



BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN

The BGP Multipath Load Sharing for eBGP and iBGP feature allows you to configure multipath load balancing with both external BGP (eBGP) and internal BGP (iBGP) paths in Border Gateway Protocol (BGP) networks that are configured to use Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). This feature provides improved load balancing deployment and service offering capabilities and is useful for multi-homed autonomous systems and Provider Edge (PE) routers that import both eBGP and iBGP paths from multihomed and stub networks.

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the [Feature Information for BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN, page 11](#).

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



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2004–2008 Cisco Systems, Inc. All rights reserved.

Contents

- [Prerequisites for BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN, page 2](#)
- [Restrictions for BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN, page 2](#)
- [Information About BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN, page 3](#)
- [How to Configure BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN, page 5](#)
- [Configuration Examples for the BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN Feature, page 7](#)
- [Additional References, page 9](#)
- [Feature Information for BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN, page 11](#)
- [Command Reference, page 10](#)

Prerequisites for BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN

Load Balancing is Configured Under CEF

Cisco Express Forwarding (CEF) or distributed CEF (dCEF) must be enabled on all participating routers.

Restrictions for BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN

Address Family Support

This feature is configured on a per VPN routing and forwarding instance (VRF) basis. This feature can be configured under only the IPv4 VRF address family.

Memory Consumption Restriction

Each BGP multipath routing table entry will use additional memory. We recommend that you do not use this feature on a router with a low amount of available memory and especially if router is carries full Internet routing tables.

Route Reflector Limitation

When multiple iBGP paths installed in a routing table, a route reflector will advertise only one paths (next hop). If a router is behind a route reflector, all routers that are connected to multihomed sites will not be advertised unless a different route distinguisher is configured for each VRF.

Information About BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN

To configure the BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN feature, you must understand the following concepts:

- [Multipath Load Sharing Between eBGP and iBGP](#), page 3
- [eBGP and iBGP Multipath Load Sharing in a BGP MPLS Network](#), page 3
- [eBGP and iBGP Multipath Load Sharing With Route Reflectors](#), page 4
- [Benefits of Multipath Load Sharing for Both eBGP and iBGP](#), page 5

Multipath Load Sharing Between eBGP and iBGP

A BGP routing process will install a single path as the best path in the routing information base (RIB) by default. The **maximum-paths** command allows you to configure BGP to install multiple paths in the RIB for multipath load sharing. BGP uses the best path algorithm to still select a single multipath as the best path and advertise the best path to BGP peers.

**Note**

The number of paths of multipaths that can be configured is documented on the **maximum-paths** command reference page.

Load balancing over the multipaths is performed by CEF. CEF load balancing is configured on a per-packet round robin or on a per session (source and destination pair) basis. For information about CEF, refer to the “[Cisco Express Forwarding Overview](#)” documentation:

The BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature is enabled only under the IPv4 VRF address family configuration mode. When enabled, this feature can perform load balancing on eBGP and/or iBGP paths that are imported into the VRF. The number of multipaths is configured on a per VRF basis. Separate VRF multipath configurations are isolated by unique route distinguisher.


**Note**

The BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature operates within the parameters of configured outbound routing policy.

eBGP and iBGP Multipath Load Sharing in a BGP MPLS Network

[Figure 1](#) shows a service provider BGP MPLS network that connects two remote networks to PE router 1 and PE router 2. PE router 1 and PE router 2 are both configured for VPNv4 unicast iBGP peering. Network 2 is a multihomed network that is connected to PE router 1 and PE router 2. Network 2 also has extranet VPN services configured with Network 1. Both Network 1 and Network 2 are configured for eBGP peering with the PE routers.

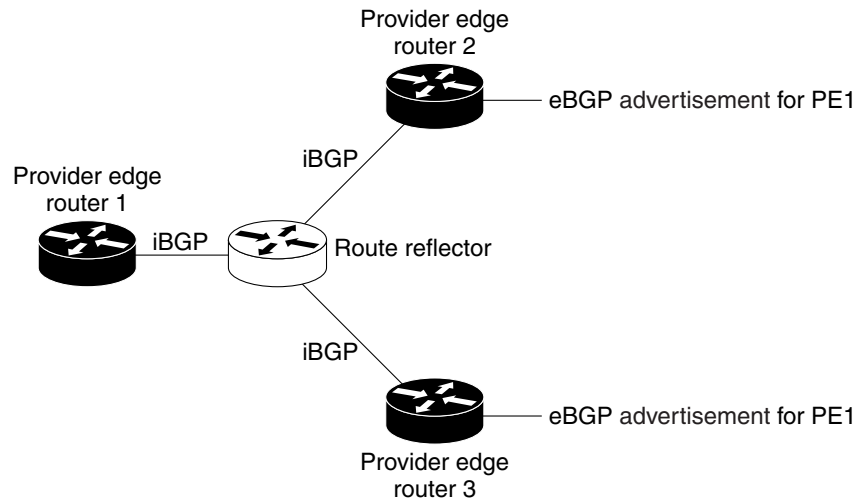
Figure 1 A Service Provider BGP MPLS Network



PE router 1 can be configured with the BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature so that both iBGP and eBGP paths can be selected as multipaths and imported into the VRF of Network 1. The multipaths will be used by CEF to perform load balancing. IP traffic that is sent from Network 2 to PE router 1 and PE router 2 will be sent across the eBGP paths as IP traffic. IP traffic that is sent across the iBGP path will be sent as MPLS traffic, and MPLS traffic that is sent across an eBGP path will be sent as IP traffic. Any prefix that is advertised from Network 2 will be received by PE router 1 through route distinguisher (RD) 21 and RD 22. The advertisement through RD 21 will be carried in IP packets, and the advertisement through RD 22 will be carried in MPLS packets. Both paths can be selected as multipaths for VRF1 and installed into the VRF1 RIB.

eBGP and iBGP Multipath Load Sharing With Route Reflectors

Figure 2 shows a topology that contains three PE routers and a route reflector, all configured for iBGP peering. PE router 2 and PE router 3 each advertise an equal preference eBGP path to PE router 1. By default, the route reflector will choose only one path and advertise PE router 1.

Figure 2 A Topology With a Route Reflector

For all equal preference paths to PE router 1 to be advertised through the route reflector, you must configure each VRF with a different RD. The prefixes received by the route reflector will be recognized differently and advertised to PE router 1.

Benefits of Multipath Load Sharing for Both eBGP and iBGP

The BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS VPN feature allows multihomed autonomous systems and PE routers to be configured to distribute traffic across both eBGP and iBGP paths.

How to Configure BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN

This section contains the following procedures:

- [Configuring Multipath Load Sharing for Both eBGP and iBGP](#), page 5
- [Verifying Multipath Load Sharing for Both eBGP and iBGP](#), page 6

Configuring Multipath Load Sharing for Both eBGP and iBGP

To configure this feature, perform the steps in this section.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **address-family ipv4** [*mdt* | *multicast* | *tunnel* | *unicast* [*vrf vrf-name*] | *vrf vrf-name*] | **ipv6** [*multicast* | *unicast*] | **vpn4** [*unicast*]

5. `maximum-paths eibgp number [import number]`
6. `end`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>enable</code> Example: Router> enable	Enables higher privilege levels, such as privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	<code>configure terminal</code> Example: Router# configure terminal	Enters global configuration mode.
Step 3	<code>router bgp autonomous-system-number</code> Example: Router(config)# router bgp 40000	Enters router configuration mode to create or configure a BGP routing process.
Step 4	<code>address-family ipv4 vrf vrf-name</code> Example: Router(config-router)# address-family ipv4 vrf RED	Places the router in address family configuration mode. <ul style="list-style-type: none"> • Separate VRF multipath configurations are isolated by unique route distinguisher.
Step 5	<code>maximum-paths eibgp number [import number]</code> Example: Router(config-router-af)# maximum-paths eibgp 6	Configures the number of parallel iBGP and eBGP routes that can be installed into a routing table. <p>Note The <code>maximum-paths eibgp</code> command can be configured only under the IPv4 VRF address family configuration mode and cannot be configured in any other address family configuration mode.</p>
Step 6	<code>end</code> Example: Router(config-router-af)# end	Exits address family configuration mode, and enters Privileged EXEC mode.

Verifying Multipath Load Sharing for Both eBGP and iBGP

To verify this feature, perform the steps in this section

SUMMARY STEPS

1. `enable`
2. `show ip bgp neighbors [neighbor-address [advertised-routes | dampened-routes | flap-statistics | paths [regexp] | received prefix-filter | received-routes | routes]]`
3. `show ip bgp vpnv4 {all | rd route-distinguisher | vrf vrf-name}`
4. `show ip route vrf vrf-name`

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables higher privilege levels, such as privileged EXEC mode. • Enter your password if prompted.
Step 2	show ip bgp neighbors [<i>neighbor-address</i> <i>advertised-routes</i> <i>dampened-routes</i> <i>flap-statistics</i> <i>paths</i> [<i>regex</i>] <i>received-prefix-filter</i> <i>received-routes</i> <i>routes</i>] Example: Router# show ip bgp neighbors	Displays information about the TCP and BGP connections to neighbors.
Step 3	show ip bgp vpnv4 { <i>all</i> <i>rd route-distinguisher</i> <i>vrf vrf-name</i> } Example: Router# show ip bgp vpnv4 vrf RED	Displays VPN address information from the BGP table. This command is used to verify that the VRF has been received by BGP.
Step 4	show ip route vrf <i>vrf-name</i> Example: Router# show ip route vrf RED	Displays the IP routing table associated with a VRF instance. The show ip route vrf command is used to verify that the VRF is in the routing table.

Configuration Examples for the BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN Feature

The following examples show how to configure and verify this feature:

- [eBGP and iBGP Multipath Load Sharing Configuration Example, page 7](#)
- [eBGP and iBGP Multipath Load Sharing Verification Examples, page 7](#)

eBGP and iBGP Multipath Load Sharing Configuration Example

This following configuration example configures a router in address-family mode to select six BGP routes (eBGP or iBGP) as multipaths:

```
Router(config)# router bgp 40000
Router(config-router)# address-family ipv4 vrf RED
Router(config-router-af)# maximum-paths eibgp 6
Router(config-router-af)# end
```

eBGP and iBGP Multipath Load Sharing Verification Examples

To verify that iBGP and eBGP routes have been configured for load sharing, use the **show ip bgp vpnv4 EXEC** command or the **show ip route vrf EXEC** command.

In the following example, the **show ip bgp vpnv4** command is entered to display multipaths installed in the VPNv4 RIB:

```
Router# show ip bgp vpnv4 all 10.22.22.0

BGP routing table entry for 10:1:22.22.22.0/24, version 19
Paths:(5 available, best #5)
Multipath:eiBGP
  Advertised to non peer-group peers:
  10.0.0.2 10.0.0.3 10.0.0.4 10.0.0.5
  22
    10.0.0.2 (metric 20) from 10.0.0.4 (10.0.0.4)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:0x0:0:0 RT:100:1 0x0:0:0
      Originator:10.0.0.2, Cluster list:10.0.0.4
  22
    10.0.0.2 (metric 20) from 10.0.0.5 (10.0.0.5)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:0x0:0:0 RT:100:1 0x0:0:0
      Originator:10.0.0.2, Cluster list:10.0.0.5
  22
    10.0.0.2 (metric 20) from 10.0.0.2 (10.0.0.2)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1 0x0:0:0
  22
    10.0.0.2 (metric 20) from 10.0.0.3 (10.0.0.3)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:0x0:0:0 RT:100:1 0x0:0:0
      Originator:10.0.0.2, Cluster list:10.0.0.3
  22
    10.1.1.12 from 10.1.1.12 (10.22.22.12)
      Origin IGP, metric 0, localpref 100, valid, external, multipath, best
      Extended Community:RT:100:1
```

In the following example, the **show ip route vrf** command is entered to display multipath routes in the VRF table:

```
Router# show ip route vrf PATH 10.22.22.0

Routing entry for 10.22.22.0/24
  Known via "bgp 1", distance 20, metric 0
  Tag 22, type external
  Last update from 10.1.1.12 01:59:31 ago
  Routing Descriptor Blocks:
  * 10.0.0.2 (Default-IP-Routing-Table), from 10.0.0.4, 01:59:31 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
  10.0.0.2 (Default-IP-Routing-Table), from 10.0.0.5, 01:59:31 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
  10.0.0.2 (Default-IP-Routing-Table), from 10.0.0.2, 01:59:31 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
  10.0.0.2 (Default-IP-Routing-Table), from 10.0.0.3, 01:59:31 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
  10.1.1.12, from 10.1.1.12, 01:59:31 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
```

Where to Go Next

For information about advertising the bandwidth of an autonomous system exit link as an extended community, refer to the “[BGP Link Bandwidth](#)” document.

Additional References

For additional information related to BGP Multipath Load sharing for Both eBGP and iBGP in an MPLS VPN, refer to the following references:

Related Documents

Related Topic	Document Title
BGP commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	Cisco IOS IP Routing Protocols Command Reference
BGP configuration tasks	BGP Feature Roadmap module
Comprehensive BGP link bandwidth configuration examples and tasks	BGP Link Bandwidth module
CEF configuration tasks	Cisco Express Forwarding Overview module

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To obtain lists of supported MIBs by platform and Cisco IOS release, and to download MIB modules, go to the Cisco MIB website on Cisco.com at the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

RFCs

RFCs	Title
RFC 1771	A Border Gateway Protocol 4 (BGP4)

RFCs	Title
RFC 2547	BGP/MPLS VPNs
RFC 2858	Multiprotocol Extensions for BGP-4

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>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.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/techsupport

Command Reference

The following commands are introduced or modified in the feature or features documented in this module. For information about these commands, see the *Cisco IOS IP Routing Protocols Command Reference* at http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html. For information about all Cisco IOS commands, go to the Command Lookup Tool at <http://tools.cisco.com/Support/CLILookup> or to the *Cisco IOS Master Commands List*.

maximum-paths eibgp

Feature Information for BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN

Table 1 lists the features in this module and provides links to specific configuration information. Only features that were introduced or modified in Cisco IOS Releases 12.2(1), 12.0(3)S, 12.2(27)SBC, 12.2(33)SRB, 12.2(33)SXH, or later release appear in the table.

For information on a feature in this technology that is not documented here, see the “[Cisco BGP Features Roadmap](#).”

Not all commands may be available in your Cisco IOS software release. For release information about a specific command, see the command reference documentation.

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which Cisco IOS and Catalyst OS software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.



Note

Table 1 lists only the Cisco IOS software release that introduced support for a given feature in a given Cisco IOS software release train. Unless noted otherwise, subsequent releases of that Cisco IOS software release train also support that feature.

Table 1 Feature Information for BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN

Feature Name	Releases	Feature Configuration Information
BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN	12.0(24)S 12.2(14)S 12.2(18)SXE 12.2(4)T	The BGP Multipath Load Sharing for eBGP and iBGP feature allows you to configure multipath load balancing with both eBGP and iBGP paths in BGP networks that are configured to use MPLS VPNs. This feature provides improved load balancing deployment and service offering capabilities and is useful for multi-homed autonomous systems and PE routers that import both eBGP and iBGP paths from multihomed and stub networks. The following command was introduced or modified by this feature: maximum-paths eibgp .

CCDE, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0903R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2004–2009 Cisco Systems, Inc. All rights reserved.

