



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.

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

Release	Modification
12.2(4)T	This feature was introduced.
12.2(14)S	This feature was integrated into Cisco IOS Release 12.2(14)S.
12.0(24)S	This feature was integrated into Cisco IOS Release 12.0(24)S.
12.2(18)SXE	This feature was integrated into Cisco IOS Release 12.2(18)SXE.

Finding Support Information for Platforms and Cisco IOS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support. Access Cisco Feature Navigator at <http://www.cisco.com/go/fn>. You must have an account on Cisco.com. If you do not have an account or have forgotten your username or password, click **Cancel** at the login dialog box and follow the instructions that appear.



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

Copyright © 2005 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 8](#)
- [Additional References, page 10](#)
- [Command Reference, page 11](#)

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 4](#)
- [eBGP and iBGP Multipath Load Sharing With Route Reflectors, page 5](#)
- [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 *Cisco IOS Switching Services Configuration Guide* documentation:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/swit_vcg.htm

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 6](#)
- [Verifying Multipath Load Sharing for Both eBGP and iBGP, page 7](#)

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 | memory | network}`
3. `router bgp as-number`
4. `address-family ipv4 [mdt | multicast | tunnel | unicast [vrf vrf-name] | vrf vrf-name] | ipv6 [multicast | unicast] | vpnv4 [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. • Enter your password if prompted.
Step 2	<code>configure {terminal memory network}</code> Example: Router# configure terminal	Enters global configuration mode.
Step 3	<code>router bgp as-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. • 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. Note The maximum-paths eibgp command can be configured only under the IPv4 VRF address family configuration mode and cannot be configured in any other address family configuration mode.
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 [*regex*] | received-prefix-filter | received-routes | routes]]
3. **show ip bgp vpnv4** {all | rd *route-distinguisher* | vrf *vrf-name*} *ip-prefix/length* [longer-prefixes] [*output-modifiers*] | *network-address* [*mask*] [longer-prefixes] [*output-modifiers*] [cidr-only] | [community [*number* | exact-match | local-as | no-advertise | no-export]] | [community-list *name* | *number* [exact-match]] | [dampening dampened-paths | flap-statistics | parameters] | [filter-list *regex-acl*] | [inconsistent-as] | [injected-paths] | [labels] | [neighbors] | [paths [*regex*]] | [peer-group [*name* [summary]]] | [quote-regex [*regex*]] | [regex *string*] | [replication [*update-group* [*ip-address*]] | [*ip-address*]] | [rib-failure] | [route-map *name*] | [summary] | [templates peer-policy [*name*] | peer-session [*name*]] | update-group [*update-group* [*ip-address*]] | [*ip-address*]]
4. **show ip route vrf** *vrf-name* [connected] [protocol [*process-number*] [tag] [*output-modifiers*]] [*ip-prefix*] [list *number* [*output-modifiers*]] [profile] [static [*output-modifiers*]] [summary [*output-modifiers*]] [supernets-only [*output-modifiers*]] [traffic-engineering [*output-modifiers*]]

DETAILED STEPS

	Command or Action	Purpose
Step 1	<pre>enable</pre> <p>Example: Router> enable</p>	Enables higher privilege levels, such as privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	<pre>show ip bgp neighbors [neighbor-address advertised-routes dampened-routes flap-statistics paths [regex] received prefix-filter received-routes routes]</pre> <p>Example: Router# show ip bgp neighbors</p>	Displays information about the TCP and BGP connections to neighbors.

	Command or Action	Purpose
Step 3	<pre>show ip bgp vpnv4 {all rd route-distinguisher vrf vrf-name} ip-prefix/length [longer-prefixes] [output-modifiers] network-address [mask] [longer-prefixes] [output-modifiers] [cidr-only] [community number exact-match local-as no-advertise no-export]] [community-list name number [exact-match]] [dampening dampened-paths flap-statistics parameters] [filter-list regex-acl] [inconsistent-as] [injected-paths] [labels] [neighbors] [paths [regex]] [peer-group [name [summary]]] [quote-regex [regex]] [regex string] [replication [update-group [ip-address]] [ip-address]] [rib-failure] [route-map name] [summary] [templates peer-policy [name] peer-session[name]] update-group [update-group [ip-address]] [ip-address]]</pre> <p>Example: Router# show ip bgp vpnv4 vrf RED</p>	Displays VPN address information from the BGP table. This command is used to verify that the VRF has been received by BGP.
Step 4	<pre>show ip route vrf vrf-name [connected] [protocol [process-number] [tag] [output-modifiers]] [ip-prefix] [list number [output-modifiers]] [profile] [static [output-modifiers]] [summary [output-modifiers]] [supernets-only [output-modifiers]] [traffic-engineering [output-modifiers]]</pre> <p>Example: Router# show ip route vrf RED</p>	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 8](#)
- [eBGP and iBGP Multipath Load Sharing Verification Examples, page 9](#)

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:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s24/s_bgplb.htm

Additional References

For additional information related to BGP Dynamic Update Peer-Groups, 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	<ul style="list-style-type: none"> <i>Cisco IOS IP Command Reference, Volume 2 of 4: Routing Protocols, Release 12.3T</i>
BGP configuration tasks	<ul style="list-style-type: none"> <i>Cisco IOS IP Configuration Guide, Release 12.3</i>
Comprehensive BGP link bandwidth configuration examples and tasks	<ul style="list-style-type: none"> <i>BGP Link Bandwidth, Release 12.0(24)S</i>
CEF configuration tasks	<ul style="list-style-type: none"> <i>Cisco IOS Switching Services Configuration Guide</i>

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	<i>A Border Gateway Protocol 4 (BGP4)</i>
RFC 2547	<i>BGP/MPLS VPNs</i>
RFC 2858	<i>Multiprotocol Extensions for BGP-4</i>

Technical Assistance

Description	Link
Technical Assistance Center (TAC) home page, containing 30,000 pages of searchable technical content, including links to products, technologies, solutions, technical tips, tools, and lots more. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/public/support/tac/home.shtml

Command Reference

This section documents the [maximum-paths eibgp](#) command.

maximum-paths eibgp

To configure multipath load sharing for external BGP (eBGP) and internal (iBGP) routes, use the **maximum-paths eibgp** command in address family configuration mode. To disable multipath load sharing for eBGP and iBGP routes, use the **no** form of this command.

maximum-paths eibgp *number* [**import** *number*]

no maximum-paths eibgp *number* [**import** *number*]

Syntax Description

<i>number</i>	Specifies the number of routes to install to the routing table. See the usage guidelines for the number of paths that can be configured with this argument.
import <i>number</i>	(Optional) Specifies the number of redundant paths that can be configured as back up multipaths for a VRF. This keyword can only be configured under a VRF in address family configuration mode.
Note	We recommend that this feature is enabled only where needed and that the number of import paths be kept to the minimum (Typically, not more than two paths). For more information, see the related note in the usage guidelines of this command reference page.

Defaults

Border Gateway Protocol (BGP) by default will install only one best path in the routing table.

Command Modes

IPv4 VRF Address family configuration

Command History

Release	Modification
12.2(4)T	This command was introduced.
12.0(24)S	This command was integrated into Cisco IOS Release 12.0(24)S.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.0(25)S	The import keyword was introduced.
12.2(13)T	The import keyword was integrated into Cisco IOS Release 12.2(13)T.
12.2(14)S	The import keyword was integrated into Cisco IOS Release 12.2(14)S.

Usage Guidelines

The **maximum-paths eibgp** command used to configure Border Gateway Protocol (BGP) multipath load sharing in an Multiprotocol Label Switching (MPLS) virtual private network (VPN) using eBGP and iBGP routes. This feature is configured under a VPN routing and forwarding instance (VRF) in address family configuration mode. The number of multipaths is configured separately for each VRF. The number of paths that can be configured is determined by the version of Cisco IOS software. The following list shows current limits:

- Cisco IOS Release 12.0S based software: 8 paths
- Cisco IOS Release 12.3T based software: 16 paths
- Cisco IOS Release 12.2S based software: 32 paths

The **maximum-paths eibgp** command cannot be configured with the **maximum-paths** or **maximum-paths ibgp** command because the **maximum-paths eibgp** command is a superset of these commands.

**Note**

The configuration of this command does not override the existing outbound routing policy.

Configuring VRF Import Paths

A VRF will import only one path (best path) per prefix from the source VRF table, unless the prefix is exported with a different route-target. If the best path goes down, the destination will not be reachable until the next import event occurs, and then a new best path will be imported into the VRF table. The import event runs every 15 seconds by default.

The **import** keyword allows you to configure the VRF table to accept multiple redundant paths in addition to the best path. An import path is a redundant path, and it can have a next hop that matches an installed multipath. This feature should be used when there are multiple paths with identical next hops available to ensure optimal convergence times. A typical application of this feature is to configure redundant paths in a network that has multiple route reflectors for redundancy.

**Note**

Configuring redundant paths with the **import** keyword can increase CPU and memory utilization significantly, especially in a network where there are many prefixes to learn and a large number of configured VRFs. It is recommended that this feature is only configured as necessary and that the minimum number of redundant paths are configured (Typically, not more than two).

Examples

In the following example, the router is configured to install 6 eBGP or iBGP routes into the VRF routing table:

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

In the following example, the router is configured to install 4 equal-cost routes and 2 import routes (backup) in the VRF routing table:

```
Router(config)# router bgp 45000
Router(config-router)# address-family ipv4 vrf GREEN
Router(config-router-af)# maximum-paths eibgp 4 import 2
```

In the following example, the router is configured to install 2 import routes in the VRF routing table:

```
Router(config)# router bgp 50000
Router(config-router)# address-family ipv4 vrf ORANGE
Router(config-router-af)# maximum-paths eibgp import 2
```

**Note**

Separate VRFs must be configured with different route distinguisher to support separate multipath configurations.

Related Commands	Command	Description
	maximum-paths	Configures the number of equal-cost routes that BGP will install in the routing table.
	maximum-paths ibgp	Configures the number of equal-cost or unequal-cost routes that BGP will install in the routing table.
	show ip bgp	Displays entries in the BGP routing table.

CCVP, the Cisco logo, and Welcome to the Human Network are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath 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. (0711R)

Copyright © 2005 Cisco Systems, Inc. All rights reserved