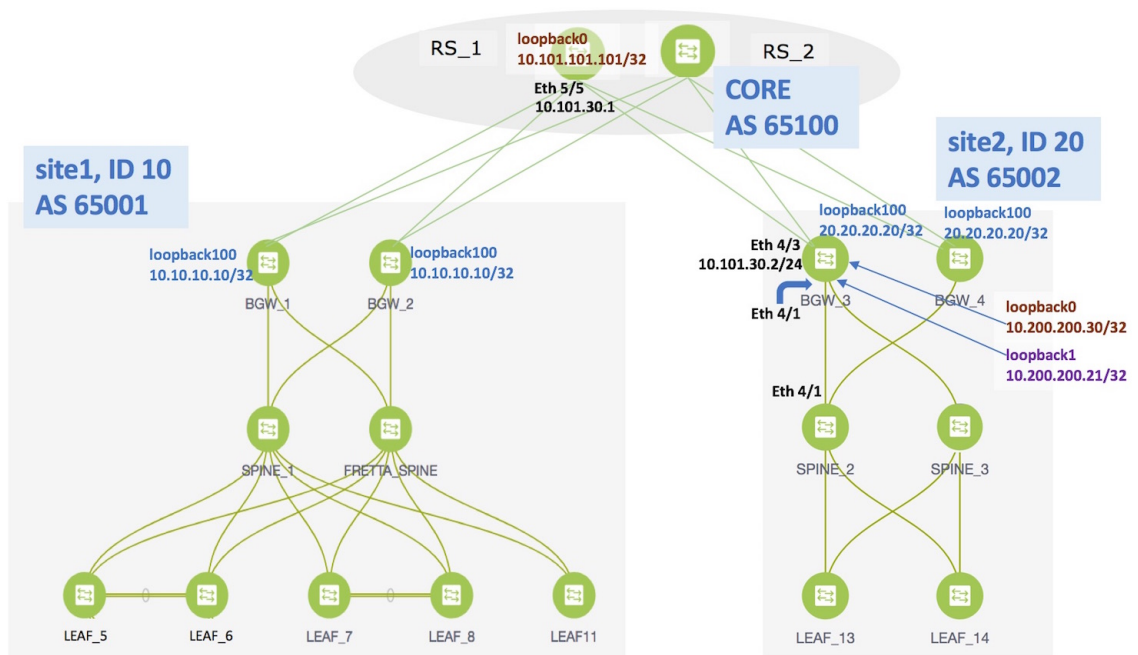
**Note** For a detailed explanation on the EVPN Multi-Site feature, see the [VXLAN BGP EVPN Multi-Site Design and Deployment](#) document.

## Limitations

- BGWs cannot form a virtual port channel (vPC) switch pair. This is a switch side limitation and not a DCNM software limitation.

## Sample Scenario

The EVPN Multi-Site feature is explained through an example scenario. Consider two VXLAN BGP EVPN fabrics, site1 and site2. This document will show you how to enable end-to-end Layer 3 and Layer 2 traffic between hosts in site1 and site2.



Network configurations for the two fabrics are provisioned through DCNM software, 10.4(2) release. VXLAN BGP EVPN configurations are configured on the switches in the two fabrics. However, server traffic between the sites is only possible through a Data Center Interconnect (DCI) function. If a server in site1 has to send

traffic to a server in site2 or vice versa, the DCI function (such as the Multi-Site feature, used for this example) should be configured on the BGWs of both the fabrics.

**Note**

The DCI functions VRF Lite and VRF Lite + Multi-Site are in the scope of this document, but MPLS L3VPN and LISP technologies are not in the scope of this document.

The steps involved to enable EVPN Multi-Site feature and traffic flow across the sites/fabrics are:

- 1 Top-Down deployment of the underlay for the IP core at the BGWs. This is a one-time configuration.
- 2 Top-Down deployment of the BGP overlay for the IP core. This is a one-time configuration for each BGW.
- 3 Deployment of networks/virtual routing and forwarding (VRF) instances on the leaf switches. This is a per network/VRF configuration.
- 4 Deployment of networks/VRFs at the BGWs. This is a per network/VRF configuration.

***EVPN Multi-Site feature***—This requires setting up the BGW base configuration for enabling the EVPN Multi-Site feature on the BGWs and the underlay peering to the external devices. This is followed by establishing overlay peering from the BGW to appropriate external devices, either BGWs in other fabrics or route servers. Both the underlay and overlay peering are established over eBGP. BGWs are special devices that allow clear control and data plane segregation from one site to another, allowing for policy enforcement points for any inter-fabric traffic. They allow the same data plane (VXLAN) and control plane (BGP EVPN) to be employed both for inter-fabric and intra-fabric traffic.

**Note**

DCNM 10.4.2 Top-Down provisioning only supports eBGP underlay.

The end-to-end configurations can be split into these 2 steps:

- 1 **EVPN Multi-Site configurations on the BGWs (BGW\_1, BGW\_2, BGW\_3 and BGW\_4).**
  - 1 EVPN Multi-Site feature on the BGWs on site1—Overlay and underlay connections between the BGWs BGW\_1 and BGW\_2, and directly connected route servers RS\_1 and RS\_2.
  - 2 EVPN Multi-Site feature on the BGWs on site2—This includes overlay and underlay connections between the BGWs BGW\_3 and BGW\_4, and directly connected route servers RS\_1 and RS\_2.
  - 3 Configurations on RS\_1 and RS\_2—These configurations are not in the scope of DCNM provisioning and this document. For completeness, it is mentioned here, and sample configurations provided in the Appendix section.

For this example, BGW\_3 EVPN Multi-Site configurations will be explained.

- 2 **Deploying Networks and VRF Instances on the leaf switches and the BGWs**

For this example, 2 networks will be configured on the BGWs in site2 (with the assumption that network deployment on leaf switches is already completed).

After successful deployment on both the sites, Layer 2 and Layer 3 traffic will flow between the two sites.

**Note**

In the DCNM GUI, the lines connecting devices managed by DCNM (for example, LEAF\_5 to SPINE\_1 and SPINE\_1 to BGW\_2) symbolize a physical cable connection, and not that the connection is functional and network traffic flows between them.

To start off with, let us consider EVPN Multi-Site provisioning on BGW\_3 through DCNM Top-Down LAN Fabric Provisioning.

## EVPN Multi-Site Configuration

### Prerequisite Configuration for EVPN Multi-Site Feature

- Manual loopback interfaces' configuration.
- Setting the BGW role to *Border Gateway*.

**Manual loopback interfaces' configuration**—Before you begin EVPN Multi-Site configuration through DCNM, you should manually configure loopback interfaces on the BGWs. There are three loopback interfaces configured on a BGW that must be reachable by fabric-internal neighbors and fabric-external neighbors. The fabric internal neighbors will learn these through the fabric interior gateway protocol (IGP). For fabric external

neighbors, these are redistributed into the IPv4 eBGP session. In order to achieve that, the 3 loopback IP addresses must be tagged as shown below:

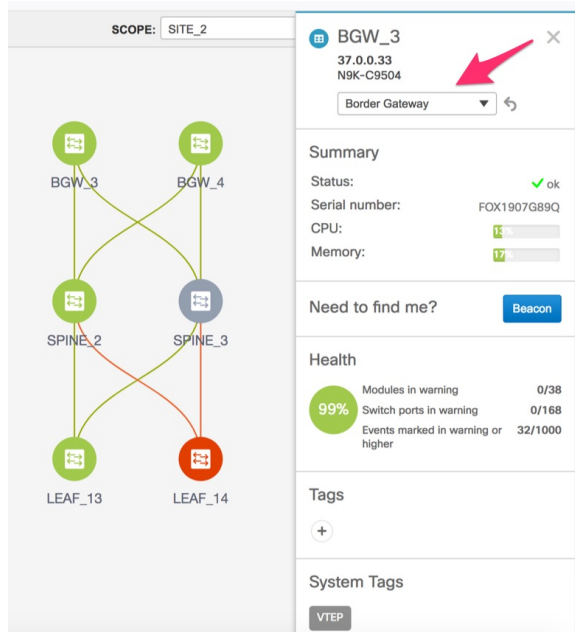
Loopback configuration	Description
<pre>interface loopback0 description RID AND BGP PEERING <b>ip address 10.200.200.30/32 tag 54321</b> ip router ospf UNDERLAY area 0.0.0.0 ip pim sparse-mode</pre>	<ul style="list-style-type: none"> <li>• This is the address used for BGP peering with external and internal neighbors.</li> <li>• In this example, Open Shortest Path First (OSPF) is shown as the fabric underlay routing protocol used for fabric neighbors.</li> <li>• The <b>ip pim sparse-mode</b> setting is needed only for intra-site multicast-based Broadcast, Unknown unicast and Multicast (BUM) replication.</li> </ul>
<pre>interface loopback1 description NVE INTERFACE (PIP VTEP) <b>ip address 10.200.200.21/32 tag 54321</b> ip router ospf UNDERLAY area 0.0.0.0 ip pim sparse-mode</pre>	<p>This is the interface used for local NVE peer address.</p>
<pre>Interface loopback100 description MULTI-SITE INTERFACE (VIP VTEP) <b>ip address 20.20.20.20/32 tag 54321</b> ip router ospf UNDERLAY area 0.0.0.0</pre>	<p>This is the EVPN Multi-Site loopback address. This is provisioned as part of Top-Down auto-configuration of the underlay/overlay, and only shown here for the sake of completeness. This does not need to be pre-provisioned.</p> <p><b>Note</b> The EVPN Multi-Site loopback IP address should be common for all the BGWs in a VXLAN BGP EVPN site.</p>

### Setting the BGW role to *Border Gateway*

After configuring the loopback interfaces on the BGWs, you should change the role of each designated BGW to *Border Gateway*, since, by default a device will be treated as a leaf switch.

To update the switch role, login to DCNM, and click **Topology** from the main menu at the left part of the screen. In the Topology screen, select the fabric/site from the **Scope** drop down box (site2 or site1 in this case), and click on the switch icon. A screen pops up with the switch information.

Change the entry from *Leaf* to *Border Gateway* as shown in the image.



Likewise, update the role for BGW\_4 in site2 and then select site1 from the scope box and update roles for BGW\_1 and BGW\_2. This is required for configuring the EVPN Multi-Site feature through the DCNM GUI.

After completing the EVPN Multi-Site specific prerequisites, start EVPN Multi-Site configuration on BGW\_3 with extensions to the route server RS\_1.

## EVPN Multi-Site Extensions from BGW\_3 to RS\_1

From the Cisco DCNM Web Client, choose **Configure > LAN Fabric Provisioning > Network Deployment**. The LAN Fabric Provisioning page appears.

Click **Continue**. The **Select a Fabric** page comes up.

## Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled or create a new fabric.

SITE\_2

 Fabric Extension Settings


OR

+ Create a new fabric

Select **site2** from the drop-down box since you are configuring BGW *BGW\_3* on site2.

Click **Fabric Extension Settings** since the purpose of this task is to allow site2 to communicate to external fabrics through RS\_1 and RS\_2. The **Fabric Extension** screen comes up.

Fabric Extension ×

Inter-Fabric Connections Selected 0 / Total 0 

Type	Source Fabric	Source Device	Source Interface	Destination Fabric	Destination Device	Destination Interface	Configuration	Status

No data available


The **Inter-Fabric Connections** section lists previously created external connections from the BGWs on site2. Each line represents a physical or logical connection between a BGW in site2 and an external device in another fabric. For each connection, the source fabric, source device, source interface, destination fabric, destination device, and destination interface are listed along with the type of external connectivity. This section is empty as this is the first time you are adding an external connection.

*To extend the fabric through EVPN Multi-Site, you should first create an underlay extension and then an overlay extension.*

### Underlay Extension from BGW\_3 to RS\_1

Click on the + icon to add a new external connection. The **Add Inter-Fabric Connection** screen appears.

## Add Inter-Fabric Connections



* Extension Type	VRF_LITE	▼	
* Base Template	BorderBase_v1	▼	
* Extension Template	FabricSetup	▼	
* Source Fabric	SITE_2		
* Destination Fabric		▼	
* Source Device		▼	① VRF_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway"
* Source Interface		▼	
* Destination Device		▼	
* Destination Interface		▼	

Previous


Next

Save &amp; Deploy

Cancel

By default, VRF\_LITE is populated in the **Extension Type** field. Change the selection to MULTISITE\_UNDERLAY.

## Add Inter-Fabric Connections



* Extension Type	MULTISITE_UNDERLAY	▼	
* Base Template	BorderBase_v1	▼	
* Extension Template	MultiSiteUnderlaySetup	▼	
* Source Fabric	SITE_2		
* Destination Fabric		▼	
* Source Device		▼	① VRF_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway"
* Source Interface		▼	
* Destination Device		▼	
* Destination Interface		▼	

Previous

Next

Save &amp; Deploy

Cancel



**Base Template**—By default, the BorderBase\_v1 base template is populated. This template is a one-time configuration pushed to the BGW.

**Extension Template**—*MultiSiteUnderlaySetup* is a setup template that contains the configuration that will be generated and pushed to the BGW to setup the corresponding inter-fabric connection.

These templates are auto-populated with corresponding pre-packaged default templates based on your selection.

**Source Fabric**—This field is pre-populated with *site2* since the EVPN Multi-Site underlay connection is between BGW\_3 in site2 and RS\_1 in the CORE fabric.

**Destination Fabric**—Choose CORE.

**Source Device and Source Interface**—Choose BGW\_3 as the source device and an Ethernet interface that needs to be connected to RS\_3.

**Destination Device and Destination Interface**—Choose RS\_1 as the destination device and the Ethernet interface that connects to the BGW BGW\_3

Note that based on the selection of the source device and source interface, the destination information will be auto-populated based on Cisco Discovery Protocol information, if available. There is extra validation performed to ensure that the destination external device is indeed part of the destination fabric.

After filling up the Fabric Interconnect section, the screen looks like this.

Add Inter-Fabric Connections
✕

1 Fabric Interconnect

→

2 Define Variables

→

3 Preview & Deploy


● ● ●

<b>* Extension Type</b>	MULTISITE_UNDERLAY ▼	
<b>* Base Template</b>	▼	
<b>* Extension Template</b>	MultiSiteUnderlaySetup ▼	
<b>* Source Fabric</b>	SITE_2	
<b>* Destination Fabric</b>	CORE ▼	
<b>* Source Device</b>	BGW_3 ▼	ⓘ VRF_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway"
<b>* Source Interface</b>	Ethernet4/3  ▼	
<b>* Destination Device</b>	RS_1 ▼	
<b>* Destination Interface</b>	Ethernet5/5 ▼	

Previous
Next
Save & Deploy
Cancel

Click **Next** to go to the **Define Variables** section.

## Add Inter-Fabric Connections



▼ Network Profile

General
MULTISITE
<p>* IF_NAME <input type="text" value="Ethernet4/3"/> ?</p> <p>* Interface IP/Mask <input type="text" value="10.101.30.2/24"/> ?</p> <p>* Neighbor IP <input type="text" value="10.101.30.1"/> ?</p> <p>* NEIGHBOR_ASN <input type="text" value="65100"/> ?</p> <p>* Extension Type <input type="text" value="MULTISITE"/> ?</p>
<p>Previous   Next   Save &amp; Deploy   Cancel</p>

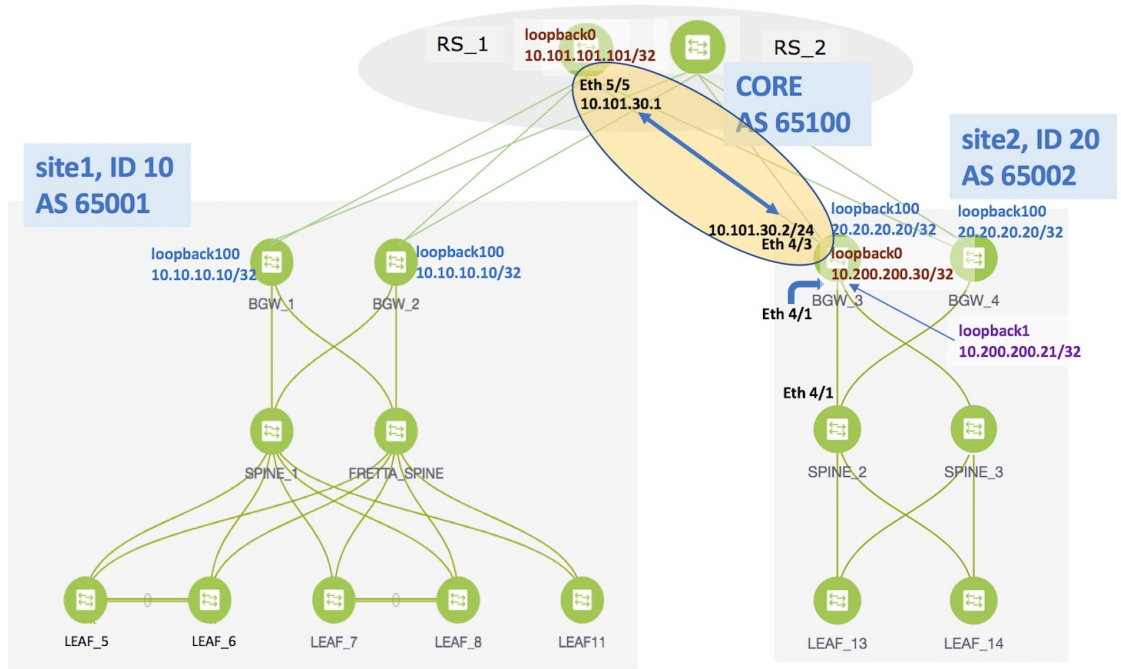
**IF\_NAME**—In this field, the interface name is auto-populated from the previous step.

**Interface IP\_MASK**—Fill up this field with the IP address of the BGW\_3 interface that connects to RS\_1.

**NEIGHBOR\_IP**—Fill up this field with the IP address of the RS\_1 interface that connects to BGW\_3.

**NEIGHBOR\_ASN**—In this field, the AS number of RS\_1 will be auto-populated.

The corresponding connection in the topology is displayed:




Click the **MULTISITE** tab.

While the **General** tab contains external connection details, this tab contains intra-fabric information such as fabric IGP, fabric facing Ethernet interface, etc.

The **MULTISITE** tab only appears the first time you create an EVPN Multi-Site underlay on a device, since the details remain the same for subsequent connections. The next time you create an EVPN Multi-Site underlay connection, only the General tab will be available.

## Add Inter-Fabric Connections



## Network Profile

General	
MULTISITE	
* Fabric Site ID	20 ?
* NVE Identifier	1 ?
* Fabric Interfaces	? E.g. e1/1-4, e2/2
* Multisite Loopback ID	100 ? [0-1023]
* MultiSite Loopback IP	? IPv4 address
* Routing Protocol	is-is ? Select IGP (ospf or is-is)
* IS-IS/OSPF Router ID	UNDERLAY ? String
* OSPF Area #	0 ? String

**Fabric Site ID**—This is the identification for the VXLAN BGP EVPN fabric site2 to which BGW\_3 belongs. When you configure the EVPN Multi-Site feature on BGW\_4 (or any other BGW on site2), the site ID will be 20. The site *site1* will be assigned with a unique ID.

**NVE Identifier**—This is the VXLAN overlay ID.

**Fabric Interfaces**—Fill up this field with the interfaces on BGW\_3 that connects to other intra-fabric device ports. Since Ethernet 4/1 connects to SPINE\_2 and Ethernet 4/2 connects to SPINE\_3 in the topology, the interfaces should be entered over here.


**Multisite Loopback ID and Multisite Loopback IP**—These are the loopback ID and IP address of this EVPN Multi-Site instance.

**Routing Protocol and Router ID**—This is the IGP and the IGP instance ID within the fabric. Note that, if the IGP used in your setup is OSPF, the field has to be updated to *OSPF*.

**OSPF AREA**—OSPF area ID within the fabric.

A fully filled screen looks like this.

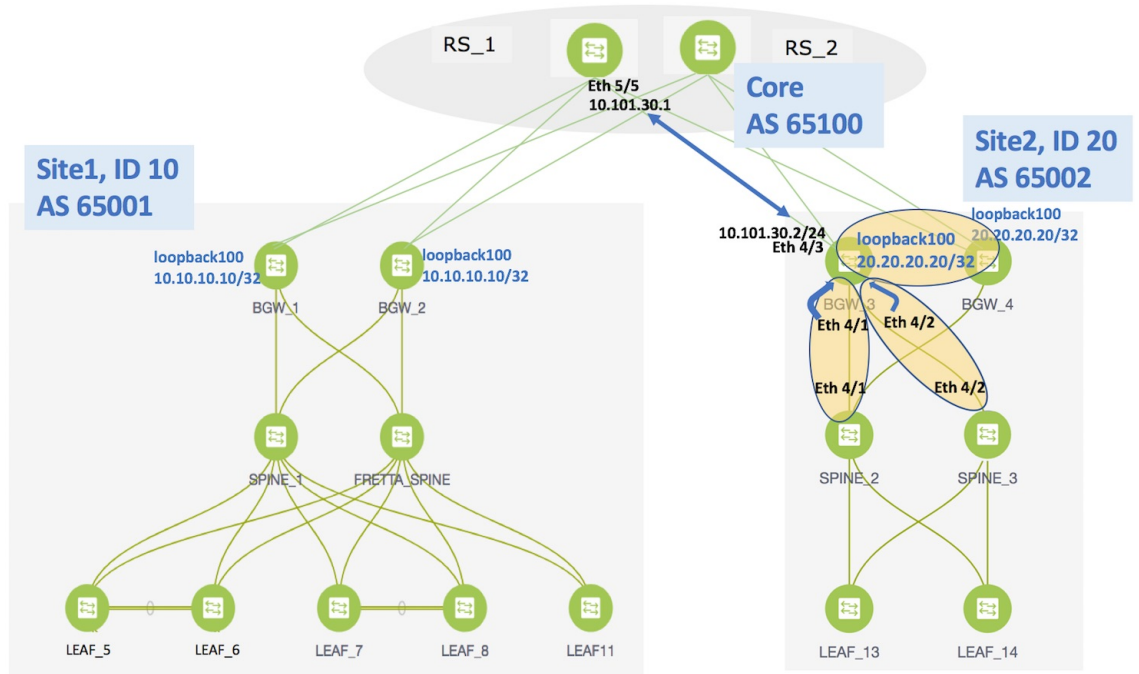
### Add Inter-Fabric Connections



#### Network Profile

General	
<b>MULTISITE</b>	
* Fabric Site ID	20
* NVE Identifier	1
* Fabric Interfaces	Eth4/1, Eth4/2
* Multisite Loopback ID	100
* MultiSite Loopback IP	20.20.20.20
* Routing Protocol	ospf
* IS-IS/OSPF Router ID	UNDERLAY
* OSPF Area #	0

The corresponding topology depiction is given below:



Now that all the information is filled in, click **Next** to go to the **Preview and Deploy** section.

## Add Inter-Fabric Connections

Switch:

Generated Configuration:

```
route-map RMAP-REDIST-DIRECT permit 10
  match tag 54321

evpn multisite border-gateway 20

interface loopback100
  description Used for EVPN Multi-Site
  ip address 20.20.20.20/32 tag 54321

  ip router ospf UNDERLAY area 0
  no shutdown

interface nve 1
  multisite border-gateway interface loopback100

interface e4/1
  evpn multisite fabric-tracking
```

Here, you can preview the configuration that will be deployed to BGW\_3. Note that no configuration will be pushed to the external device itself.

Click **Save and Deploy** to complete the task. This results in the configuration getting pushed to BGW\_3. The external connection will appear in the Fabric Extension screen.

Fabric Extension ×

Inter-Fabric Connections Selected 0 / Total 1

Type	Source Fabric	Source Device	Source Interf...	Destination ...	Destination De...	Destination Interf...	Configuration	Status
MULTISITE_UNDERLAY	SITE_2	BGW_3	Ethernet4/3	CORE	RS_1	Ethernet5/5	<a href="#">View Config</a>	DEPLOYED

The view doesn't auto-refresh, hence the refresh button on the top right part of the screen needs to be clicked to trigger refresh. You can check the status of the deployment (Deployment Pending, Deployed, Failed) in the **Status** column. In this case, the status changes from *Deployment Pending* to *Deployed* after you click on the refresh button.

In case of FAILED or UNDEPLOYMENT FAILED status, use the hyperlink in the **Status** column to check the error messages for failure.

To view the configurations, click on *View Config* in the **Configuration** field.

After the underlay configuration, you need to configure the overlay configuration from BGW\_3 to RS\_1 (the external device connected to BGW\_3), as shown in the next section.

## Overlay Extension from BGW\_3 to RS\_1




### Note

You can have multiple underlay connections to an external device but only one overlay connection from BGW\_3 to each external device.

In the **Fabric Extension** page, click on the + icon to add an external overlay connection. The **Add Inter-Fabric Connection** screen appears.

By default, VRF\_LITE is populated in the **Extension Type** field. Change the selection to MULTISITE\_OVERLAY. The screen changes accordingly.

### Add Inter-Fabric Connections



* Extension Type	<input type="text" value="MULTISITE_OVERLAY"/>	
* Base Template	<input type="text" value="BorderBase_v1"/>	
* Extension Template	<input type="text" value="MultiSiteOverlaySetup"/>	
* Source Fabric	<input type="text" value="SITE_2"/>	
* Destination Fabric	<input type="text"/>	
* Source Device	<input type="text"/>	<small>ⓘ VRF_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway"</small>
* Source Interface	<input type="text"/>	
* Destination Device	<input type="text"/>	
* Destination Interface	<input type="text"/>	

**Base Template**—*BorderBase\_v1* is auto-populated in this field. The *BorderBase\_v1* base template is a one-time configuration pushed to the BGW.

**Extension Template**—*MultiSiteOverlaySetup* is a setup template that contains the configuration that will be generated and pushed to the BGW to setup the corresponding inter-fabric connection. These templates are auto-populated with corresponding pre-packaged default templates based on your selection.

**Source Fabric**—This field is pre-populated with site2 since you are deploying the configurations in site2.

**Destination Fabric**—For the destination fabric, select the fabric that contains RS\_1, CORE.

**Source Device**—Choose BGW\_3 since the overlay connection is from BGW\_3 to RS\_1.

**Source Interface**—Typically, a loopback interface is created for the overlay. Choose the loopback interface amongst the 3 loopback interfaces you created as prerequisites. In this example, loopback0 is the BGP peer address.

**Destination Device**—Choose RS\_1 since the overlay connection is from BGW\_3 to RS\_1.

**Destination Interface**—Choose the destination interface. Choose the interface which is the BGP peer address. Note that the destination interface is not used in generating the configuration.

After filling up the Fabric Interconnect section, the screen looks like this.

### Add Inter-Fabric Connections ✕

1 Fabric Interconnect

→

2 Define Variables

→

3 Preview & Deploy

● ● ●


<b>* Extension Type</b>	MULTISITE_OVERLAY ▼	
<b>* Base Template</b>	▼	
<b>* Extension Template</b>	MultiSiteOverlaySetup ▼	
<b>* Source Fabric</b>	SITE_2 ▼	
<b>* Destination Fabric</b>	CORE ▼	
<b>* Source Device</b>	BGW_3 ▼	<small>ⓘ VRF_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway"</small>
<b>* Source Interface</b>	Loopback0 ▼	
<b>* Destination Device</b>	RS_1 ▼	
<b>* Destination Interface</b>	Loopback0 ▼	

---

Click **Next** to go to the **Define Variables** section.



## Add Inter-Fabric Connections



## ▼ Network Profile

General	
* IF_NAME	<input type="text" value="Loopback0"/> ?
* Overlay Neighbor IP	<input type="text" value="10.101.101.101"/> ? IPv4 address
* NEIGHBOR_ASN	<input type="text" value="65100"/> ?


**IF\_NAME**—In this field, the source interface is auto-populated from the previous step.

**Overlay Neighbor IP**—Enter the IP address on RS\_1 that the overlay peers with. This is typically a loopback address.

**NEIGHBOR\_ASN**—This field is populated with the RS\_1's AS Number.

Click **Next** to go to the **Preview and Deploy** section.

## Add Inter-Fabric Connections

Switch: 

Generated Configuration:

```
router bgp 65002
 neighbor 10.101.101.101 remote-as 65100
  update-source Loopback0
  ebgp-multihop 5
  peer-type fabric-external
  address-family l2vpn evpn
    send-community
    send-community extended
  rewrite-evpn-rt-asn
```

Here, you can preview the overlay configuration that will be deployed to BGW\_3. In this section, you can see that an overlay connection is being established from Loopback0 on BGW\_3 to the neighbor with AS Number 65100.

Note that no configuration will be pushed to the external device itself.

Click **Save and Deploy** to complete the task. This results in the configuration getting pushed to BGW\_3. The external connection will appear in the Fabric Extension screen.

Fabric Extension ×

Inter-Fabric Connections Selected 0 / Total 2

Type	Source Fa...	Source Device	Source Interface	Destination Fa...	Destination De...	Destination Int...	Configuration	Status
<input type="radio"/> MULTISITE_OVERLAY	SITE_2	BGW_3	Loopback0	CORE	RS_1	Loopback0	<a href="#">View Config</a>	DEPLOYED
<input type="radio"/> MULTISITE_UNDERLAY	SITE_2	BGW_3	Ethernet4/3	CORE	RS_1	Ethernet5/5	<a href="#">View Config</a>	DEPLOYED

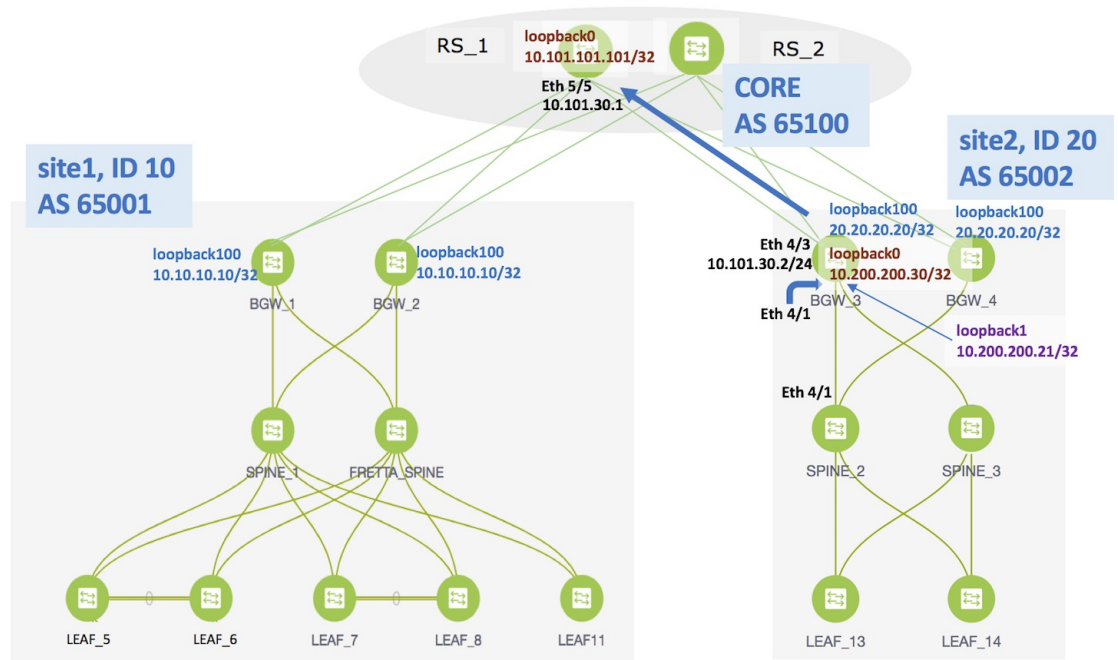
The view doesn't auto-refresh, hence the refresh button on the top right part of the screen needs to be clicked to trigger refresh. You can check the status of the deployment (Pending, Deployed, Failed) in the **Status** column. In case of FAILED or UNDEPLOYMENT FAILED status, use the hyperlink in the **Status** column to check the error messages for failure.

**Note**

Extensions will need to be deleted and then reconfigured in case of deployment failures. Currently there is no option to edit or redeploy an overlay or underlay extension.

## Other EVPN Multi-Site Configurations

At this stage, overlay and underlay EVPN Multi-Site configurations are provisioned on BGW\_3 towards RS\_1 (as shown by the arrow in the figure).



To complete EVPN Multi-Site configurations between site1 and site2 using DCNM, you should also configure as follows:

- **On site2**
  - EVPN Multi-Site configurations from BGW\_3 to RS\_2.
  - EVPN Multi-Site configurations from BGW\_4 to RS\_1 and RS\_2.
- **On site1**
  - EVPN Multi-Site configurations from BGW\_1 to RS\_1 and RS\_2.
  - EVPN Multi-Site configurations from BGW\_2 to RS\_1 and RS\_2.
- **On the route servers**
  - Apart from the DCNM provisioning on the BGWs of site1 and site2, you should enable appropriate configurations on RS\_1 and RS\_2 for connectivity between the route servers and the BGWs. Sample RS\_1 configurations are provided in the Appendix for your reference.

As noted earlier, the end-to-end Multi-Site configurations through DCNM Top-Down provisioning include these 2 steps:

(1) **Multi-Site configurations on the BGWs (BGW\_1, BGW\_2, BGW\_3 and BGW\_4).**

(2) **Deploying Networks and VRF Instances on the leaf switches and the BGWs.**

At this stage, the first step explanation is complete. In the next part of the document, the networks' configuration (second step), is explained. After appropriate network configurations on the leaf switches and BGWs, server traffic will flow across the 2 sites for the deployed and extended networks and VRFs.

## Deploying Networks and VRF Instances

Typically, you create a fabric in DCNM, then create and deploy networks and VRFs on devices within the fabric on leaf switches, and then configure the BGWs for external connectivity. Though the focus of the document is external connectivity with EVPN Multi-Site configurations on BGWs using DCNM, for completeness and right context, network deployment on the BGWs is explained in this section. When EVPN Multi-Site deployment is completed, server traffic from these networks and VRFs on site2 will pass through a BGW (BGW\_3 or BGW\_4) towards site1.

### Deploying Networks on the BGWs

*Before you begin*—In this scenario, we will deploy two networks in site2, *MyNetwork\_10000* and *MyNetwork\_10001*, on the BGWs BGW\_3 and BGW\_4. You should ensure that you have already deployed the networks that you want to extend to site1 on the leaf switches ( LEAF\_13 and LEAF\_14 in this case).

After deploying the 2 networks on the leaf switches and the BGWs, the networks will be extended to site1. To know how to create a new fabric, network, and VRF, see LAN Fabric Provisioning section in the [DCNM user guide](#).

In the Select a Fabric page, click the **Continue** button at the top right part of the screen.

(After Multi-Site overlays and underlays are created, the DCNM GUI automatically takes you to the Select a Fabric page).

## Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled or create a new fabric.

SITE\_2

 Fabric Extension Settings

OR


[+ Create a new fabric](#)

After clicking **Continue**, the **Networks** page comes up.

We will deploy two new networks *MyNetwork\_10000* and *MyNetwork\_10001* on the BGWs. To do that, select the checkboxes (in the extreme left column).

Fabric Selection > Network Selection > Network Deployment > VRF View | Continue

Fabric Selected: SITE\_2

Selected 2 / Total 4  

<input type="checkbox"/>	Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
<input checked="" type="checkbox"/>	MyNetwork_10000	10000	MyVRF_200000	10.1.10.1/24	10:1:A::1/48	UNDEPLOYED	
<input checked="" type="checkbox"/>	MyNetwork_10001	10001	MyVRF_200000	10.1.11.1/24		UNDEPLOYED	11
<input type="checkbox"/>	MyNetwork_10002	10002	MyVRF_200002	10.1.12.1/24	10:1:C::1/48	UNDEPLOYED	
<input type="checkbox"/>	MyNetwork_10003	10003	NA	10.1.13.1/24		UNDEPLOYED	

Click the **Continue** button at the top right part of the screen. The Network Deployment page (Topology View) comes up. You can deploy networks on multiple switches simultaneously, but with the same role. So, deploy the selected networks on the BGWs.

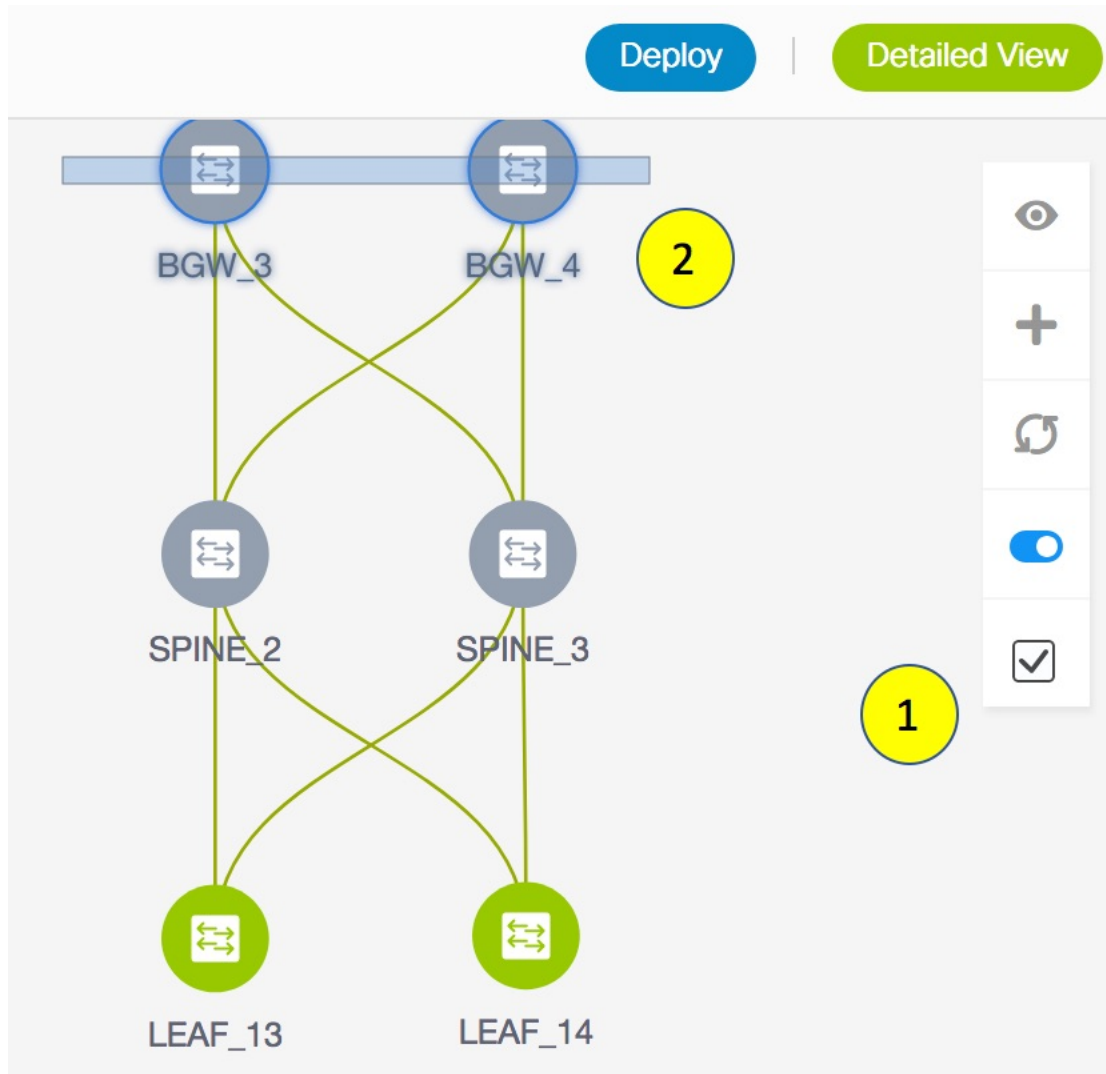


### Note

In the image, you can see that the networks are deployed on the leaf switches (green color indicates *deployed* status). Note that the color code (and hence the deployment state) on switches is contextual and specific to the selection. In this scenario, the deployed state only depicts that networks MYNetwork10000 and MYNetwork10001 are deployed on leaf switches LEAF\_13 and LEAF\_14. It does not display information about other (networks and VRFs) deployment instances, if any.

Select the multi-select check box at the bottom of the panel of options available at the right part of the page. (displayed as step 1 in the image).

Then, click your mouse (or track pad) and drag the cursor across BGW\_3 and BGW\_4. (step 2).



Immediately, the **Switches Deploy** screen (for networks) appears.

## Switches Deploy

*Fabric Name: SITE\_2*

MyNetwork\_10000

MyNetwork\_10001

*Deploy Options:* *Select the row and click on the cell to edit and save changes*

<input type="checkbox"/>	Switch	▲	VLAN	Extend	Status
<input type="checkbox"/>	BGW_3		10	MULTISITE	NA
<input type="checkbox"/>	BGW_4		10	MULTISITE	UNDEPLOYED

Save

A tab is displayed for each network. Click the checkbox next to the **Switch** column. Both the BGW check boxes will be selected automatically and the **Extension Details** section will appear at the bottom part of the screen.

In the **Extension Details** section, select the **Switch** checkbox (or ensure that you select the check box in each row) and click **Save** (bottom right part of your screen).

## Switches Deploy

Fabric Name: *SITE\_2*

MyNetwork\_10000

MyNetwork\_10001

## Deploy Options:

Select the row and click on the cell to edit and save changes

<input checked="" type="checkbox"/>	Switch	▲	VLAN	Extend	Status
<input checked="" type="checkbox"/>	BGW_3		10	MULTISITE	NA
<input checked="" type="checkbox"/>	BGW_4		10	MULTISITE	UNDEPLOYED

 Extension Details

<input checked="" type="checkbox"/>	Switch	▲	Type	IF_NAME
<input checked="" type="checkbox"/>	BGW_3		MULTISITE	Loopback0
<input checked="" type="checkbox"/>	BGW_3		MULTISITE	Loopback0
<input checked="" type="checkbox"/>	BGW_4		MULTISITE	Loopback0

Save

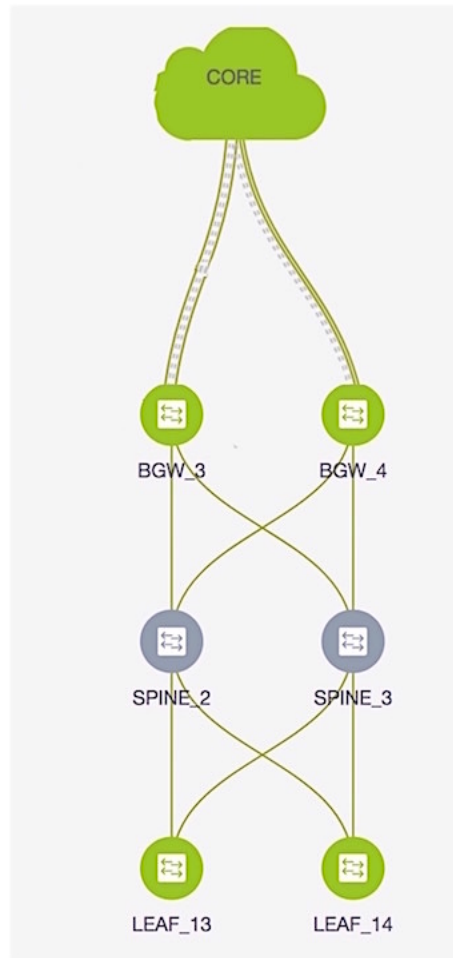
After saving the details in this screen, the Network Deployment screen (Topology view) appears.

BGW\_3 and BGW\_4 will be displayed in blue color, indicating pending deployment. If you want to check your configurations again, click on the Preview (eye) icon.

After you verify that the configurations that are generated from the profiles are correct for the selected switches, click the **Deploy** button (on the top right part of the screen) to deploy the MYNetwork10000 and MYNetwork10001 network configurations on BGW\_3 and BGW\_4.

DCNM shows the deployment status in the topology by highlighting the switch icons with different colors, yellow for *In Progress* and green for *Deployed*.



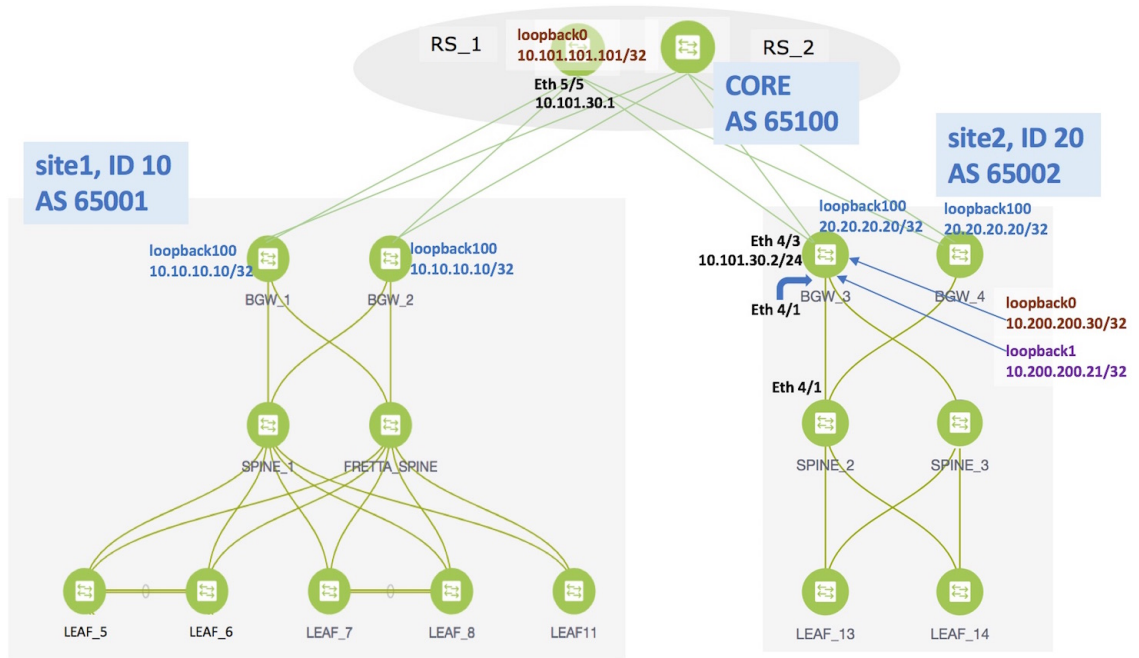


From the snapshot, you can see that the 2 networks MYNetwork10000 and MYNetwork10001 have been implemented on the leaf switches and BGWs.

After configurations in site2 are complete, configure the following in site1 too.

## Configurations in site1

Provision the networks MYNetwork10000 and MYNetwork10001 on the leaf switches (LEAF\_5, LEAF\_6, LEAF\_7, LEAF\_8, LEAF\_11) and the BGWs (BGW\_1 and BGW\_2).



As noted in the EVPN Multi-Site configuration section, enable the following for end-to-end configuration:

- Since DCNM does not provision configurations for RS\_1 and RS\_2 (devices directly connected to the BGWs), enable appropriate configurations on these devices.
- Configure the EVPN Multi-Site feature on the site1 BGWs (as explained in this document) so that server traffic from the 2 networks can flow to site2 and back.

## Additional References

Document Title and Link	Document Description
<a href="#">VXLAN EVPN Multi-Site Design and Deployment White Paper</a>	This document explains Multi-Site design and deployment in detail.
<a href="#">Configuring VXLAN EVPN Multi-Site</a>	This document explains manual configurations for the Multi-Site solution.

# Appendix

## Route Server Configurations

**RS\_1 configuration example for the overlay**—The following configurations are enabled on RS1, and reproduced here for reference.


**Note**

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

```
switch(config)#

route-map ALL-PATHS permit 100
  set path-selection all advertise
route-map RMAP-REDIST-DIRECT permit 10
  match tag 12345
route-map UNCHANGED permit 10
  set ip next-hop unchanged

switch(config)#

interface loopback0
  ip address 10.101.101.101/32 tag 12345
line vty
router bgp 65100
  router-id 10.101.101.101
  address-family ipv4 unicast
    redistribute direct route-map RMAP-REDIST-DIRECT
    maximum-paths 4
    additional-paths send
    additional-paths receive
    additional-paths selection route-map ALL-PATHS
  address-family l2vpn evpn
    retain route-target all
  template peer OVERLAY-PEERING
    update-source loopback0
    ebgp-multihop 5
  address-family l2vpn evpn
    send-community both
    route-map UNCHANGED out
  neighbor 10.100.100.10
    inherit peer OVERLAY-PEERING
    remote-as 65001
  address-family l2vpn evpn
    rewrite-evpn-rt-asn
    route-map UNCHANGED out
  neighbor 10.100.100.20
    inherit peer OVERLAY-PEERING
    remote-as 65001
  address-family l2vpn evpn
    rewrite-evpn-rt-asn
    route-map UNCHANGED out
  neighbor 10.101.11.2
    remote-as 65101
    update-source Ethernet5/1
  address-family ipv4 unicast
    next-hop-self
  neighbor 10.101.12.2
    remote-as 65101
    update-source Ethernet5/2
  address-family ipv4 unicast
```

```
        next-hop-self
neighbor 10.101.13.2
  remote-as 65102
  update-source Ethernet5/3
  address-family ipv4 unicast
    next-hop-self
neighbor 10.101.14.2
  remote-as 65102
  update-source Ethernet5/4
  address-family ipv4 unicast
    next-hop-self
neighbor 10.101.30.2
  remote-as 65002
  update-source Ethernet5/5
  address-family ipv4 unicast
    next-hop-self
neighbor 10.101.40.2
  remote-as 65002
  update-source Ethernet5/6
  address-family ipv4 unicast
    next-hop-self
neighbor 10.200.200.30
  remote-as 65002
  update-source loopback0
  ebgp-multihop 5
  address-family l2vpn evpn
    rewrite-evpn-rt-asn
    send-community both
    route-map UNCHANGED out
neighbor 10.200.200.40
  remote-as 65002
  update-source loopback0
  ebgp-multihop 5
  address-family l2vpn evpn
    rewrite-evpn-rt-asn
    send-community both
    route-map UNCHANGED out
```