



BGP Hide Local-Autonomous System

The BGP Hide Local-Autonomous System feature simplifies the task of changing the autonomous system number in a Border Gateway Protocol (BGP) network. Without this feature, this task can be difficult because, during the transition, internal BGP (iBGP) peers will reject external routes from peers with a local autonomous system number in the autonomous system number path to prevent routing loops. This feature allows you to transparently change the autonomous system number for the entire BGP network and ensure that routes can be propagated throughout the autonomous system, while the autonomous system number transition is incomplete.

Feature History for BGP Hide Local-Autonomous System

Release	Modification
12.0(18)S	This feature was introduced in Cisco IOS Release 12.0(18)S.
12.2(8)T	This feature was integrated into Cisco IOS Release 12.2(8)T.
12.2(14)S	This feature was integrated into Cisco IOS Release 12.2(14)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 Hide Local-Autonomous System, page 2](#)
- [Restrictions for BGP Hide Local-Autonomous System, page 2](#)
- [Information About BGP Hide Local-Autonomous System, page 2](#)
- [How to Configure BGP Hide Local-Autonomous System, page 3](#)
- [Additional References, page 6](#)
- [Command Reference, page 7](#)



Corporate Headquarters:

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

Copyright © 2003 Cisco Systems, Inc. All rights reserved.

Prerequisites for BGP Hide Local-Autonomous System

This document assumes that BGP is enabled and peering has been established in all participating networks.

Restrictions for BGP Hide Local-Autonomous System

- This feature can be configured for only external BGP (eBGP) peers.
- This feature should be deconfigured after the transition to the new autonomous system number is completed to minimize the possible creation of routing loops.

Information About BGP Hide Local-Autonomous System

To configure the BGP Hide Local-Autonomous System feature, you must understand the following concepts:

- [Changing the Autonomous System Number in a BGP Network, page 2](#)
- [Configuring the BGP Hide Local-Autonomous System Feature, page 2](#)
- [Benefits of the BGP Hide Local-Autonomous System Feature, page 3](#)

Changing the Autonomous System Number in a BGP Network

Changing the autonomous system number may be necessary when 2 separate BGP networks are combined under a single autonomous system. This typically occurs when one ISP purchases another ISP. The **neighbor local-as** command is used initially to configure BGP peers to support 2 local autonomous system numbers to maintain peering between 2 separate BGP networks. This configuration allows the ISP to immediately make the transition without any impact on existing customer configurations.

When the customer configurations have been updated, The next step is to complete the transition from the old autonomous system number to the new autonomous system number. However, when the **neighbor local-as** command is configured on a BGP peer, the local autonomous system number is automatically prepended to all routes that are learned from eBGP peers by default. This behavior, however, makes changing the autonomous system number for a service provider or large BGP network difficult because routes, with the prepended autonomous system number, will be rejected by internal BGP (iBGP) peers that are configured with the same autonomous system number. For example, if you configure an iBGP peer with the **neighbor 10.0.0.2 local-as 20** statement, all routes that are learned from the 10.0.0.2 external peer will automatically have the autonomous system number 20 prepended. Internal routers that are configured with the autonomous number 20 will detect these routes as routing loops and reject them. This behavior requires you to change the autonomous system number for all iBGP peers at the same time.

Configuring the BGP Hide Local-Autonomous System Feature

The BGP Hide Local-Autonomous System feature introduces the **no-prepend** keyword to the **neighbor local-as** command. The use of the **no-prepend** keyword will allow you to configure a BGP speaker to not prepend the local autonomous system number to any routes that are received from eBGP peers. This

feature can be used to help transparently change the autonomous system number of a BGP network and ensure that routes are propagated throughout the autonomous system, while the autonomous system number transition is incomplete. Because the local autonomous system number is not prepended to these routes, external routes will not be rejected by internal peers during the transition from one autonomous system number to another.

**Caution**

BGP prepends the autonomous system number from each BGP network that a route traverses. This behavior is designed to maintain network reachability information and to prevent routing loops from occurring. Configuring this feature incorrectly could create routing loops. So, the configuration of this feature should only be attempted by an experienced network operator.

Benefits of the BGP Hide Local-Autonomous System Feature

You can use the BGP Hide Local-Autonomous System feature to transparently change the autonomous system number of a BGP network and ensure that routes can be propagated throughout the autonomous system while the autonomous system number transition is incomplete.

How to Configure BGP Hide Local-Autonomous System

This section contains the following procedures:

- [Configuring BGP to Not Prepend the Local Autonomous System Number to Routes Learned From External Peers, page 3](#)
- [Verifying the Configuration of the BGP Hide Local-Autonomous Feature, page 5](#)

Configuring BGP to Not Prepend the Local Autonomous System Number to Routes Learned From External Peers

To configure a router that is running BGP with the BGP Hide Local-Autonomous System feature to not prepend the local autonomous system number to routes that are received from external peers, use the following steps.

Configuring the no-prepend Keyword

The **no-prepend** keyword should be used only to change the autonomous system number in a BGP network and should be deconfigured after the transition is complete because routing loops can be created if this feature is used incorrectly.

**Caution**

BGP prepends the autonomous system number from each BGP network that a route traverses. This behavior is designed to maintain network reachability information and to prevent routing loops from occurring. Configuring this feature incorrectly could create routing loops. So, the configuration of this feature should only be attempted by an experienced network operator.

Restrictions

- This feature can only be configured for eBGP peers.
- This feature should be deconfigured after the transition to the new autonomous system number is completed to minimize the possible creation of routing loops.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *as-number*
4. **address-family** {*ipv4* | *ipv6* | *vpnvp4*} [**multicast** | **unicast** | **vrf** {*vrf-name*}]
5. **network** *ip-address* [*network-mask*] [**route-map** *map-name*] [**backdoor**]
6. **neighbor** *ip-address* **remote-as** *as-number*
7. **neighbor** *ip-address* **local-as** *as-number* **no-prepend**
8. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>as-number</i> Example: Router(config)# router bgp 100	Enters router configuration mode, and creates a BGP routing process.
Step 4	address-family <i>ipv4</i> <i>ipv6</i> <i>vpnvp4</i> [multicast unicast vrf { <i>vrf-name</i> }] Example: Router(config-router-af)# address-family ipv4 unicast	Enter address family configuration mode to configure BGP peers to accept address family specific configurations. <ul style="list-style-type: none"> • The example command creates an IPv4 unicast address family session.
Step 5	network <i>ip-address</i> [<i>network-mask</i>] [route-map <i>map-name</i>] [backdoor] Example: Router(config-router-af)# network 10.1.1.1 remote-as 100	Specifies the networks to be advertised by the BGP and multiprotocol BGP routing processes.

	Command or Action	Purpose
Step 6	neighbor <i>ip-address</i> remote-as <i>as-number</i> Example: Router(config-router-af)# neighbor 10.1.1.1 remote-as 100	Establishes peering with the specified neighbor and configures the neighbor as internal to the local autonomous system.
Step 7	neighbor <i>ip-address</i> local-as <i>as-number</i> [no-prepend] Example: Router(config-router-af)# neighbor 10.1.1.1 local-as 300 no-prepend	Allows the customization of the autonomous system number for eBGP peer groupings. <ul style="list-style-type: none"> Using the no-prepend keyword configures the router to not prepend the local autonomous system number to routes that are received from external peers.
Step 8	end Example: Router(config-router)# end	Exits address-family configuration mode, and enters Privileged EXEC mode.

Examples

The following example configures the router to not prepend autonomous system number 300 to routes that are received from external peers:

```
router bgp 100
 network 10.1.1.0
 neighbor 10.1.1.1 remote-as 100
 neighbor 10.1.1.1 local-as 300 no-prepend
end
```

What to Do Next

You can verify that this feature is configured correctly with the **show ip bgp neighbors** command. Go to the [Verifying the Configuration of the BGP Hide Local-Autonomous Feature](#) section for instructions and example output.

Verifying the Configuration of the BGP Hide Local-Autonomous Feature

To verify that the local autonomous system number is not prepended to received external routes, use the **show ip bgp neighbors** command. The output of this command will display the local autonomous system number and then “no-prepend” for received external routes when this feature is configured.

The following example shows that autonomous system number 300 will not be prepended to the 10.1.1.1 peer:

```
Router# show ip bgp neighbors
BGP neighbor is 10.1.1.1, remote AS 100, local AS 300 no-prepend, external link
BGP version 4, remote router ID 10.1.1.1
BGP state = Established, up for 00:00:49
Last read 00:00:49, hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received(new)
  Address family IPv4 Unicast: advertised and received
  IPv4 MPLS Label capability:
Received 10 messages, 1 notifications, 0 in queue
```

Sent 10 messages, 0 notifications, 0 in queue
 Default minimum time between advertisement runs is 30 seconds

Additional References

The following sections provide references related to BGP Prefix-Based Outbound Route Filtering feature.

Related Documents

Related Topic	Document Title
The BGP Hide Local-Autonomous System feature is an extension of the BGP routing protocol. For more information about configuring BGP, autonomous systems, and route filtering, refer to the “Configuring BGP” chapter of the Release 12.2 <i>Cisco IOS IP Configuration Guide</i> and <i>Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols</i> .	<ul style="list-style-type: none"> • Cisco IOS IP Configuration Guide, Release 12.3 • Cisco IOS IP Command Reference, Volume 2 of 4: Routing Protocols, Release 12.3 • Configuring the BGP Local-AS Feature (

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

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 modified commands. All other commands used with this feature are documented in the Cisco IOS command reference publications.

- [neighbor local-as](#)
- [show ip bgp neighbors](#)

neighbor local-as

To allow customization of the autonomous system number for external BGP (eBGP) peer groupings, use the **neighbor local-as** command in address family or router configuration mode. To disable this function, use the **no** form of this command.

neighbor {*ip-address* | *peer-group-name*} **local-as** *as-number* [**no-prepend**]

no neighbor {*ip-address* | *peer-group-name*} **local-as** *as-number*

Syntax Description

<i>ip-address</i>	IP address of the local BGP-speaking neighbor.
<i>peer-group-name</i>	Name of a BGP peer group.
<i>as-number</i>	Valid autonomous system number from 1 to 65535. Do not specify the autonomous system number to which the neighbor belongs.
no-prepend	Configures the router to not prepend the local autonomous system number to any routes received from an external peer.

Defaults

The local autonomous system number is prepended to all external routes unless the **no-prepend** keyword is used.

Command Modes

Address family configuration
Router configuration

Command History

Release	Modification
12.0(5)S	This command was introduced.
12.0(5)T	Address family configuration mode was added.
12.0(18)S	The no-prepend keyword was added.
12.2(8)T	
12.2(14)S	

Usage Guidelines

Each Border Gateway Protocol (BGP) peer or peer group can be configured to have a local autonomous system value for the purpose of peering. In the case of peer groups, the local autonomous system value is valid for all peers within the peer group. This feature, however, cannot be customized for individual peers within the peer group.

If this command is configured, you cannot use the local BGP autonomous system number or the autonomous system number of the remote peer.

This command is valid only if the peer is a true eBGP peer. This feature does not work for two peers in different subautonomous systems of a confederation.

The **no-prepend** keyword should be used only to change the autonomous system number in a BGP network and should be deconfigured after the transition has been completed.

**Caution**

BGP prepends the autonomous system number from each BGP network that a route traverses. This behavior is designed to maintain network reachability information and to prevent routing loops from occurring. Configuring this command, using the **no-prepend** keyword, incorrectly could create routing loops. So, the configuration of this feature should only be attempted by an experienced network operator.

Examples

The following address family configuration example shows the customization of neighbor 172.20.1.1 configured to use autonomous system number 300 for the purpose of peering:

```
router bgp 109
address-family ipv4 multicast
 network 172.20.0.0
 neighbor 172.20.1.1 local-as 300
```

The following configuration example shows the customization of neighbor 172.20.1.1 configured to not prepend autonomous system number 300 to routes that are received from eBGP peers:

```
router bgp 109
address-family ipv4 unicast
 network 172.20.0.0
 neighbor 172.20.1.1 remote-as 200
 neighbor 172.20.1.1 local-as 300 no-prepend
```

Related Commands

Command	Description
address-family ipv4 (BGP)	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard IPv4 address prefixes.
address-family vpnv4	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard VPNv4 address prefixes.
show ip bgp neighbors	Displays information about BGP neighbors.
show ip bgp peer-group	Displays information about BGP peer groups.

show ip bgp neighbors

To display information about TCP and Border Gateway Protocol (BGP) connections to neighbors, use the **show ip bgp neighbors** EXEC command in privileged EXEC mode.

show ip bgp neighbors [*neighbor-address*] [**received-routes** | **routes** | **advertised-routes** | {**paths** *regex*} | **dampened-routes**]

Syntax Description

<i>neighbor-address</i>	(Optional) Address of the neighbor whose routes you have learned from. If you omit this argument, all neighbors are displayed.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes that are received and accepted. The information that is provided when this keyword is used is a subset of the output from the received-routes keyword.
advertised-routes	(Optional) Displays all the routes the router has advertised to the neighbor.
paths <i>regex</i>	(Optional) Regular expression that is used to match the paths received.
dampened-routes	(Optional) Displays the dampened routes to the neighbor at the IP address specified.

Command Modes

Privileged EXEC

Command History

Release	Modification
10.0	This command was introduced.
11.2	The received-routes keyword was added.
12.0(18)S	The no-prepend configuration option was added to the display output.
12.2(8)T	
12.2(14)S	

Examples

The following is sample output from the **show ip bgp neighbors** command:

```
Router# show ip bgp neighbors 172.16.232.178
```

```
BGP neighbor is 172.16.232.178, remote AS 35, local AS 2 no-prepend, external link
BGP version 4, remote router ID 192.168.3.3
BGP state = Established, up for 1w1d
Last read 00:00:53, hold time is 180, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received
  Address family IPv4 Unicast: advertised and received
  Address family IPv4 Multicast: advertised and received
Received 12519 messages, 0 notifications, 0 in queue
Sent 12523 messages, 0 notifications, 0 in queue
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 30 seconds
```

```

For address family: IPv4 Unicast
  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor
  Inbound path policy configured
  Outbound path policy configured

  Route map for incoming advertisements is uni-in
  Route map for outgoing advertisements is uni-out
  3 accepted prefixes consume 108 bytes
  Prefix advertised 6, suppressed 0, withdrawn 0

For address family: IPv4 Multicast
  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
  Inbound path policy configured
  Outbound path policy configured
  Route map for incoming advertisements is mul-in
  Route map for outgoing advertisements is mul-out
  3 accepted prefixes consume 108 bytes
  Prefix advertised 6, suppressed 0, withdrawn 0

Connections established 2; dropped 1
Last reset 1w1d, due to Peer closed the session
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 172.16.232.178, Local port: 179
Foreign host: 172.16.232.179, Foreign port: 11002

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x2CF49CF8):
Timer           Starts      Wakeups      Next
Retrans         12518         0           0x0
TimeWait        0             0           0x0
AckHold         12514        12281        0x0
SendWnd         0             0           0x0
KeepAlive       0             0           0x0
GiveUp          0             0           0x0
PmtuAger        0             0           0x0
DeadWait        0             0           0x0

iss: 273358651  snduna: 273596614  sndnxt: 273596614  sndwnd: 15434
irs: 190480283  rcvnxt: 190718186  rcvwnd: 15491  delrcvwnd: 893

SRTT: 300 ms, RTTO: 607 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: passive open, nagle, gen tcbs

Datagrams (max data segment is 1460 bytes):
Rcvd: 24889 (out of order: 0), with data: 12515, total data bytes: 237921
Sent: 24963 (retransmit: 0), with data: 12518, total data bytes: 237981

```

Table 1 describes the significant fields shown in the display.

Table 1 *show ip bgp neighbors Field Descriptions*

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number. If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external.
remote AS	Autonomous system of the neighbor.
local AS	Local autonomous system.
no-prepend	Indicates that the local autonomous system number will not be prepended to routes received from this peer.
external link	Indicates that this peer is an external BGP (eBGP) peer.
BGP version	BGP version being used to communicate with the remote router; the router ID (an IP address) of the neighbor is also specified.
remote router ID	IP address of the neighbor.
BGP state	Internal state of this BGP connection.
up for	Amount of time that the underlying TCP connection has been in existence.
Last read	Time that BGP last read a message from this neighbor.
hold time	Maximum amount of time, in seconds, that can elapse between messages from the peer.
keepalive interval	Time period between sending keepalive packets, which help ensure that the TCP connection is up.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor.
Route refresh	Indicates that the neighbor supports dynamic soft reset using the route refresh capability.
Address family IPv4 Unicast:	IP Version 4 unicast-specific properties of this neighbor.
Address family IPv4 Multicast:	IP Version 4 multicast-specific properties of this neighbor.
Received	Number of total BGP messages received from this peer, including keepalives.
notifications	Number of error messages received from the peer.
Sent	Total number of BGP messages that have been sent to this peer, including keepalives.
notifications	Number of error messages the router has sent to this peer.
Route refresh request:	Number of route refresh requests sent and received from this neighbor.
advertisement runs	Value of the minimum advertisement interval.
For address family:	Address family to which the following fields refer.
BGP table version	Indicates that the neighbor has been updated with this version of the primary BGP routing table.
neighbor version	Number used by the software to track the prefixes that have been sent and those that must be sent to this neighbor.

Table 1 *show ip bgp neighbors Field Descriptions (continued)*

Field	Description
Community attribute	Displays the community attributes that were sent if the neighbor send-community command is configured for this neighbor.
Inbound path policy	Indicates if an inbound policy is configured.
Outbound path policy	Indicates if an outbound policy is configured.
mul-in	Name of an inbound route map for the multicast address family.
mul-out	Name of an outbound route map for the multicast address family.
accepted prefixes	Number of prefixes accepted.
Prefix advertised	Number of prefixes advertised.
suppressed	Number of prefixes suppressed.
withdrawn	Number of prefixes withdrawn.
Connections established	Number of times the router has established a TCP connection and the two peers have agreed to speak BGP with each other.
dropped	Number of times that a good connection has failed or been taken down.
Last reset	Elapsed time, in seconds, since this peering session was last reset.
Connection state	State of the BGP peer.
unread input bytes	Number of bytes of packets still to be processed.
Local host, Local port	Peering address of the local router, plus the port.
Foreign host, Foreign port	Peering address of the neighbor.
Event Timers	Table displays the number of starts and wakeups for each timer.
iss	Initial send sequence number.
snduna	Last send sequence number the local host sent but for which it has not received an acknowledgment.
sndnxt	Sequence number the local host will send next.
sndwnd	TCP window size of the remote host.
irs	Initial receive sequence number.
rcvnxt	Last receive sequence number the local host has acknowledged.
rcvwnd	TCP window size of the local host.
delrcvwnd	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT	A calculated smoothed round-trip timeout.
RTTO	Round-trip timeout.
RTV	Variance of the round-trip time.
KRTT	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.
minRTT	Smallest recorded round-trip timeout (hard wire value used for calculation).

Table 1 *show ip bgp neighbors Field Descriptions (continued)*

Field	Description
maxRTT	Largest recorded round-trip timeout.
ACK hold	Time the local host will delay an acknowledgment in order to carry additional data with the acknowledgment.
Flags	IP precedence of the BGP packets.
Datagrams: Rcvd	Number of update packets received from a neighbor.
with data	Number of update packets received with data.
total data bytes	Total bytes of data.
Sent	Number of update packets sent.
with data	Number of update packets with data sent.
total data bytes	Total number of data bytes.

The following is sample output from the **show ip bgp neighbors** command with the **advertised-routes** keyword:

```
Router# show ip bgp neighbors 172.16.232.178 advertised-routes
```

```
BGP table version is 27, local router ID is 172.16.232.181
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i110.0.0.0	172.16.232.179	0	100	0	?
*> 200.2.2.0	0.0.0.0	0		32768	i

The following is sample output from the **show ip bgp neighbors** command with the **routes** keyword:

```
Router# show ip bgp neighbors 172.16.232.178 routes
```

```
BGP table version is 27, local router ID is 172.16.232.181
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.0.0.0	172.16.232.178	40		0	10 ?
*> 20.0.0.0	172.16.232.178	40		0	10 ?

Table 2 describes the significant fields shown in the displays.

Table 2 *show ip bgp neighbors advertised-routes and routes Field Descriptions*

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is damped. h—The table entry history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
Metric	The value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

The following is sample output from the **show ip bgp neighbors** command entered with a specific IP address, the **paths** keyword, and a regular expression:

```
Router# show ip bgp neighbors 171.69.232.178 paths ^10
```

```
Address      Refcount Metric Path
0x60E577B0      2      40 10 ?
```

Table 3 describes the significant fields shown in the display.

Table 3 *show ip bgp neighbors paths Field Descriptions*

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.
Metric	The Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	The autonomous system path for that route, followed by the origin code for that route.

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 © 2003 Cisco Systems, Inc. All rights reserved.