



BGP Link Bandwidth

The Border Gateway Protocol (BGP) Link Bandwidth feature is used to advertise the bandwidth of an autonomous system exit link as an extended community. This feature is configured for links between directly connected external BGP (eBGP) neighbors. The link bandwidth extended community attribute is propagated to iBGP peers when extended community exchange is enabled. This feature is used with BGP multipath features to configure load balancing over links with unequal bandwidth.

History for the BGP Link Bandwidth Feature

Release	Modification
12.2(2)T	This feature was introduced.
12.2(14)S	This feature was integrated into Cisco IOS Release 12.0(14)S.
12.2(11)T	This feature was integrated in Cisco IOS Release 12.2(11)T.
12.0(24)S	This feature was integrated into Cisco IOS Release 12.0(24)S.

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.

Contents

- [Prerequisites for BGP Link Bandwidth, page 2](#)
- [Restrictions for BGP Link Bandwidth, page 2](#)
- [Information About BGP Link Bandwidth, page 2](#)
- [How to Configure BGP Link Bandwidth, page 3](#)
- [Configuration Examples for BGP Link Bandwidth, page 5](#)
- [Additional References, page 9](#)
- [Command Reference, page 11](#)



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

Copyright © 2005 Cisco Systems, Inc. All rights reserved.

Prerequisites for BGP Link Bandwidth

- BGP load balancing or multipath load balancing must be configured before this feature is enabled.
- BGP extended community exchange must be enabled between iBGP neighbors to which the link bandwidth attribute is to be advertised.
- Cisco Express Forwarding (CEF) or distributed CEF (dCEF) must be enabled on all participating routers.

Restrictions for BGP Link Bandwidth

- This feature can be configured only under IPv4 and VPNv4 address family sessions.
- BGP can originate the link bandwidth community only for directly connected links to eBGP neighbors.
- Both iBGP and eBGP load balancing are supported in IPv4 and VPNv4 address families. However, eiBGP load balancing is supported only in VPNv4 address-family.

Information About BGP Link Bandwidth

To configure the BGP Link Bandwidth feature, you must understand the following concept:

- [BGP Link Bandwidth Overview, page 2](#)
- [Link Bandwidth Extended Community Attribute, page 3](#)
- [Benefits of the BGP Link Bandwidth Feature, page 3](#)

BGP Link Bandwidth Overview

The BGP Link Bandwidth feature used to enable multipath load balancing for external links with unequal bandwidth capacity. This feature is enabled under an IPv4 or VPNv4 address family sessions by entering the **bgp dmzlink-bw** command. This feature supports both iBGP, eBGP multipath load balancing, and eiBGP multipath load balancing in Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). When this feature is enabled, routes learned from directly connected external neighbor are propagated through the internal BGP (iBGP) network with the bandwidth of the source external link.

The link bandwidth extended community indicates the preference of an autonomous system exit link in terms of bandwidth. This extended community is applied to external links between directly connected eBGP peers by entering the **neighbor dmzlink-bw** command. The link bandwidth extended community attribute is propagated to iBGP peers when extended community exchange is enabled with the **neighbor send-community** command.

Link Bandwidth Extended Community Attribute

The link bandwidth extended community attribute is a 4-byte value that is configured for a link that on the demilitarized zone (DMZ) interface that connects two single hop eBGP peers. The link bandwidth extended community attribute is used as a traffic sharing value relative to other paths while forwarding traffic. Two paths are designated as equal for load balancing if the weight, local-pref, as-path length, Multi Exit Discriminator (MED), and Interior Gateway Protocol (IGP) costs are the same.

Benefits of the BGP Link Bandwidth Feature

The BGP Link Bandwidth feature allows BGP to be configured to send traffic over multiple iBGP or eBGP learned paths where the traffic that is sent is proportional to the bandwidth of the links that are used to exit the autonomous system. The configuration of this feature can be used with eBGP and iBGP multipath features to enable unequal cost load balancing over multiple links. Unequal cost load balancing over links with unequal bandwidth was not possible in BGP before the BGP Link Bandwidth feature was introduced.

How to Configure BGP Link Bandwidth

This section contains the following procedures:

- [Configuring BGP Link Bandwidth, page 3](#)
- [Verifying BGP Link Bandwidth Configuration, page 5](#)

Configuring BGP Link Bandwidth

To configure the BGP Link Bandwidth 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**] | **vpn4** [**unicast**]
5. **bgp** **dmzlink-bw**
6. **neighbor** *ip-address* **dmzlink-bw**
7. **neighbor** *ip-address* **send-community** [**both** | **extended** | **standard**]
8. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable 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	configure { terminal memory network } Example: Router# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>autonomous-system-number</i> Example: Router(config)# router bgp 50000	Enters router configuration mode to create or configure a BGP routing process.
Step 4	address-family ipv4 [mdt multicast tunnel unicast [vrf vrf-name] vrf vrf-name] ipv6 [multicast unicast] vpn4 [unicast] Example: Router(config-router)# address-family ipv4	Places the router in address family configuration mode. <ul style="list-style-type: none">The BGP Link Bandwidth feature is supported only under the IPv4 and VPNv4 address families.
Step 5	bgp dmzlink-bw Example: Router(config-router-af)# bgp dmzlink-bw	Configures BGP to distribute traffic proportionally to the bandwidth of the link. <ul style="list-style-type: none">This command must be entered on each router that contains an external interface that is to be used for multipath load balancing.
Step 6	neighbor ip-address dmzlink-bw Example: Router(config-router-af)# neighbor 172.16.1.1 dmzlink-bw	Configures BGP to include the link bandwidth attribute for routes learned from the external interface specified IP address. <ul style="list-style-type: none">This command must be configured for each eBGP link that is to be configured as a multipath. Enabling this command allows the bandwidth of the external link to be propagated through the link bandwidth extended community.
Step 7	neighbor ip-address send-community [both extended standard] Example: Router(config-router-af)# neighbor 10.10.10.1 send-community extended	(Optional) Enables community and/or extended community exchange with the specified neighbor. <ul style="list-style-type: none">This command must be configured for iBGP peers to which the link bandwidth extended community attribute is to be propagated.
Step 8	end Example: Router(config-router-af)# end	Exits address family configuration mode, and enters Privileged EXEC mode.

Verifying BGP Link Bandwidth Configuration

To verify the BGP Link Bandwidth feature, perform the steps in this section.

SUMMARY STEPS

1. `enable`
2. `show ip bgp ip-address [longer-prefixes [injected] | shorter-prefixes [mask-length]]`
3. `show ip route [[ip-address [mask] [longer-prefixes]] | [protocol [process-id]] | [list access-list-number | access-list-name] | [static download]]`

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>show ip bgp ip-address [longer-prefixes [injected] shorter-prefixes [mask-length]]</code> Example: Router# show ip bgp 10.0.0.0	Displays information about the TCP and BGP connections to neighbors. <ul style="list-style-type: none"> • The output displays the status of the link bandwidth configuration. The bandwidth of the link is shown in kilobytes.
Step 3	<code>show ip route [[ip-address [mask] [longer-prefixes]] [protocol [process-id]] [list access-list-number access-list-name] [static download]]</code> Example: Router# show ip route 10.0.0.0	Displays the current state of the routing table. <ul style="list-style-type: none"> • The output displays traffic share values, including the weights of the links that are used to direct traffic proportionally to the bandwidth of each link.

Configuration Examples for BGP Link Bandwidth

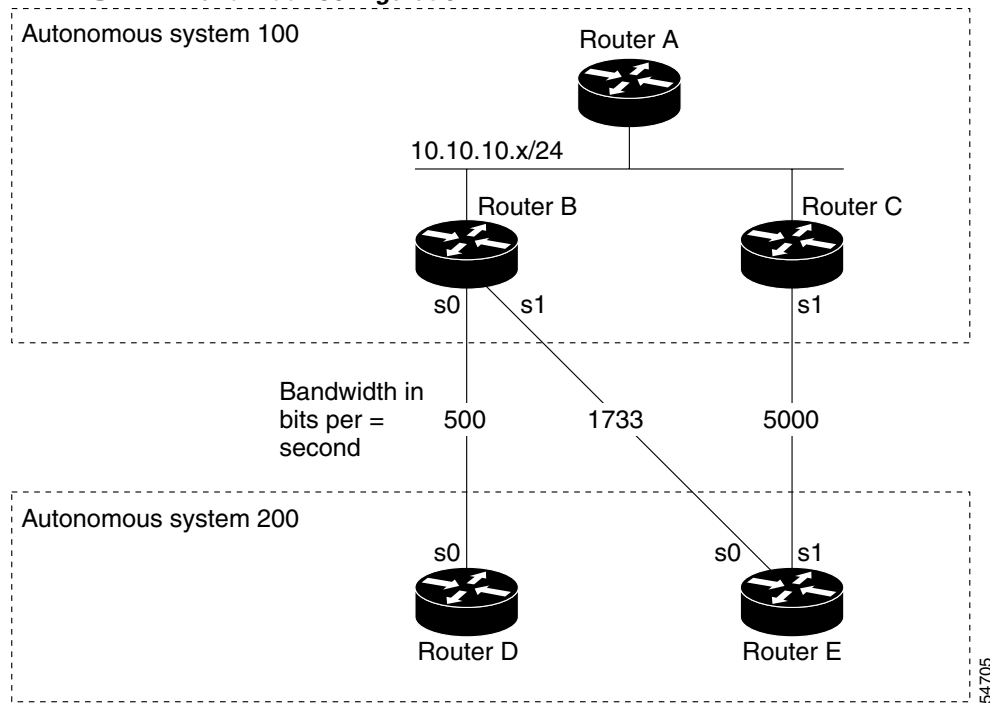
The following examples show how to configure and verify this feature:

- [BGP Link Bandwidth Configuration Example, page 6](#)
- [Verifying BGP Link Bandwidth, page 8](#)

BGP Link Bandwidth Configuration Example

In the following examples, the BGP Link Bandwidth feature is configured so BGP will distribute traffic proportionally to the bandwidth of each external link. Figure 1 shows two external autonomous systems connected by three links that each carry a different amount of bandwidth (unequal cost links). Multipath load balancing is enabled and traffic is balanced proportionally.

Figure 1 BGP Link Bandwidth Configuration



Router A Configuration

In the following example, Router A is configured to support iBGP multipath load balancing and to exchange the BGP extended community attribute with iBGP neighbors:

```
Router A(config)# router bgp 100
Router A(config-router)# neighbor 10.10.10.2 remote-as 100
Router A(config-router)# neighbor 10.10.10.3 remote-as 100
Router A(config-router)# address-family ipv4
Router A(config-router)# bgp dmzlink-bw
Router A(config-router-af)# neighbor 10.10.10.2 activate
Router A(config-router-af)# neighbor 10.10.10.2 send-community both
Router A(config-router-af)# neighbor 10.10.10.3 activate
Router A(config-router-af)# neighbor 10.10.10.3 send-community both
Router A(config-router-af)# maximum-paths ibgp 6
```

Router B Configuration

In the following example, Router B is configured to support multipath load balancing, to distribute Router D and Router E link traffic proportionally to the bandwidth of each link, and to advertise the bandwidth of these links to iBGP neighbors as an extended community:

```
Router B(config)# router bgp 100
Router B(config-router)# neighbor 10.10.10.1 remote-as 100
Router B(config-router)# neighbor 10.10.10.3 remote-as 100
Router B(config-router)# neighbor 172.16.1.1 remote-as 200
Router B(config-router)# neighbor 172.16.1.1 ebgp-multihop 1
Router B(config-router)# neighbor 172.16.2.2 remote-as 200
Router B(config-router)# neighbor 172.16.2.2 ebgp-multihop 1
Router B(config-router)# address-family ipv4
Router B(config-router-af)# bgp dmzlink-bw
Router B(config-router-af)# neighbor 10.10.10.1 activate
Router B(config-router-af)# neighbor 10.10.10.1 next-hop-self
Router B(config-router-af)# neighbor 10.10.10.1 send-community both
Router B(config-router-af)# neighbor 10.10.10.3 activate
Router B(config-router-af)# neighbor 10.10.10.3 next-hop-self
Router B(config-router-af)# neighbor 10.10.10.3 send-community both
Router B(config-router-af)# neighbor 172.16.1.1 activate
Router B(config-router-af)# neighbor 172.16.1.1 dmzlink-bw
Router B(config-router-af)# neighbor 172.16.2.2 activate
Router B(config-router-af)# neighbor 172.16.2.2 dmzlink-bw
Router B(config-router-af)# maximum-paths ibgp 6
Router B(config-router-af)# maximum-paths 6
```

Router C Configuration

In the following example, Router C is configured to support multipath load balancing and to advertise the bandwidth of the link with Router E to iBGP neighbors as an extended community:

```
Router C(config)# router bgp 100
Router C(config-router)# neighbor 10.10.10.1 remote-as 100
Router C(config-router)# neighbor 10.10.10.2 remote-as 100
Router C(config-router)# neighbor 172.16.3.30 remote-as 200
Router C(config-router)# neighbor 172.16.3.30 ebgp-multihop 1
Router C(config-router)# address-family ipv4
Router C(config-router-af)# bgp dmzlink-bw
Router C(config-router-af)# neighbor 10.10.10.1 activate
Router C(config-router-af)# neighbor 10.10.10.1 send-community both
Router C(config-router-af)# neighbor 10.10.10.1 next-hop-self
Router C(config-router-af)# neighbor 10.10.10.2 activate
Router C(config-router-af)# neighbor 10.10.10.2 send-community both
Router C(config-router-af)# neighbor 10.10.10.2 next-hop-self
Router C(config-router-af)# neighbor 172.16.3.3 activate
Router C(config-router-af)# neighbor 172.16.3.3 dmzlink-bw
Router C(config-router-af)# maximum-paths ibgp 6
Router C(config-router-af)# maximum-paths 6
```

Verifying BGP Link Bandwidth

The examples in this section show the verification of this feature on Router A and Router B.

Router B

In the following example, the **show ip bgp** command is entered on Router B to verify that two unequal cost best paths have been installed into the BGP routing table. The bandwidth for each link is displayed with each route.

```
Router B# show ip bgp 192.168.1.0
BGP routing table entry for 192.168.1.0/24, version 48
Paths: (2 available, best #2)
Multipath: eBGP
  Advertised to update-groups:
    1          2
200
  172.16.1.1 from 172.16.1.2 (192.168.1.1)
    Origin incomplete, metric 0, localpref 100, valid, external, multipath, best
    Extended Community: 0x0:0:0
    DMZ-Link Bw 278 kbytes
200
  172.16.2.2 from 172.16.2.2 (192.168.1.1)
    Origin incomplete, metric 0, localpref 100, valid, external, multipath, best
    Extended Community: 0x0:0:0
    DMZ-Link Bw 625 kbytes
```

Router A

In the following example, the **show ip bgp** command is entered on Router A to verify that the link bandwidth extended community has been propagated through the iBGP network to Router A. The output shows that a route for each exit link (on Router B and Router C) to autonomous system 200 has been installed as a best path in the BGP routing table.

```
Router A# show ip bgp 192.168.1.0
BGP routing table entry for 192.168.1.0/24, version 48
Paths: (3 available, best #3)
Multipath: eBGP
  Advertised to update-groups:
    1          2
200
  172.16.1.1 from 172.16.1.2 (192.168.1.1)
    Origin incomplete, metric 0, localpref 100, valid, external, multipath
    Extended Community: 0x0:0:0
    DMZ-Link Bw 278 kbytes
200
  172.16.2.2 from 172.16.2.2 (192.168.1.1)
    Origin incomplete, metric 0, localpref 100, valid, external, multipath, best
    Extended Community: 0x0:0:0
    DMZ-Link Bw 625 kbytes
200
  172.16.3.3 from 172.16.3.3 (192.168.1.1)
    Origin incomplete, metric 0, localpref 100, valid, external, multipath, best
    Extended Community: 0x0:0:0
    DMZ-Link Bw 2500 kbytes
```

Router A

In the following example, the **show ip route** command is entered on Router A to verify the multipath routes that are advertised and the associated traffic share values:

```
Router A# show ip route 192.168.1.0
Routing entry for 192.168.1.0/24
  Known via "bgp 100", distance 200, metric 0
  Tag 200, type internal
  Last update from 192.168.1.1 00:01:43 ago
  Routing Descriptor Blocks:
  * 192.168.1.1, from 192.168.1.1, 00:01:43 ago
    Route metric is 0, traffic share count is 13
    AS Hops 1, BGP network version 0
    Route tag 200
  192.168.2.2, from 192.168.2.2, 00:01:43 ago
    Route metric is 0, traffic share count is 30
    AS Hops 1, BGP network version 0
    Route tag 200
  192.168.3.3, from 179.168.3.3, 00:01:43 ago
    Route metric is 0, traffic share count is 120
    AS Hops 1, BGP network version 0
    Route tag 200
```

Where to Go Next

For information about the BGP Multipath Load Sharing for Both eBGP and iBGP in an MPLS-VPN feature, refer to the following document:

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

For more information about the iBGP Multipath Load Sharing feature, refer to the following document:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t2/ftbgpls.htm>

Additional References

The following sections provide references related to BGP Link Bandwidth feature.

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"> Cisco IOS IP Command Reference, Volume 2 of 4: Routing Protocols, Release 12.3T
BGP configuration tasks	<ul style="list-style-type: none"> Cisco IOS IP Configuration Guide, Release 12.3
CEF configuration tasks	<ul style="list-style-type: none"> Cisco IOS Switching Services Configuration Guide, 12.3

Standards

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

MIBs

MIB	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

RFC	Title
draft-ramachandra-bgp-ext-communities-09.txt	<i>BGP Extended Communities Attribute</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, and tools. Registered Cisco.com users can log in from this page to access even more content.	TAC Home Page: http://www.cisco.com/public/support/tac/home.shtml BGP Support Page: http://www.cisco.com/cgi-bin/Support/browse/psp_view.pl?p=Inter networking:BGP

Command Reference

This section documents the new commands.

- [bgp dmzlink-bw](#)
- [neighbor dmzlink-bw](#)

bgp dmzlink-bw

To configure BGP to distribute traffic proportionally over external links with unequal bandwidth when multipath load balancing is enabled, use the **bgp dmzlink-bw** command in address family configuration mode. To disable traffic distribution proportional to the link bandwidth, use the **no** form of this command.

bgp dmzlink-bw

no bgp dmzlink-bw

Syntax Description This command has no keywords or arguments.

Defaults No default behavior or values

Command Modes Address family configuration

Command History

Release	Modification
12.2(2)T	This command was introduced.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(11)T	This command was integrated into Cisco IOS Release 12.2(11)T.
12.0(24)S	This command was integrated into Cisco IOS Release 12.0(24)S.

Usage Guidelines

The **bgp dmzlink-bw** command is used to configure BGP to distribute traffic proportionally to the bandwidth of external links. This command is configured for multipath load balancing between directly connected external BGP (eBGP) neighbors. This feature is used with BGP multipath features to configure load balancing over links with unequal bandwidth. The **neighbor dmzlink-bw** command must also be configured for each external link through which multipath load balancing is configured to advertise the link bandwidth as an extended community. The **neighbor send-community** command is configured to exchange the link bandwidth extended community with internal BGP (iBGP) peers.

Examples

In the following example, the BGP Link Bandwidth feature is configured to allow multipath load balancing to distribute link traffic proportionally to the bandwidth of each external link, and to advertise the bandwidth of these links to iBGP peers as an extended community:

```
Router(config)# router bgp 100
Router(config-router)# neighbor 10.10.10.1 remote-as 100
Router(config-router)# neighbor 10.10.10.1 update-source Loopback 0
Router(config-router)# neighbor 10.10.10.3 remote-as 100
Router(config-router)# neighbor 10.10.10.3 update-source Loopback 0
Router(config-router)# neighbor 172.16.1.1 remote-as 200
Router(config-router)# neighbor 172.16.1.1 ebgp-multihop 1
Router(config-router)# neighbor 172.16.2.2 remote-as 200
Router(config-router)# neighbor 172.16.2.2 ebgp-multihop 1
Router(config-router)# address-family ipv4
```

```

Router(config-router-af)# bgp dmzlink-bw
Router(config-router-af)# neighbor 10.10.10.1 activate
Router(config-router-af)# neighbor 10.10.10.1 next-hop-self
Router(config-router-af)# neighbor 10.10.10.1 send-community both
Router(config-router-af)# neighbor 10.10.10.3 activate
Router(config-router-af)# neighbor 10.10.10.3 next-hop-self
Router(config-router-af)# neighbor 10.10.10.3 send-community both
Router(config-router-af)# neighbor 172.16.1.1 activate
Router(config-router-af)# neighbor 172.16.1.1 dmzlink-bw
Router(config-router-af)# neighbor 172.16.2.2 activate
Router(config-router-af)# neighbor 172.16.2.2 dmzlink-bw
Router(config-router-af)# maximum-paths ibgp 6
Router(config-router-af)# maximum-paths 6

```

Related Commands

Command	Description
neighbor dmzlink-bw	Configures BGP to advertise the bandwidth of links that are used to exit an autonomous system.
neighbor send-community	Specifies that a communities attribute should be sent to a BGP neighbor.
maximum-paths	Controls the maximum number of parallel routes an IP routing protocol can support.
maximum-paths eibgp	Controls the maximum number of eBGP or iBGP paths that can be configured in an MPLS VPN.
maximum-paths ibgp	Controls the maximum number of parallel iBGP routes that can be installed in a routing table.

neighbor dmzlink-bw

To configure BGP to advertise the bandwidth of links that are used to exit an autonomous system, use the **neighbor dmzlink-bw** command in address family configuration mode. To disable link bandwidth advertisement, use the **no** form of this command.

neighbor ip-address dmzlink-bw

no neighbor ip-address dmzlink-bw

Syntax Description

ip-address The IP address that identifies the external interface.

Defaults

No default behavior or values

Command Modes

Address family configuration

Command History

Release	Modification
12.2(2)T	This command was introduced.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(11)T	This command was integrated into Cisco IOS Release 12.2(11)T.
12.0(24)S	This command was integrated into Cisco IOS Release 12.0(24)S.

Usage Guidelines

The **neighbor dmzlink-bw** command is used to configure BGP to advertise the bandwidth of the specified external interface as an extended community. This command is configured for links between directly connected external BGP (eBGP) neighbors. The link bandwidth extended community attribute is propagated to iBGP peers when extended community exchange is enabled with the **neighbor send-community** command. This feature is used with BGP multipath features to configure load balancing over links with unequal bandwidth. This feature is not enabled until the **bgp dmzlink-bw** command is entered under the address family session for each router that has a directly connected external link.

Examples

In the following example, the BGP Link Bandwidth feature is configured to allow multipath load balancing to distribute link traffic proportionally to the bandwidth of each external link, and to advertise the bandwidth of these links to iBGP peers as an extended community:

```
Router(config)# router bgp 100
Router(config-router)# neighbor 10.10.10.1 remote-as 100
Router(config-router)# neighbor 10.10.10.1 update-source Loopback 0
Router(config-router)# neighbor 10.10.10.3 remote-as 100
Router(config-router)# neighbor 10.10.10.3 update-source Loopback 0
Router(config-router)# neighbor 172.16.1.1 remote-as 200
Router(config-router)# neighbor 172.16.1.1 ebgp-multihop 1
Router(config-router)# neighbor 172.16.2.2 remote-as 200
Router(config-router)# neighbor 172.16.2.2 ebgp-multihop 1
```

```

Router(config-router)# address-family ipv4
Router(config-router-af)# bgp dmzlink-bw
Router(config-router-af)# neighbor 10.10.10.1 activate
Router(config-router-af)# neighbor 10.10.10.1 next-hop-self
Router(config-router-af)# neighbor 10.10.10.1 send-community both
Router(config-router-af)# neighbor 10.10.10.3 activate
Router(config-router-af)# neighbor 10.10.10.3 next-hop-self
Router(config-router-af)# neighbor 10.10.10.3 send-community both
Router(config-router-af)# neighbor 172.16.1.1 activate
Router(config-router-af)# neighbor 172.16.1.1 dmzlink-bw
Router(config-router-af)# neighbor 172.16.2.2 activate
Router(config-router-af)# neighbor 172.16.2.2 dmzlink-bw
Router(config-router-af)# maximum-paths ibgp 6
Router(config-router-af)# maximum-paths 6

```

Related Commands

Command	Description
bgp dmzlink-bw	Configures BGP to distribute traffic proportionally over external links with unequal bandwidth when multipath load balancing is enabled.
neighbor send-community	Specifies that a communities attribute should be sent to a BGP neighbor.
maximum-paths	Controls the maximum number of parallel routes an IP routing protocol can support.
maximum-paths eibgp	Controls the maximum number of eBGP or iBGP paths that can be configured in an MPLS VPN.
maximum-paths ibgp	Controls the maximum number of parallel iBGP routes that can be installed in a 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.