



BGP PBB EVPN Route Reflector Support

The BGP PBB EVPN Route Reflector Support feature provides Border Gateway Protocol (BGP) route reflector functionality for Ethernet VPN (EVPN) and provider backbone bridging (PBB) EVPN of Layer 2 VPN address family. EVPN enables customer MAC addresses as routable addresses and distributes them in BGP to avoid any data plane MAC address learning over the Multiprotocol Label Switching (MPLS) core network. The route reflector is enhanced to store the received EVPN updates without configuring EVPN explicitly on the route reflector and then advertises these updates to other provider edge (PE) devices so that the PEs do not need to have a full mesh of BGP sessions.

- [Prerequisites for BGP PBB EVPN Route Reflector Support, on page 1](#)
- [Information About BGP PBB EVPN Route Reflector Support, on page 1](#)
- [How to Configure BGP PBB EVPN Route Reflector Support, on page 2](#)
- [Configuration Examples for BGP PBB EVPN Route Reflector Support, on page 4](#)
- [Additional References for BGP PBB EVPN Route Reflector Support, on page 5](#)
- [Feature Information for BGP PBB EVPN Route Reflector Support, on page 5](#)

Prerequisites for BGP PBB EVPN Route Reflector Support

- Before you configure the BGP PBB EVPN Route Reflector Support feature, you must configure the RT filter unicast address family type to support for EVPN address family. For more information, see the "Configuring BGP: RT Constrained Route Distribution" module in the *IP Routing: BGP Configuration Guide*.
- The EVPN Subsequent Address Family Identifier (SAFI) needs to be enabled globally before you enable it under the BGP neighbor.

Information About BGP PBB EVPN Route Reflector Support

EVPN Overview

Ethernet VPN (EVPN) allows Multiprotocol Label Switching (MPLS) networks to provide multipoint Layer 2 VPN (L2VPN) services.

In EVPN, the customer MAC addresses are learned in the data plane over links connecting customer devices (CE) to the provider edge (PE) devices. The MAC addresses are then distributed over the Multiprotocol Label

Switching (MPLS) core network using Border Gateway Protocol (BGP) with an MPLS label identifying the service instance. A single MPLS label per EVPN instance is sufficient as long as the receiving PE device performs a MAC lookup in the disposition path. Receiving PE devices inject these routable MAC addresses into their Layer 2 routing information base (RIB) and forwarding information base (FIB) along with their associated adjacencies.

EVPN defines a BGP Network Layer Reachability Information (NLRI) that advertises different route types and route attributes. The EVPN NLRI is carried in BGP using BGP multiprotocol extensions with an Address Family Identifier (AFI) and a Subsequent Address Family Identifier (SAFI). BGP drops unsupported route types and does not propagate them to neighbors.

BGP EVPN Autodiscovery Support on Route Reflector

By default, routes received from an internal BGP (iBGP) peer are not sent to another iBGP peer unless a full mesh configuration is formed between all Border Gateway Protocol (BGP) devices within an autonomous system (AS). Configuring a route reflector allows a device to advertise or reflect the iBGP learned routes to other iBGP speakers.

Ethernet VPN (EVPN) Autodiscovery supports BGP route reflectors. A BGP route reflector can be used to reflect BGP EVPN prefixes without EVPN being explicitly configured on the route reflector. The route reflector does not participate in autodiscovery; that is, no pseudowires are set up between the route reflector and the provider edge (PE) devices. A route reflector reflects EVPN prefixes to other PE devices so that these PE devices do not need to have a full mesh of BGP sessions. The network administrator configures only the BGP EVPN address family on a route reflector.

BGP uses the Layer 2 VPN (L2VPN) routing information base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. The prefix and path information is stored in the L2VPN database, which allows BGP to make decisions about the best path. When BGP distributes the endpoint provisioning information in an update message to all its BGP neighbors, this endpoint information is used to configure a pseudowire mesh to support L2VPN-based services.

EVPN Address Family

BGP supports Layer 2 VPN (L2VPN) EVPN address family under router configuration mode to carry L2VPN EVPN autodiscovery and signaling Network Layer Reachability Information (NLRI) to Border Gateway Protocol (BGP) neighbors. This address family is allowed on both internal BGP (iBGP) and external BGP (eBGP) neighbors under default virtual routing and forwarding (VRF) for both IPv4 and IPv6 neighbors. The EVPN SAFI is not supported under VRF and VRF neighbors.

How to Configure BGP PBB EVPN Route Reflector Support

Configuring BGP PBB EVPN Route Reflector

Perform this task on the Border Gateway Protocol (BGP) route reflector to configure the device as a BGP route reflector and configure the specified neighbor as its client and to display the information from the BGP routing table.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *as-number*
4. **address-family** *l2vpn [vpls | evpn]*
5. **neighbor** *{ip-address | peer-group-name}* **activate**
6. **neighbor** *{ip-address | ipv6-address | peer-group-name}* **route-reflector-client**
7. **end**
8. **show bgp l2vpn evpn all**
9. **debug bgp l2vpn evpn updates**
10. **clear bgp l2vpn evpn**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>as-number</i> Example: Device(config)# router bgp 1	Configures a BGP routing process and enters router configuration mode.
Step 4	address-family <i>l2vpn [vpls evpn]</i> Example: Device(config-router)# address-family l2vpn evpn	Specifies the L2VPN address family and enters address family configuration mode. The optional evpn keyword specifies that EVPN endpoint provisioning information is to be distributed to BGP peers.
Step 5	neighbor <i>{ip-address peer-group-name}</i> activate Example: Device(config-router-af)# neighbor 10.0.0.2 activate	Enables PBB EVPN with the specified BGP neighbor.
Step 6	neighbor <i>{ip-address ipv6-address peer-group-name}</i> route-reflector-client Example: Device(config-router-af)# neighbor 10.0.0.2 route-reflector-client	Configures the local device as the BGP route reflector and the specified neighbor as one of its clients.

	Command or Action	Purpose
Step 7	end Example: Device(config-router-af)# end	Exits address family configuration mode and returns to privileged EXEC mode.
Step 8	show bgp l2vpn evpn all Example: Device# show bgp l2vpn evpn all	(Optional) Displays the complete L2VPN EVPN database.
Step 9	debug bgp l2vpn evpn updates Example: Device# debug bgp l2vpn evpn updates events	(Optional) Specifies debugging messages for BGP update.
Step 10	clear bgp l2vpn evpn Example: Device# clear bgp l2vpn evpn *	(Optional) Specifies that all current BGP sessions will be reset.

Configuration Examples for BGP PBB EVPN Route Reflector Support

Example: Configuring BGP PBB EVPN Route Reflector

In the following example, the local device is a route reflector. It passes learned iBGP routes to the neighbor at 10.0.0.2:

```
Device# configure terminal
Device(config)# router bgp 1
Device(config-router)# address-family l2vpn evpn
Device(config-router-af)# neighbor 10.0.0.2 activate
Device(config-router-af)# neighbor 10.0.0.2 route-reflector-client
Device(config-router-af)# exit address-family
```

In the following example, the **show bgp l2vpn evpn all route-type 1** command displays the Ethernet autodiscovery route information:

```
show bgp l2vpn evpn all route-type 1
```

```
BGP routing table entry for
[1][100.100.100.100:11111][AAAABBBBCCCCDDDEEEE][23456789][101234]/25, version 2
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
    1          2          3
Refresh Epoch 1
Local, (Received from a RR-client)
  19.0.101.1 from 19.0.101.1 (19.0.101.1)
    Origin IGP, localpref 100, valid, internal, best
    Extended Community: RT:100:101 EVPN LABEL:0x1:Label-101234
    rx pathid: 0, tx pathid: 0x0
```

Additional References for BGP PBB EVPN Route Reflector Support

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
BGP commands	Cisco IOS IP Routing: BGP Command Reference

Standards and RFCs

Standard/RFC	Title
RFC 4456	<i>BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)</i>
RFC 4684	<i>Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)</i>

Technical Assistance

Description	Link
The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies. To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds. Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.	http://www.cisco.com/support

Feature Information for BGP PBB EVPN Route Reflector Support

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for BGP PBB EVPN Route Reflector Support

Feature Name	Releases	Feature Information
BGP PBB EVPN Route Reflector Support		The BGP PBB EVPN Route Reflector Support feature provides Border Gateway Protocol (BGP) route reflector functionality for Ethernet VPN (EVPN) and provider backbone bridging (PBB) EVPN of Layer 2 VPN address family. EVPN enables customer MAC addresses as routable addresses and distributes them in BGP to avoid any data plane MAC address learning over the Multiprotocol Label Switching (MPLS) core network. The route reflector is enhanced to store the received EVPN updates without configuring EVPN explicitly on the route reflector and then advertises these updates to other provider edge (PE) devices so that the PEs do not need to have a full mesh of BGP sessions. The following command was modified: address-family l2vpn.