Configuring Multi-Site

This chapter contains the following sections:

About VXLAN EVPN Multi-Site

The VXLAN EVPN Multi-Site solution interconnects two or more BGP-based Ethernet VPN (EVPN) sites/fabrics (overlay domains) in a scalable fashion over an IP-only network. This solution uses border gateways (BGWs) in anycast or vPC mode to terminate and interconnect two sites. The BGWs provide the network control boundary that is necessary for traffic enforcement and failure containment functionality.

In the BGP control plane, BGP sessions between the BGWs rewrite the next hop information of EVPN routes and reoriginate them.

VXLAN Tunnel Endpoints (VTEPs) are only aware of their overlay domain internal neighbors, including the BGWs. All routes external to the fabric have a next hop on the BGWs for Layer 2 and Layer 3 traffic.

The BGW is the node that interacts with nodes within a site and with nodes that are external to the site. For example, in a leaf-spine data center fabric, it can be a leaf, a spine, or a separate device acting as a gateway to interconnect the sites.

The VXLAN EVPN Multi-Site feature can be conceptualized as multiple site-local EVPN control planes and IP forwarding domains interconnected via a single common EVPN control and IP forwarding domain. Every EVPN node is identified with a unique site-scope identifier. A site-local EVPN domain consists of EVPN nodes with the same site identifier. BGWs on one hand are also part of the site-specific EVPN domain and on the other hand a part of a common EVPN domain to interconnect with BGWs from other sites. For a given site, these BGWs facilitate site-specific nodes to visualize all other sites to be reachable only via them. This means:

- Site-local bridging domains are interconnected only via BGWs with bridging domains from other sites.
- Site-local routing domains are interconnected only via BGWs with routing domains from other sites.

- Site-local flood domains are interconnected only via BGWs with flood domains from other sites.

Selective Advertisement is defined as the configuration of the per-tenant information on the BGW. Specifically, this means IP VRF or MAC VRF (EVPN instance). In cases where external connectivity (VRF-lite) and EVPN Multi-Site coexist on the same BGW, the advertisements are always enabled.

Guidelines and Limitations for VXLAN EVPN Multi-Site

VXLAN EVPN Multi-Site has the following configuration guidelines and limitations:

- Cisco Nexus 9332C and 9364C are supported as border gateways.
- VXLAN EVPN Multi-Site is not supported on Cisco Nexus 9500 platform switches with -R line cards.
- Support for VXLAN EVPN Multi-Site functionality on the Cisco Nexus N9K-C9336C-FX and N9K-C93240YC-FX2 is added. N9K-C9348GC-FXP does not support VXLAN EVPN Multi-Site functionality.
- VXLAN EVPN Multi-Site and Tenant Routed Multicast (TRM) is supported between source and receivers deployed in the same site.
- The Multi-Site border gateway allows the co-existence of Multi-Site extensions (Layer 2 unicast/multicast and Layer 3 unicast) as well as Layer 3 unicast and multicast external connectivity.
- The following switches support VXLAN EVPN Multi-Site:
  - Cisco Nexus 9300-EX, 9300-FX, and 9500 platform switches with X9700-EX line cards

Note The Cisco Nexus 9348GC-FXP switch does not support VXLAN EVPN Multi-Site functionality.

  - Cisco Nexus 9396C switch and Cisco Nexus 9500 platform switches with X9700-FX line cards
  - Cisco Nexus 9336C-FX2 switch
- In a VXLAN EVPN multisite deployment, when you use the ttag feature, make sure that the ttag is stripped (**ttag-strip**) on BGW's DCI interfaces that connect to the cloud. To elaborate, if the ttag is attached to non-Nexus 9000 devices that do not support ether-type 0x8905, stripping of ttag is required. However, BGW back-to-back model of DCI does not require ttag stripping.
- The number of border gateways per site is limited to four.
- Beginning with Cisco NX-OS Release 9.2(1), Border Gateways (BGWs) in a vPC topology are supported.
- Support for Multicast Flood Domain between inter-site/fabric border gateways is not supported.
- Multicast Underlay between sites is not supported.
- iBGP EVPN Peering between border gateways of different fabrics/sites is not supported.
- The **peer-type fabric-external** command configuration is required only for VXLAN Multi-site BGWs (this command must not be used when peering with non-Cisco equipment).

Note The **peer-type fabric-external** command configuration is not required for pseudo BGWs.

- Anycast mode can support up to four border gateway's per site.
- Anycast mode can only support Layer 3 services attached to local interfaces.
- In Anycast mode, BUM is replicated to each border-leaf and DF election, between border leafs of a particular site decides which border leaf would forward the traffic inter-site traffic (Fabric to DCI and vice versa) for that site.
- In Anycast mode, all the Layer 3 services are advertised in BGP via EVPN Type-5 routes with their physical IP as the next hop.
- vPC mode can support only two border gateways.
- vPC mode can support both Layer 2 hosts and Layer 3 services on local interfaces.
- In vPC mode, BUM is replicated to either of the border-gateway’s for traffic coming from external site and hence both the border gateways are forwarders for site external to site internal (DCI to Fabric) direction.
- In vPC mode, BUM is replicated to either of the border gateways for traffic coming from the local site leaf for a VLAN using Ingress Replication (IR) underlay. Both border gateways are forwarders for site internal to site external (Fabric to DCI) direction for VLANs using the IR underlay.
- In vPC mode, BUM is replicated to both border gateways for traffic coming from the local site leaf for a VLAN using the multicast underlay. Therefore, a decapper/forwarder election happens and the decapsulation winner/forwarder only forwards the site-local traffic to external site border-gateways for VLANs using the multicast underlay.
- In vPC mode, all the Layer 3 services/attachments are advertised in BGP via EVPN Type-5 routes with their virtual IP as next hop. If the VIP/PIP feature is configured, they are advertised with PIP as the next hop.
- If different Anycast Gateway MAC addresses are configured across sites, ARP suppression must be enabled for all VLANs that have been extended.
- Bind NVE to a loopback address that is separate from loopback addresses that are required by Layer 3 protocols. A best practice is to use a dedicated loopback address for the NVE source interface (PIP VTEP) and Multi-Site source interface (anycast and virtual IP VTEP).
- PIM BiDir is not supported for fabric underlay multicast replication with VXLAN Multi-Site.
- PIM is not supported on multisite VXLAN DCI links.
- FEX is not supported on a vPC BGW and Anycast BGW.
- To improve the convergence in case of fabric link failure and avoid issues in case of fabric link flapping, ensure to configure multi-hop BFD between loopbacks of spines and BGWs.

  In the specific scenario where a BGW node becomes completely isolated from the fabric due to all its fabric links failing, the use of multi-hop BFD ensures that the BGP sessions between the spines and the isolated BGW can be immediately brought down, without relying on the configured BGP hold-time value.

Enabling VXLAN EVPN Multi-Site

This procedure enables the VXLAN EVPN Multi-Site feature. Multi-Site is enabled on the BGWs only. The site-id must be the same on all BGWs in the fabric/site.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal`	Enters global configuration mode.
Step 2	**evpn multisite border-gateway** *ms-id* **Example:** `switch(config)# evpn multisite border-gateway 100`	Configures the site ID for a site/fabric. The range of values for *ms-id* is 1 to 2,814,749,767,110,655. The *ms-id* must be the same in all BGWs within the same fabric/site.
Step 3	**interface nve 1** **Example:** `switch(config-evpn-msite-bgw)# interface nve 1`	Creates a VXLAN overlay interface that terminates VXLAN tunnels. **Note** Only one NVE interface is allowed on the switch.
Step 4	**source-interface loopback** *src-if* **Example:** `switch(config-if-nve)# source-interface loopback 0`	The source interface must be a loopback interface that is configured on the switch with a valid /32 IP address. This /32 IP address must be known by the transient devices in the transport network and the remote VTEPs. This requirement is accomplished by advertising it through a dynamic routing protocol in the transport network.
Step 5	**host-reachability protocol bgp** **Example:** `switch(config-if-nve)# host-reachability protocol bgp`	Defines BGP as the mechanism for host reachability advertisement.
Step 6	**multisite border-gateway interface loopback** *vi-num* **Example:** `switch(config-if-nve)# multisite border-gateway interface loopback 100`	Defines the loopback interface used for the BGW virtual IP address (VIP). The border-gateway interface must be a loopback interface that is configured on the switch with a valid /32 IP address. This /32 IP address must be known by the transient devices in the transport network and the remote VTEPs. This requirement is accomplished by advertising it through a dynamic routing protocol in the transport network. This loopback must be

	Command or Action	Purpose
		different than the source interface loopback. The range of *vi-num* is from 0 to 1023.
Step 7	**no shutdown** **Example:** `switch(config-if-nve)# no shutdown`	Negates the **shutdown** command.
Step 8	**exit** **Example:** `switch(config-if-nve)# exit`	Exits the NVE configuration mode.
Step 9	**interface loopback** *loopback-number* **Example:** `switch(config)# interface loopback 0`	Configures the loopback interface.
Step 10	**ip address** *ip-address* **Example:** `switch(config-if)# ip address 198.0.2.0/32`	Configures the IP address for the loopback interface.

Multi-Site with vPC Support

About Multi-Site with vPC Support

The BGWs can be in a vPC complex. In this case, it is possible to support dually-attached directly-connected hosts that might be bridged or routed as well as dually-attached firewalls or service attachments. The vPC BGWs have vPC-specific multihoming techniques and do not rely on EVPN Type 4 routes for DF election or split horizon.

Guidelines and Limitations for Multi-Site with vPC Support

Multi-Site with vPC support has the following configuration guidelines and limitations:

- 4000 VNIs for vPC are not supported.
- For BUM with continued VIP use, the MCT link is used as transport upon core isolation or fabric isolation, and for unicast traffic in fabric isolation.
- The routes to remote Multisite BGW loopback addresses must always prioritize the DCI link path over the iBGP protocol between vPC Border Gateway switches configured using the backup SVI. The backup SVI should be used strictly in the event of a DCI link failure.

Configuring Multi-Site with vPC Support

This procedure describes the configuration of Multi-Site with vPC support:

- Configure vPC domain.
- Configure port channels.
- Configuring vPC Peer Link.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal`	Enters global configuration mode.
Step 2	**feature vpc** **Example:** `switch(config)# feature vpc`	Enables vPCs on the device.
Step 3	**feature interface-vlan** **Example:** `switch(config)# feature interface-vlan`	Enables the interface VLAN feature on the device.
Step 4	**feature lacp** **Example:** `switch(config)# feature lacp`	Enables the LACP feature on the device.
Step 5	**feature pim** **Example:** `switch(config)# feature pim`	Enables the PIM feature on the device.
Step 6	**feature ospf** **Example:** `switch(config)# feature ospf`	Enables the OSPF feature on the device.
Step 7	**ip pim rp-address** *address* **group-list** *range* **Example:** `switch(config)# ip pim rp-address 100.100.100.1 group-list 224.0.0/4`	Defines a PIM RP address for the underlay multicast group range.
Step 8	**vpc domain** *domain-id* **Example:** `switch(config)# vpc domain 1`	Creates a vPC domain on the device and enters vpn-domain configuration mode for configuration purposes. There is no default. The range is from 1 to 1000.

	Command or Action	Purpose
Step 9	**peer switch** **Example:** `switch(config-vpc-domain)# peer switch`	Defines the peer switch.
Step 10	**peer gateway** **Example:** `switch(config-vpc-domain)# peer gateway`	Enables Layer 3 forwarding for packets destined to the gateway MAC address of the vPC.
Step 11	**peer-keepalive destination** *ip-address* **Example:** `switch(config-vpc-domain)# peer-keepalive destination 172.28.230.85`	Configures the IPv4 address for the remote end of the vPC peer-keepalive link. **Note** The system does not form the vPC peer link until you configure a vPC peer-keepalive link. The management ports and VRF are the defaults.
Step 12	**ip arp synchronize** **Example:** `switch(config-vpc-domain)# ip arp synchronize`	Enables IP ARP synchronize under the vPC domain to facilitate faster ARP table population following device reload.
Step 13	**ipv6 nd synchronize** **Example:** `switch(config-vpc-domain)# ipv6 nd synchronize`	Enables IPv6 ND synchronization under the vPC domain to facilitate faster ND table population following device reload.

Step 14	Create the vPC peer-link.
	Example:
	`switch(config)# interface port-channel 1`
	`switch(config)# switchport`
	`switch(config)# switchport mode trunk`
	`switch(config)# switchport trunk allowed vlan 1,10,100-200`
	`switch(config)# mtu 9216`
	`switch(config)# vpc peer-link`
	`switch(config)# no shut`
	
	`switch(config)# interface Ethernet 1/1, 1/21`
	`switch(config)# switchport`
	`switch(config)# mtu 9216`
	`switch(config)# channel-group 1 mode active`
	`switch(config)# no shutdown`
		Creates the vPC peer-link port-channel interface and adds two member interfaces to it.

Step 15	**system nve infra-vlans** *range* **Example:**	Defines a non-VXLAN-enabled VLAN as a backup routed path.

	Command or Action	Purpose
	`switch(config)# system nve infra-vlans 10`	
Step 16	**vlan** *number* **Example:** `switch(config)# vlan 10`	Creates the VLAN to be used as an infra-VLAN.

Step 17	Create the SVI.
	Example:
	`switch(config)# interface vlan 10`
	`switch(config)# ip address 10.10.10.1/30`
	`switch(config)# ip router ospf process UNDERLAY area 0`
	`switch(config)# ip pim sparse-mode`
	`switch(config)# no ip redirects`
	`switch(config)# mtu 9216`
	`switch(config)# no shutdown`
		Creates the SVI used for the backup routed path over the vPC peer-link.

Step 18	(Optional) **delay restore interface-vlan** *seconds* **Example:** `switch(config-vpc-domain)# delay restore interface-vlan 45`	Enables the delay restore timer for SVIs. We recommend tuning this value when the SVI/VNI scale is high. For example, when the SCI count is 1000, we recommend that you set the delay restore to 45 seconds.
Step 19	**evpn multisite border-gateway** *ms-id* **Example:** `switch(config)# evpn multisite border-gateway 100`	Configures the site ID for a site/fabric. The range of values for *ms-id* is 1 to 281474976710655. The *ms-id* must be the same in all BGWs within the same fabric/site.
Step 20	**interface nve 1** **Example:** `switch(config-evpn-msite-bgw)# interface nve 1`	Creates a VXLAN overlay interface that terminates VXLAN tunnels. **Note** Only one NVE interface is allowed on the switch.
Step 21	**source-interface loopback** *src-if* **Example:** `switch(config-if-nve)# source-interface loopback 0`	Defines the source interface, which must be a loopback interface with a valid /32 IP address. This /32 IP address must be known by the transient devices in the transport network and the remote VTEPs. This requirement is accomplished by advertising the address through a dynamic routing protocol in the transport network.
Step 22	**host-reachability protocol bgp** **Example:** `switch(config-if-nve)# host-reachability protocol bgp`	Defines BGP as the mechanism for host reachability advertisement.

	Command or Action	Purpose
Step 23	**multisite border-gateway interface loopback** *vi-num* **Example:** `switch(config-if-nve)# multisite border-gateway interface loopback 100`	Defines the loopback interface used for the BGW virtual IP address (VIP). The BGW interface must be a loopback interface that is configured on the switch with a valid /32 IP address. This /32 IP address must be known by the transient devices in the transport network and the remote VTEPs. This requirement is accomplished by advertising the address through a dynamic routing protocol in the transport network. This loopback must be different than the source interface loopback. The range of *vi-num* is from 0 to 1023.
Step 24	**no shutdown** **Example:** `switch(config-if-nve)# no shutdown`	Negates the **shutdown** command.
Step 25	**exit** **Example:** `switch(config-if-nve)# exit`	Exits the NVE configuration mode.
Step 26	**interface loopback** *loopback-number* **Example:** `switch(config)# interface loopback 0`	Configures the loopback interface.
Step 27	**ip address** *ip-address* **Example:** `switch(config-if)# ip address 198.0.2.0/32`	Configures the primary IP address for the loopback interface.
Step 28	**ip address** *ip-address* **secondary** **Example:** `switch(config-if)# ip address 198.0.2.1/32 secondary`	Configures the secondary IP address for the loopback interface.
Step 29	**ip pim sparse-mode** **Example:** `switch(config-if)# ip pim sparse-mode`	Configures PIM sparse mode on the loopback interface.

Configuring Peer Link as Transport in Case of Link Failure

This procedure describes the configuration of an SVI interface configured with a high IGP cost to ensure it is only used as a backup link.

Note This configuration is required to use the peer link as a backup link during fabric and/or DCI link failures.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal`	Enters global configuration mode.
Step 2	**system nve infra-vlans** *vlan-range* **Example:** `switch(config)# system nve infra-vlans 10`	Specifies VLANs used by all SVI interfaces for uplink and vPC peer-links in VXLAN as infra-VLANs. You should not configure certain combinations of infra-VLANs. For example, 2 and 514, 10 and 522, which are 512 apart.
Step 3	**interface** *vlan-id* **Example:** `switch(config)# interface vlan10`	Configures the interface.
Step 4	**no shutdown** **Example:** `switch(config-if)# no shutdown`	Negates the **shutdown** command.
Step 5	**mtu** *value* **Example:** `switch(config-if)# mtu 9216`	Sets the maximum transmission unit (MTU).
Step 6	**no ip redirects** **Example:** `switch(config-if)# no ip redirects`	Prevents the device from sending redirects.
Step 7	**ip address** *ip-address/length* **Example:** `switch(config-if)# ip address 35.1.1.2/24`	Configures an IP address for this interface.
Step 8	**no ipv6 redirects** **Example:** `switch(config-if)# no ipv6 redirects`	Disables the ICMP redirect messages on BFD-enabled interfaces.
Step 9	**ip ospf cost** *cost* **Example:** `switch(config-if)# ip ospf cost 100`	Configures the OSPF cost metric for this interface.

	Command or Action	Purpose
Step 10	**ip ospf network point-to-point** **Example:** `switch(config-if)# ip ospf network point-to-point`	Specifies the OSPF point-to-point network.
Step 11	**ip router ospf** *instance* **area** *area-number* **Example:** `switch(config-if)# ip router ospf 1 area 0.0.0.0`	Configures the routing process for the IP on an interface and specifies an area.
Step 12	**ip pim sparse-mode** **Example:** `switch(config-if)# ip pim sparse-mode`	Configures sparse-mode PIM on an interface.

Verifying the Multi-Site with vPC Support Configuration

To display Multi-Site with vPC support information, enter one of the following commands:

show vpc brief	Displays general vPC and CC status.
show vpc consistency-parameters global	Displays the status of those parameters that must be consistent across all vPC interfaces.
show vpc consistency-parameters vni	Displays configuration information for VNIs under the NVE interface that must be consistent across both vPC peers.

Output example for the **show vpc brief** command:

```
switch# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 1
Peer status                       : peer adjacency formed ok     (<--- peer up)
vPC keep-alive status             : peer is alive
Configuration consistency status  : success (<----- CC passed)
Per-vlan consistency status       : success                       (<---- per-VNI CCpassed)
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 1
Peer Gateway                      : Enabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 240s)
Delay-restore status              : Timer is off.(timeout = 30s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled
[...]
```

Output example for the **show vpc consistency-parameters global** command:

```
switch# show vpc consistency-parameters global

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
[...]
Nve1 Adm St, Src Adm St,    1     Up, Up, 2.1.44.5, CP,  Up, Up, 2.1.44.5, CP,
Sec IP, Host Reach, VMAC          TRUE, Disabled,        TRUE, Disabled,
Adv, SA,mcast l2, mcast           0.0.0.0, 0.0.0.0,      0.0.0.0, 0.0.0.0,
l3, IR BGP,MS Adm St, Reo         Disabled, Up,          Disabled, Up,
                                  200.200.200.200        200.200.200.200
[...]
```

Output example for the **show vpc consistency-parameters vni** command:

```
switch(config-if-nve-vni)# show vpc consistency-parameters vni

    Legend:
        Type 1 : vPC will be suspended in case of mismatch

Name                        Type  Local Value            Peer Value
-------------               ----  ---------------------- -----------------------
Nve1 Vni, Mcast, Mode,      1     11577, 234.1.1.1,      11577, 234.1.1.1,
Type, Flags                       Mcast, L2, MS IR       Mcast, L2, MS IR
Nve1 Vni, Mcast, Mode,      1     11576, 234.1.1.1,      11576, 234.1.1.1,
Type, Flags                       Mcast, L2, MS IR       Mcast, L2, MS IR
[...]
```

Configuring VNI Dual Mode

This procedure describes the configuration of the BUM traffic domain for a given VLAN. Support exists for using multicast or ingress replication inside the fabric/site and ingress replication across different fabrics/sites.

Note If you have multiple VRFs and only one is extended to ALL leaf switches, you can add a dummy loopback to that one extended VRF and advertise through BGP. Otherwise, you'll need to check how many VRFs are extended and to which switches, and then add a dummy loopback to the respective VRFs and advertise them as well. Therefore, use the **advertise-pip** command to prevent potential user errors in the future.

For more information about configuring multicast or ingress replication for a large number of VNIs, see Example of VXLAN BGP EVPN (EBGP).

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal`	Enters global configuration mode.
Step 2	**interface nve 1** **Example:**	Creates a VXLAN overlay interface that terminates VXLAN tunnels.

	Command or Action	Purpose
	`switch(config)# interface nve 1`	**Note** Only one NVE interface is allowed on the switch.
Step 3	**member vni** *vni-range* **Example:** `switch(config-if-nve)# member vni 200`	Configures the virtual network identifier (VNI). The range for *vni-range* is from 1 to 16,777,214. The value of *vni-range* can be a single value like 5000 or a range like 5001-5008. **Note** Enter one of the Step 4 or Step 5 commands.
Step 4	**mcast-group** *ip-addr* **Example:** `switch(config-if-nve-vni)# mcast-group 255.0.4.1`	Configures the NVE Multicast group IP prefix within the fabric.
Step 5	**ingress-replication protocol bgp** **Example:** `switch(config-if-nve-vni)# ingress-replication protocol bgp`	Enables BGP EVPN with ingress replication for the VNI within the fabric.
Step 6	**multisite ingress-replication** **Example:** `switch(config-if-nve-vni)# multisite ingress-replication`	Defines the Multi-Site BUM replication method for extending the Layer 2 VNI.

Configuring Fabric/DCI Link Tracking

This procedure describes the configuration to track all DCI-facing interfaces and site internal/fabric facing interfaces. Tracking is mandatory and is used to disable reorigination of EVPN routes either from or to a site if all the DCI/fabric links go down.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal`	Enters global configuration mode.
Step 2	**interface ethernet** *port* **Example:** `switch(config)# interface ethernet1/1`	Enters interface configuration mode for the DCI or fabric interface. **Note** Enter one of the following commands in Step 3 or Step 4.

	Command or Action	Purpose
Step 3	**evpn multisite dci-tracking** **Example:** `switch(config-if)# evpn multisite dci-tracking`	Configures DCI interface tracking.
Step 4	(Optional) **evpn multisite fabric-tracking** **Example:** `switch(config-if)# evpn multisite fabric-tracking`	Configures EVPN Multi-Site fabric tracking. The **evpn multisite fabric-tracking** is mandatory for anycast BGWs and vPC BGW fabric links.
Step 5	**ip address** *ip-addr* **Example:** `switch(config-if)# ip address 192.1.1.1`	Configures the IP address.
Step 6	**no shutdown** **Example:** `switch(config-if)# no shutdown`	Negates the **shutdown** command.

Configuring Fabric External Neighbors

This procedure describes the configuration of fabric external/DCI neighbors for communication to other site/fabric BGWs.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal`	Enters global configuration mode.
Step 2	**router bgp** *as-num* **Example:** `switch(config)# router bgp 100`	Configures the autonomous system number. The range for *as-num* is from 1 to 4,294,967,295.
Step 3	**neighbor** *ip-addr* **Example:** `switch(config-router)# neighbor 100.0.0.1`	Configures a BGP neighbor.
Step 4	**remote-as** *value* **Example:** `switch(config-router-neighbor)# remote-as 69000`	Configures remote peer's autonomous system number.

	Command or Action	Purpose
Step 5	**peer-type fabric-external** **Example:** `switch(config-router-neighbor)# peer-type fabric-external`	Enables the next hop rewrite for Multi-Site. Defines site external BGP neighbors for EVPN exchange. The default for **peer-type** is **fabric-internal**. **Note** The **peer-type fabric-external** command is required only for VXLAN Multi-Site BGWs. It is not required for pseudo BGWs.
Step 6	**address-family l2vpn evpn** **Example:** `switch(config-router-neighbor)# address-family l2vpn evpn`	Configures the address family Layer 2 VPN EVPN under the BGP neighbor.
Step 7	**rewrite-evpn-rt-asn** **Example:** `switch(config-router-neighbor)# rewrite-evpn-rt-asn`	Rewrites the route target (RT) information to simplify the MAC-VRF and IP-VRF configuration. BGP receives a route, and as it processes the RT attributes, it checks if the AS value matches the peer AS that is sending that route and replaces it. Specifically, this command changes the incoming route target's AS number to match the BGP-configured neighbor's remote AS number. You can see the modified RT value in the receiver router.