iBGP Multipath Load Sharing

Feature History

Release	Modification
12.2(2)T	This feature was introduced.
12.0(16)ST	This feature was integrated into Cisco IOS Release 12.0(16)ST.
12.2(14)S	This feature was integrated into Cisco IOS Release 12.2(14)S. This feature was enhanced to support multiple redundant paths.
12.0(22)S	This feature was integrated into Cisco IOS Release 12.0(22)S.
12.2(13)T	This feature was enhanced to support multiple redundant paths.
12.0(25)S	This feature was enhanced to support multiple redundant paths.

This feature module describes the iBGP Multipath Load Sharing feature. It includes the following sections:

Feature Overview

When a Border Gateway Protocol (BGP) speaking router with no local policy configured receives multiple network layer reachability information (NLRI) from the internal BGP (iBGP) for the same destination, the router will choose one iBGP path as the best path. The best path is then installed in the IP routing table of the router. For example, in Figure 1, although there are three paths to autonomous system 200, Router 2 determines that one of the paths to autonomous system 200 is the best path and uses this path only to reach autonomous system 200.

Figure 1 Non-MPLS Topology with One Best Path

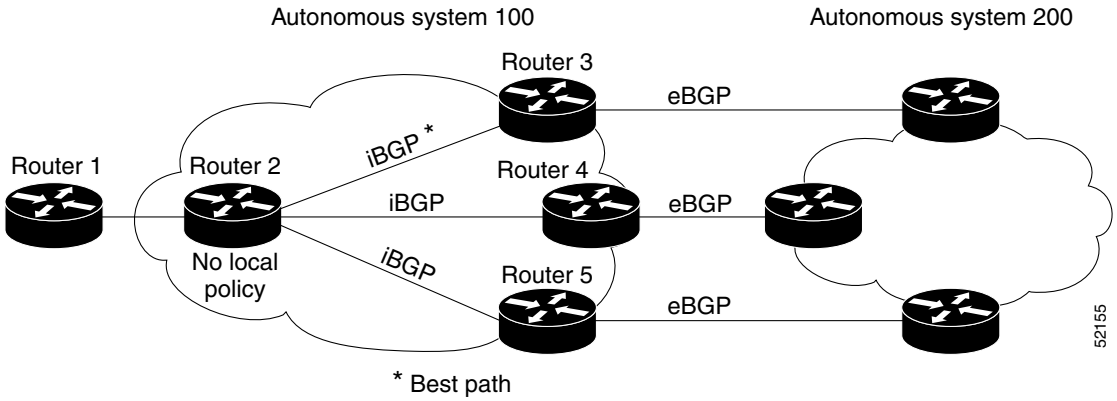

The iBGP Multipath Load Sharing feature enables the BGP speaking router to select multiple iBGP paths as the best paths to a destination. The best paths or multipaths are then installed in the IP routing table of the router. For example, on router 2 in Figure 2, the paths to routers 3, 4, and 5 are configured as multipaths and can be used to reach autonomous system 200, thereby equally sharing the load to autonomous system 200.

Figure 2 Non-MPLS Topology with Three Multipaths

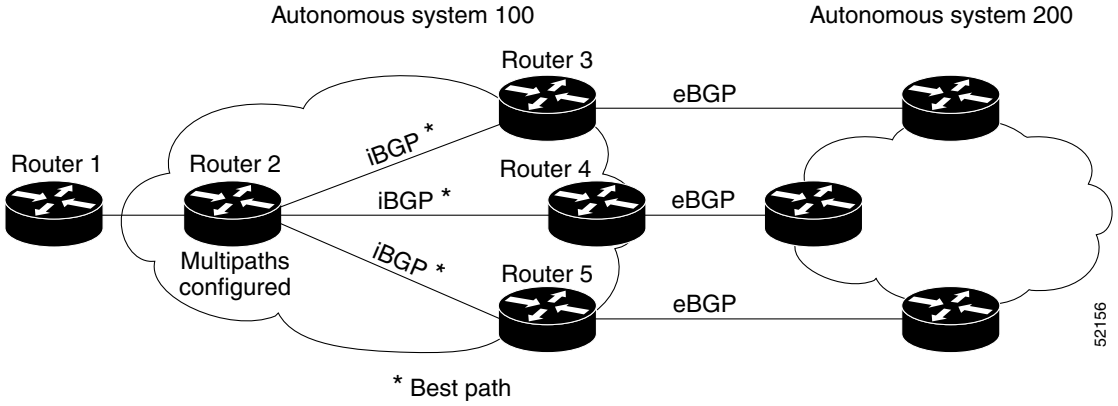

The iBGP Multipath Load Sharing feature functions similarly in a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) with a service provider backbone. For example, on router PE1 in Figure 3, the paths to routers PE2, PE3, and PE4 can be selected as multipaths and can be used to equally share the load to site 2.

Figure 3 MPLS VPN with Three Multipaths

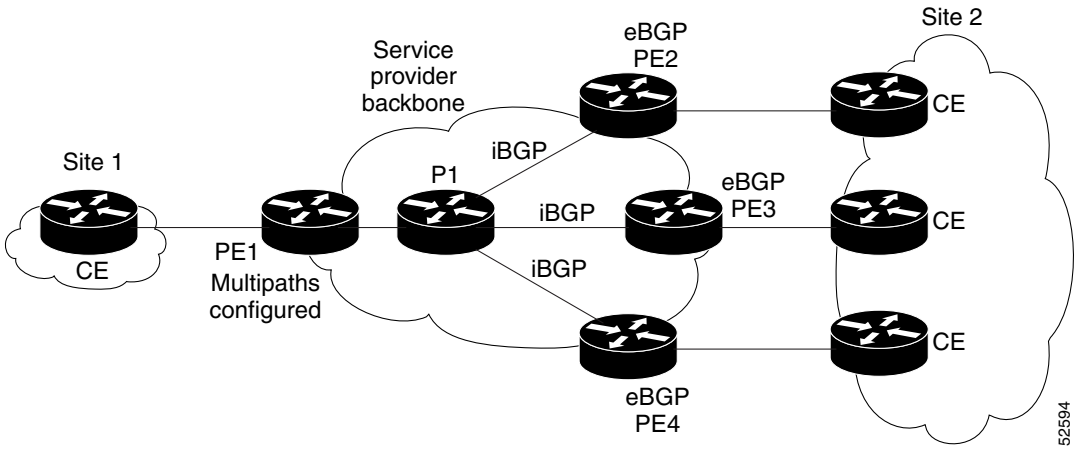

For multiple paths to the same destination to be considered as multipaths, the following criteria must be met:

- All attributes must be the same. The attributes include weight, local preference, autonomous system path (entire attribute and not just length), origin code, Multi Exit Discriminator (MED), and Interior Gateway Protocol (IGP) distance.
- The next hop router for each multipath must be different.

Even if the criteria are met and multiple paths are considered multipaths, the BGP speaking router will still designate one of the multipaths as the best path and advertise this best path to its neighbors.

Benefits

Configuring multiple iBGP best paths enables a router to evenly share the traffic destined for a particular site.

Restrictions

Route Reflector Limitation

With multiple iBGP paths installed in a routing table, a route reflector will advertise only one of the paths (one next hop).

Memory Consumption Restriction

Each IP routing table entry for a BGP prefix that has multiple iBGP paths uses approximately 350 bytes of additional memory. We recommend not using this feature on a router with a low amount of available memory and especially when the router is carrying a full Internet routing table.

Related Features and Technologies

The iBGP Multipath Load Sharing feature is similar to BGP multipath support for external BGP (eBGP) paths; however, the iBGP Multipath Load Sharing feature is applied to internal rather than eBGP paths. BGP multipath support for eBGP paths is documented in the "Configuring BGP" chapter of the *Cisco IOS IP Routing Configuration Guide* and in the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*.

The iBGP Multipath Load Sharing feature is related to the BGP Link Bandwidth feature, which is documented in the "New Features in Release 12.2(14)S" area of Cisco.com.

Related Documents

For related information on this feature, refer to the following documents:

- *Cisco IOS IP Routing Configuration Guide*, Release 12.2.
- *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*
- *BGP Link Bandwidth*

For more information on MPLS VPNs, refer to the following documents:

- *Cisco IOS Switching Services Configuration Guide*, Release 12.2.
- *Cisco IOS Switching Services Command Reference*, Release 12.2.

Supported Platforms

The iBGP Multipath Load Sharing feature is supported for the following platforms in Cisco IOS Release 12.2(14)S:

- Cisco 7200 series
- Cisco 7400 series
- Cisco 7500 series

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that support specific platforms. To get updated information regarding platform support for this feature, access Cisco Feature Navigator. Cisco Feature Navigator dynamically updates the list of supported platforms as new platform support is added for the feature.

Cisco Feature Navigator is a web-based tool that enables you to determine which Cisco IOS software images support a specific set of features and which features are supported in a specific Cisco IOS image. You can search by feature or release. Under the release section, you can compare releases side by side to display both the features unique to each software release and the features in common.

To access Cisco Feature Navigator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions at http://www.cisco.com/register.

Cisco Feature Navigator is updated regularly when major Cisco IOS software releases and technology releases occur. For the most current information, go to the Cisco Feature Navigator home page at the following URL:

http://www.cisco.com/go/fn

Availability of Cisco IOS Software Images

Platform support for particular Cisco IOS software releases is dependent on the availability of the software images for those platforms. Software images for some platforms may be deferred, delayed, or changed without prior notice. For updated information about platform support and availability of software images for each Cisco IOS software release, refer to the online release notes or, if supported, Cisco Feature Navigator.

Supported Standards, MIBs, and RFCs

Standards

No new or modified standards are supported by this feature.

MIBs

No new or modified MIBs are supported 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

No new or modified RFCs are supported by this feature.

Configuration Tasks

See the following sections for configuration tasks for the iBGP Multipath Load Sharing feature. Each task in the list is identified as either required or optional.

- Configuring iBGP Multipath Load Sharing (required)
- Verifying iBGP Multipath Load Sharing (optional)

Configuring iBGP Multipath Load Sharing

To configure the iBGP Multipath Load Sharing feature, use the following command in router configuration mode:

Command	Purpose
`Router(config-router)# maximum-paths ibgp` *maximum-number*	Controls the maximum number of parallel iBGP routes that can be installed in a routing table.

Verifying iBGP Multipath Load Sharing

To verify that the iBGP Multipath Load Sharing feature is configured correctly, perform the following steps:

Step 1 Enter the **show ip bgp** *network-number* EXEC command to display attributes for a network in a non-MPLS topology, or the **show ip bgp vpnv4 all** *ip-prefix* EXEC command to display attributes for a network in an MPLS VPN:

```
Router# show ip bgp 10.22.22.0

BGP routing table entry for 10.22.22.0/24, version 119
Paths:(6 available, best #1)
Multipath:iBGP
Flag:0x820
  Advertised to non peer-group peers:
  10.1.12.12
  22
    10.2.3.8 (metric 11) from 10.1.3.4 (100.0.0.5)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      Originator:100.0.0.5, Cluster list:100.0.0.4
  22
    10.2.1.9 (metric 11) from 10.1.1.2 (100.0.0.9)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Originator:100.0.0.9, Cluster list:100.0.0.2
  22
    10.2.5.10 (metric 11) from 10.1.5.6 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Originator:100.0.0.10, Cluster list:100.0.0.6
  22
    10.2.4.10 (metric 11) from 10.1.4.5 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Originator:100.0.0.10, Cluster list:100.0.0.5
  22
    10.2.6.10 (metric 11) from 10.1.6.7 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Originator:100.0.0.10, Cluster list:100.0.0.7

Router# show ip bgp vpnv4 all 10.22.22.0

BGP routing table entry for 100:1:10.22.22.0/24, version 50
Paths:(6 available, best #1)
Multipath:iBGP
  Advertised to non peer-group peers:
  200.1.12.12
  22
    10.22.7.8 (metric 11) from 10.11.3.4 (100.0.0.8)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      Extended Community:RT:100:1
      Originator:100.0.0.8, Cluster list:100.1.1.44
  22
    10.22.1.9 (metric 11) from 10.11.1.2 (100.0.0.9)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.9, Cluster list:100.1.1.22
  22
    10.22.6.10 (metric 11) from 10.11.6.7 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.10, Cluster list:100.0.0.7
  22
```

```
    10.22.4.10 (metric 11) from 10.11.4.5 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.10, Cluster list:100.0.0.5
  22
    10.22.5.10 (metric 11) from 10.11.5.6 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.10, Cluster list:100.0.0.6
```

Step 2 In the display resulting from the **show ip bgp** *network-number* EXEC command or the **show ip bgp vpnv4 all** *ip-prefix* EXEC command, verify that the intended multipaths (those that meet the criteria discussed in the "Feature Overview" section) are marked as "multipaths." Notice that one of the multipaths is marked as "best."

Step 3 Enter the **show ip route** *ip-address* EXEC command to display routing information for a network in a non-MPLS topology or the **show ip route vrf** *vrf-name ip-prefix* EXEC command to display routing information for a network in an MPLS VPN:

```
Router# show ip route 10.22.22.0

Routing entry for 10.22.22.0/24
  Known via "bgp 1", distance 200, metric 0
  Tag 22, type internal
  Last update from 10.2.6.10 00:00:03 ago
  Routing Descriptor Blocks:
  * 10.2.3.8, from 10.1.3.4, 00:00:03 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.2.1.9, from 10.1.1.2, 00:00:03 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.2.5.10, from 10.1.5.6, 00:00:03 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.2.4.10, from 10.1.4.5, 00:00:03 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.2.6.10, from 10.1.6.7, 00:00:03 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1

Router# show ip route vrf PATH 10.22.22.0

Routing entry for 10.22.22.0/24
  Known via "bgp 1", distance 200, metric 0
  Tag 22, type internal
  Last update from 10.22.5.10 00:01:07 ago
  Routing Descriptor Blocks:
  * 10.22.7.8 (Default-IP-Routing-Table), from 10.11.3.4, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.22.1.9 (Default-IP-Routing-Table), from 10.11.1.2, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.22.6.10 (Default-IP-Routing-Table), from 10.11.6.7, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.22.4.10 (Default-IP-Routing-Table), from 10.11.4.5, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    10.22.5.10 (Default-IP-Routing-Table), from 10.11.5.6, 00:01:07 ago
      Route metric is 0, traffic share count is 1
```

```
AS Hops 1
```

Step 4 Verify that the paths marked as "multipath" in the display resulting from the **show ip bgp** *ip-prefix* EXEC command or the **show ip bgp vpnv4 all** *ip-prefix* EXEC command are included in the routing information. (The routing information is displayed after performing Step 3.)

Monitoring and Maintaining iBGP Multipath Load Sharing

To display iBGP Multipath Load Sharing information, use the following commands in EXEC mode, as needed:

Command	Purpose
`Router# show ip bgp ip-prefix`	Displays attributes and multipaths for a network in a non-MPLS topology.
`Router# show ip bgp vpnv4 all ip-prefix`	Displays attributes and multipaths for a network in an MPLS VPN.
`Router# show ip route ip-prefix`	Displays routing information for a network in a non-MPLS topology.
`Router# show ip route vrf vrf-name ip-prefix`	Displays routing information for a network in an MPLS VPN.

Configuration Examples

This section provides the following configuration examples:

- Non-MPLS Topology Example
- MPLS VPN Topology Example

Both examples assume that the appropriate attributes for each path are equal and that the next hop router for each multipath is different.

Non-MPLS Topology Example

The following example shows how to set up the iBGP Multipath Load Sharing feature in a non-MPLS topology (see Figure 4).

Figure 4 Non-MPLS Topology Example

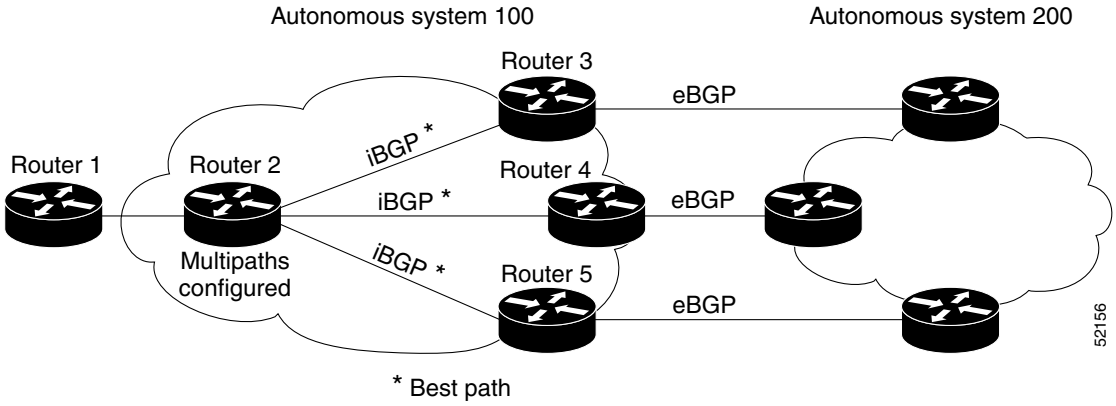

Router 2 Configuration

```
router bgp 100
maximum-paths ibgp 3
```

MPLS VPN Topology Example

The following example shows how to set up the iBGP Multipath Load Sharing feature in an MPLS VPN topology (see Figure 5).

Figure 5 MPLS VPN Topology Example

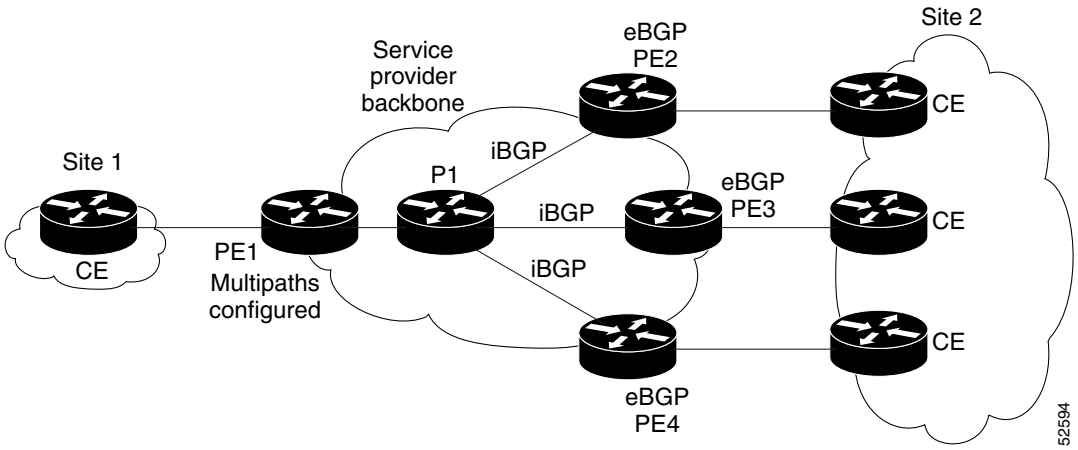

Router PE1 Configuration

```
router bgp 100
address-family ipv4 unicast vrf site2
 maximum-paths ibgp 3
```

Command Reference

This section documents new and modified commands. All other commands used with this feature are documented in the Cisco IOS Release 12.2 command reference publications.

New Commands

- **maximum-paths ibgp**

Modified Commands

- **show ip bgp**
- **show ip bgp vpnv4**
- **show ip route**
- **show ip route vrf**

maximum-paths ibgp

To control the maximum number of parallel internal Border Gateway Protocol (BGP) routes that can be installed in a routing table, use the **maximum-paths ibgp** command in router configuration mode. To restore the default value, use the **no** form of this command.

maximum-paths ibgp *maximum-number*

no maximum-paths ibgp *maximum-number*

Syntax Description

maximum-number	A number from 1 through 8. The maximum number of parallel routes an IP routing protocol installs in a routing table.

Defaults

The default for BGP is one path. The default for all other IP routing protocols is four paths.

Command Modes

Address- family configuration
Router configuration

Command History

Release	Modification
12.2(2)T	This command was introduced.
12.0(16)ST	This command was integrated into Cisco IOS Release 12.0(16)ST.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S. This command was enhanced to support multiple redundant paths.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(13)T	This command was enhanced to support multiple redundant paths.
12.0(25)S	This command was enhanced to support multiple redundant paths.

Usage Guidelines

The iBGP Multipath Load Sharing feature is enabled when the following conditions are met:

- The **maximum-paths ibgp** command must be set to a value greater than one.
- All attributes must be the same. The attributes include weight, local preference, autonomous system path (entire attribute and not just length), origin code, Multi Exit Discriminator (MED), and Interior Gateway Protocol (IGP) distance.
- The next hop router for each multipath must be different.

Even if the criteria are met and multiple paths are considered multipaths, a BGP speaking router will still designate one of the multipaths as the best path and advertise this best path to its neighbors.

Examples

The following example allows a maximum of three parallel iBGP paths in a non-Multiprotocol Label Switching (MPLS) topology:

```
router bgp 100
maximum-paths ibgp 3
```

The following example allows a maximum of three parallel iBGP paths in an MPLS Virtual Private Network (VPN) topology:

```
router bgp 100
address-family ipv4 unicast vrf red
 maximum-paths ibgp 3
```

Related Commands

Command	Description
maximum-paths	Controls the maximum number of parallel routes an IP routing protocol can support.

show ip bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show ip bgp** command in EXEC mode.

show ip bgp [*ip-prefix*] [*network-mask*] [**longer-prefixes**]

Syntax Description

ip-prefix	(Optional) An IP address prefix, entered to display a particular network in the BGP routing table.
network-mask	(Optional) Displays all BGP routes matching the address/mask pair.
longer-prefixes	(Optional) Displays a route and more specific routes.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.0	The display of prefix advertisement statistics was added.
12.0(6)T	The display of a message indicating support for route refresh capability was added.
12.2(2)T	The output of the **show ip bgp** *network-number* command was enhanced to display multipaths and a best path to the specified network.
12.2(14)S	The output of the **show ip bgp** *network-number* command was enhanced to display multipaths and a best path to the specified network.

Examples

The following is sample output from the **show ip bgp** command in privileged EXEC mode:

```
Router# show ip bgp

BGP table version is 5, local router ID is 10.0.33.34
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
*> 10.1.0.0         0.0.0.0                  0         32768 ?
*  10.2.0.0         10.0.33.35              10             0 35 ?
*>                  0.0.0.0                  0         32768 ?
*  10.0.0.0         10.0.33.35              10             0 35 ?
*>                  0.0.0.0                  0         32768 ?
*> 192.168.0.0/16   10.0.33.35              10             0 35 ?
```

Table 1 describes the significant fields shown in the display.

Table 1 show ip bgp 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. *—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 iBGP session.
Origin codes	Indicates the 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	If shown, the value of the interautonomous system metric.
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** command in privileged EXEC mode when you specify the **longer-prefixes** keyword:

```
Router# show ip bgp 192.168.0.0 255.255.0.0 longer-prefixes

BGP table version is 1738, local router ID is 192.168.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 198.168.0.0     198.168.72.30          8896         32768 ?
*                  198.168.72.30                           0 109 108 ?
*> 198.168.1.0     198.168.72.30          8796         32768 ?
*                  198.168.72.30                           0 109 108 ?
*> 198.168.11.0    198.168.72.30         42482         32768 ?
*                  198.168.72.30                           0 109 108 ?
*> 198.168.14.0    198.168.72.30          8796         32768 ?
*                  198.168.72.30                           0 109 108 ?
```

```
*> 198.168.15.0     198.168.72.30        8696         32768 ?
*                   198.168.72.30                         0 109 108 ?
*> 198.168.16.0     198.168.72.30        1400         32768 ?
*                   198.168.72.30                         0 109 108 ?
*> 198.168.17.0     198.168.72.30        1400         32768 ?
*                   198.168.72.30                         0 109 108 ?
*> 198.168.18.0     198.168.72.30        8876         32768 ?
*                   198.168.72.30                         0 109 108 ?
*> 198.168.19.0     198.168.72.30        8876         32768 ?
*                   198.168.72.30                         0 109 108 ?
```

The following is sample output from the **show ip bgp** command in privileged EXEC mode that shows information including multipaths and a best path to network 10.22.22.0:

```
Router# show ip bgp 10.22.22.0

BGP routing table entry for 10.22.22.0/24, version 119
Paths:(6 available, best #1)
Multipath:iBGP
Flag:0x820
  Advertised to non peer-group peers:
  10.1.12.12
  22
    10.2.3.8 (metric 11) from 10.1.3.4 (100.0.0.5)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      Originator:100.0.0.5, Cluster list:100.0.0.4
  22
    10.2.1.9 (metric 11) from 10.1.1.2 (100.0.0.9)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Originator:100.0.0.9, Cluster list:100.0.0.2
  22
    10.2.5.10 (metric 11) from 10.1.5.6 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Originator:100.0.0.10, Cluster list:100.0.0.6
  22
    10.2.4.10 (metric 11) from 10.1.4.5 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Originator:100.0.0.10, Cluster list:100.0.0.5
  22
    10.2.6.10 (metric 11) from 10.1.6.7 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Originator:100.0.0.10, Cluster list:100.0.0.7
```

Note If a prefix has not been advertised to any peer, the display shows “Not advertised to any peer.”

Table 2 describes the significant fields shown in the display.

Table 2 show ip bgp Field Descriptions

Field	Description
BGP routing table version	Internal version number of the table. This number is incremented whenever the table changes.
Paths:	Number of the autonomous system paths to the specified network. If multiple paths exist, one of the multipaths is designated the best path.
Multipath:	Indicates the maximum number of paths configured (iBGP or eBGP).
Flag:	Record of attribute changes in best path transition.

Table 2 show ip bgp Field Descriptions (continued)

Field	Description
Advertised to non peer-group peers: 10.1.12.12	IP address of the BGP peers to which the specified route is advertised.
10.2.3.8 (metric 11) from 10.1.3.4 (100.0.0.5)	Indicates the next hop IP address and metric value and the IP address and router ID of the interface that sent the update.
Origin	Indicates the origin of the entry. It can be one of the following values: IGP—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. incomplete—Entry originated from other than an IGP or Exterior Gateway Protocol (EGP) and was advertised with the **redistribute** router configuration command. EGP—Entry originated from an EGP.
metric	If shown, the value of the interautonomous system metric.
localpref	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
valid	Indicates that the route is usable and has a valid set of attributes.
internal/external	The field is *internal* if the path is learned via iBGP. The field is *external* if the path is learned via eBGP.
multipath	One of multiple paths to the specified network.
best	If multiple paths exist, one of the multipaths is designated the best path and advertised the neighbors.
Originator:	The router ID of the route originating router when route reflector is used.
Cluster list:	The router ID of all the route reflectors through which the specified route has passed.

Related Commands

Command	Description
clear ip bgp	Resets a BGP connection or session.
neighbor soft-reconfiguration	Configures the Cisco IOS software to start storing updates.

show ip bgp vpnv4

To display Virtual Private Network (VPN) address information from the Border Gateway Protocol (BGP) table, use the **show ip bgp vpnv4** command in EXEC mode.

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**] [**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**neighbors**] [**paths** [*line*]] [**peer-group**] [**quote-regexp**] [**regexp**] [**summary**] [**tags**]

Syntax Description

all	Displays the complete VPNv4 database.
rd *route-distinguisher*	Displays NLRIs that have a matching route distinguisher.
vrf *vrf-name*	Displays NLRIs associated with the named virtual routing and forwarding instance (VRF).
ip-prefix/length	(Optional) The IP prefix address (in dotted decimal format) and the length of the mask (0 to 32). The slash mark must be included.
longer-prefixes	(Optional) Displays the entry, if any, that exactly matches the specified prefix parameter and all entries that match the prefix in a "longest-match" sense. That is, prefixes for which the specified prefix is an initial substring.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
network-address	(Optional) The IP address of a network in the BGP routing table.
mask	(Optional) The mask of the network address, in dotted decimal format.
cidr-only	(Optional) Displays only routes that have nonnatural net masks.
community	(Optional) Displays routes matching this community.
community-list	(Optional) Displays routes matching this community list.
dampened-paths	(Optional) Displays paths suppressed on account of dampening (BGP route from peer is up and down).
filter-list	(Optional) Displays routes conforming to the filter list.
flap-statistics	(Optional) Displays flap statistics of routes.
inconsistent-as	(Optional) Displays only routes that have inconsistent autonomous systems of origin.
neighbors	(Optional) Displays details about TCP and BGP neighbor connections.
paths	(Optional) Displays path information.
line	(Optional) A regular expression to match the BGP autonomous system paths.
peer-group	(Optional) Displays information about peer groups.
quote-regexp	(Optional) Displays routes matching the autonomous system path "regular expression."
regexp	(Optional) Displays routes matching the autonomous system path regular expression.

summary	(Optional) Displays BGP neighbor status.
tags	(Optional) Displays incoming and outgoing BGP labels for each NLRI.

Defaults

No default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	The output of the **show ip bgp vpnv4 all** *ip-prefix* command was enhanced to display attributes including multipaths and a best path to the specified network.
12.2(14)S	The output of the **show ip bgp vpnv4 all** *ip-prefix* command was enhanced to display attributes including multipaths and a best path to the specified network.

Usage Guidelines

Use this command to display VPNv4 information from the BGP database. The **show ip bgp vpnv4 all** command displays all available VPNv4 information. The **show ip bgp vpnv4 summary** command displays BGP neighbor status.

Examples

The following example shows output for all available VPNv4 information in a BGP routing table:

```
Router# show ip bgp vpnv4 all

BGP table version is 18, local router ID is 10.14.14.14
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
Route Distinguisher: 100:1 (vrf1)
*> 10.1.0.0         10.5.0.1 0 0 101 i
*>i10.2.0.0         10.13.13.13 0    100 0 102 i
*> 10.5.0.0         10.5.0.1 0 0 101 i
*>i10.6.0.0         10.13.13.13 0    100 0 102 i
```

Table 3 describes the significant fields shown in the display.

Table 3 show ip bgp vpnv4 Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.

Table 3 show ip bgp vpnv4 Field Descriptions

Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows how to display a table of labels for NLRIs that have a route distinguisher value of 100:1.

```
Router# show ip bgp vpnv4 rd 100:1 tags

    Network    Next Hop      In tag/Out tag
Route Distinguisher: 100:1 (vrf1)
   10.2.0.0         10.20.0.60      34/notag
   10.0.0.0         10.20.0.60      35/notag
   10.12.0.0        10.20.0.60      26/notag
                    10.20.0.60      26/notag
   10.13.0.0        10.15.0.15      notag/26
```

Table 4 describes the significant fields shown in the display.

Table 4 show ip bgp vpnv4 rd tags Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Specifies the BGP next hop address.
In	Displays the label (if any) assigned by this router.
Out	Displays the label assigned by the BGP next hop router.

The following example shows VPNv4 routing entries for the VRF named vrf1:

```
Router# show ip bgp vpnv4 vrf vrf1

BGP table version is 18, local router ID is 10.14.14.14
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
Route Distinguisher: 100:1 (vrf1)
*> 10.11.0.0         50.0.0.1 0 0 101 i
*>i10.12.0.0         13.13.13.13 0    100 0 102 i
*> 10.50.0.0         50.0.0.1 0 0 101 i
*>i10.51.0.0         10.13.13.13 0    100 0 102 i
```

Table 5 describes the significant fields shown in the display.

Table 5 show ip bgp vpnv4 vrf Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.

Table 5 show ip bgp vpnv4 vrf Field Descriptions

Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows attributes for network 10.22.22.0 that includes multipaths and a best path:

```
Router# show ip bgp vpnv4 all 10.22.22.0

BGP routing table entry for 100:1:10.22.22.0/24, version 50
Paths:(6 available, best #1)
Multipath:iBGP
  Advertised to non peer-group peers:
  200.1.12.12
  22
    10.22.7.8 (metric 11) from 10.11.3.4 (100.0.0.8)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      Extended Community:RT:100:1
      Originator:100.0.0.8, Cluster list:100.1.1.44
  22
    10.22.1.9 (metric 11) from 10.11.1.2 (100.0.0.9)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.9, Cluster list:100.1.1.22
  22
    10.22.6.10 (metric 11) from 10.11.6.7 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.10, Cluster list:100.0.0.7
  22
    10.22.4.10 (metric 11) from 10.11.4.5 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.10, Cluster list:100.0.0.5
  22
    10.22.5.10 (metric 11) from 10.11.5.6 (100.0.0.10)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:100.0.0.10, Cluster list:100.0.0.6
```

Table 6 describes the significant fields shown in the display.

Table 6 show ip bgp vpn4 all 10.22.22.0 Field Descriptions

Field	Description
BGP routing table ... version	Internal version number of the table. This number is incremented whenever the table changes.
Paths:	Number of autonomous system paths to the specified network. If multiple paths exist, one of the multipaths is designated the best path.
Multipath:	Indicates the maximum number of paths paths configured (iBGP or eBGP).
Advertised to non peer-group peers: 200.1.12.12 22	IP address of the BGP peers to which the specified route is advertised.

Table 6 show ip bgp vpn4 all 10.22.22.0 Field Descriptions (continued)

Field	Description
10.22.7.8 (metric 11) from 10.11.3.4 (100.0.0.8)	Indicates the next hop address and the address of the gateway that sent the update.
Origin	Indicates the origin of the entry. It can be one of the following values: IGP—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. incomplete—Entry originated from other than an IGP or Exterior Gateway Protocol (EGP) and was advertised with the redistribute router configuration command. EGP—Entry originated from an EGP.
metric	If shown, the value of the interautonomous system metric.
localpref	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
valid	Indicates that the route is usable and has a valid set of attributes.
internal/external	The field is *internal* if the path is learned via iBGP. The field is *external* if the path is learned via eBGP.
multipath	One of multiple paths to the specified network.
best	If multiple paths exist, one of the multipaths is designated the best path and advertised the neighbors.
Extended Community:RT:100:1	Route target value associated with the specified route.
Originator:	The router ID of the route originating router when route reflector is used.
Cluster list:	The router ID of all the route reflectors through which the specified route has passed.

Related Commands

Command	Description
show ip vrf	Displays the set of defined VRFs and associated interfaces.

show ip route

To display the current state of the routing table, use the **show ip route** command in EXEC mode.

show ip route [*ip-address* [*mask*] [**longer-prefixes**]] | [*protocol* [*process-id*]]

Syntax Description

ip-address	(Optional) Address about which routing information should be displayed.
mask	(Optional) Argument for a subnet mask.
longer-prefixes	(Optional) Specifies that routes matching the *ip-address* and *mask* pair only should be displayed.
protocol	(Optional) The name of a routing protocol; or the keyword **connected**, **static**, or **summary**. If you specify a routing protocol, use one of the following keywords: **bgp**, **egp**, **eigrp**, **hello**, **igrp**, **isis**, **ospf**, and **rip**.
process-id	(Optional) The number used to identify a process of the specified protocol.

Command Modes

EXEC

Command History

Release	Modification
9.2	This command was introduced.
10.0	The “D—EIGRP, EX—EIGRP, N1—OSPF NSSA external type 1 route” and “N2—OSPF NSSA external type 2 route” codes were added to the command output.
10.3	The *process-id* argument was added.
11.0	The **longer-prefixes** keyword was added.
11.1	The “U—per-user static route” code was added to the command output.
11.2	The “o—on-demand routing” code was added to the command output.
11.3	The output of the **show ip route** *ip-address* command was enhanced to display the origination of an IP route in Intermediate System-to-Intermediate System (IS-IS) networks.
12.0(1)T	The “M—mobile” code was added to the command output.
12.0(3)T	The “P—periodic downloaded static route” code was added to the command output.
12.0(4)T	The “ia—IS-IS” code was added to the command output.
12.2(2)T	The output of the **show ip route** *ip-address* command was enhanced to display information on the multipaths to the specified network.
12.2(14)S	The output of the **show ip route** *ip-address* command was enhanced to display information on the multipaths to the specified network.

Examples

The following is sample output from the **show ip route** command when entered without an address:

```
Router# show ip route

Codes: I - IGRP derived, R - RIP derived, O - OSPF derived,
       C - connected, S - static, E - EGP derived, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route

Gateway of last resort is 10.119.254.240 to network 10.140.0.0

O E2 10.150.0.0 [160/5] via 10.119.254.6, 0:01:00, Ethernet2
E    10.67.131.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
O E2 10.68.132.0 [160/5] via 10.119.254.6, 0:00:59, Ethernet2
O E2 10.130.0.0 [160/5] via 10.119.254.6, 0:00:59, Ethernet2
E    10.128.0.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.129.0.0 [200/129] via 10.119.254.240, 0:02:22, Ethernet2
E    10.65.129.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.131.0.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.75.139.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
E    10.16.208.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.84.148.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
E    10.31.223.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.44.236.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
E    10.141.0.0 [200/129] via 10.119.254.240, 0:02:22, Ethernet2
E    10.140.0.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
```

The following is sample output that includes IS-IS Level 2 routes learned:

```
Router# show ip route

Codes: I - IGRP derived, R - RIP derived, O - OSPF derived,
       C - connected, S - static, E - EGP derived, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route

Gateway of last resort is not set

10.89.0.0 is subnetted (mask is 255.255.255.0), 3 subnets
C       10.89.64.0 255.255.255.0 is possibly down,
          routing via 0.0.0.0, Ethernet0
i L2    10.89.67.0 [115/20] via 10.89.64.240, 0:00:12, Ethernet0
i L2    10.89.66.0 [115/20] via 10.89.64.240, 0:00:12, Ethernet0
```

Table 7 describes the significant fields shown in the displays.

Table 7 show ip route Field Descriptions

Field	Description
O	Indicates the protocol that derived the route. It can be one of the following values: I—Interior Gateway Routing Protocol (IGRP) derived R—Routing Information Protocol (RIP) derived O—Open Shortest Path First (OSPF) derived C—connected S—static E—Exterior Gateway Protocol (EGP) derived B—Border Gateway Protocol (BGP) derived D—Enhanced Interior Gateway Routing Protocol (EIGRP) EX—EIGRP external i—IS-IS derived ia—IS-IS M—mobile P—periodic downloaded static route U—per-user static route o—on-demand routing
E2	Type of route. It can be one of the following values: *—Indicates the last path used when a packet was forwarded. It pertains only to the nonfast-switched packets. However, it does not indicate which path will be used next when forwarding a nonfast-switched packet, except when the paths are equal cost. IA—OSPF interarea route E1—OSPF external type 1 route E2—OSPF external type 2 route L1—IS-IS Level 1 route L2—IS-IS Level 2 route N1—OSPF not so stubby area (NSSA) external type 1 route N2—OSPF NSSA external type 2 route
150.150.0.0	Indicates the address of the remote network.
[160/5]	The first number in the brackets is the administrative distance of the information source; the second number is the metric for the route.
via 131.119.254.6	Specifies the address of the next router to the remote network.
0:01:00	Specifies the last time the route was updated (in hours:minutes:seconds).
Ethernet2	Specifies the interface through which the specified network can be reached.

When you specify that you want information about a specific network displayed, more detailed statistics are shown. The following is sample output from the **show ip route** command when entered with the address 10.0.0.1:

```
Router# show ip route 10.0.0.1

Routing entry for 10.0.0.1/32
    Known via "isis", distance 115, metric 20, type level-1
    Redistributing via isis
    Last update from 10.191.255.251 on Fddi1/0, 00:00:13 ago
    Routing Descriptor Blocks:
    * 10.22.22.2, from 10.191.255.247, via Serial2/3
       Route metric is 20, traffic share count is 1
       10.191.255.251, from 10.191.255.247, via Fddi1/0
       Route metric is 20, traffic share count is 1
```

When an IS-IS router advertises its link-state information, it includes one of its own IP addresses to be used as the originator IP address. When other routers calculate IP routes, they can store the originator IP address with each route in the routing table.

The following example shows the output from the **show ip route** command when looking at an IP route generated by IS-IS. Each path that is shown under the Routing Descriptor Blocks report displays two IP addresses. The first address (10.22.22.2) is the next hop address. The second is the originator IP address from the advertising IS-IS router. This address helps you determine where a particular IP route has originated in your network. In the example the route to 10.0.0.1/32 was originated by a router with IP address 10.191.255.247.

```
Router# show ip route 10.0.0.1

Routing entry for 10.0.0.1/32
    Known via "isis", distance 115, metric 20, type level-1
    Redistributing via isis
    Last update from 10.191.255.251 on Fddi1/0, 00:00:13 ago
    Routing Descriptor Blocks:
    * 10.22.22.2, from 10.191.255.247, via Serial2/3
       Route metric is 20, traffic share count is 1
       10.191.255.251, from 10.191.255.247, via Fddi1/0
       Route metric is 20, traffic share count is 1
```

Table 8 describes the significant fields shown when using the **show ip route** command with an IP address (previous displays).

Table 8 show ip route with Address Field Descriptions

Field	Description
Routing entry for[*ip-address* [*mask*]	Network number and mask.
Known via...	Indicates how the route was derived.
Tag	Integer that is used to implement the route.
type	Indicates the IS-IS route type (Level1 or Level2).
Redistributing via...	Indicates the redistribution protocol.
Last update from *ip-address*	Indicates the IP address of a router that is the next hop to the remote network and the router interface on which the last update arrived.
Routing Descriptor Blocks:	Displays the next hop IP address followed by the information source.

Table 8 show ip route with Address Field Descriptions (continued)

Field	Description
Route metric	This value is the best metric for this routing descriptor block.
traffic share count	Number of uses for this routing descriptor block.

The following is sample output using the **longer-prefixes** keyword. When the **longer-prefixes** keyword is included, the address and mask pair becomes the prefix, and any address that matches that prefix is displayed. Therefore, multiple addresses are displayed.

In the following example, the logical AND operation is performed on the source address 10.0.0.0 and the mask 10.0.0.0, resulting in 10.0.0.0. Each destination in the routing table is also logically ANDed with the mask and compared to that result of 10.0.0.0. Any destinations that fall into that range are displayed in the output.

```
Router# show ip route 10.0.0.0 10.0.0.0 longer-prefixes

Codes: I - IGRP derived, R - RIP derived, O - OSPF derived,
       C - connected, S - static, E - EGP derived, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route

Gateway of last resort is not set

S    10.134.0.0 is directly connected, Ethernet0
S    10.131.0.0 is directly connected, Ethernet0
S    10.129.0.0 is directly connected, Ethernet0
S    10.128.0.0 is directly connected, Ethernet0
S    10.49.246.0 is directly connected, Ethernet0
S    10.160.97.0 is directly connected, Ethernet0
S    10.153.88.0 is directly connected, Ethernet0
S    10.76.141.0 is directly connected, Ethernet0
S    10.75.138.0 is directly connected, Ethernet0
S    10.44.237.0 is directly connected, Ethernet0
S    10.31.222.0 is directly connected, Ethernet0
S    10.16.209.0 is directly connected, Ethernet0
S    10.145.0.0 is directly connected, Ethernet0
S    10.141.0.0 is directly connected, Ethernet0
S    10.138.0.0 is directly connected, Ethernet0
S    10.128.0.0 is directly connected, Ethernet0
     10.19.0.0 255.255.255.0 is subnetted, 1 subnets
C       10.19.64.0 is directly connected, Ethernet0
     10.69.0.0 is variably subnetted, 2 subnets, 2 masks
C       10.69.232.32 255.255.255.240 is directly connected, Ethernet0
S       10.69.0.0 255.255.0.0 is directly connected, Ethernet0
```

Related Commands

Command	Description
show interfaces tunnel	Displays a list of tunnel interface information.
show ip route summary	Displays the current state of the routing table in summary format.

show ip route vrf

To display the IP routing table associated with a VPN routing and forwarding instance (VRF), use the **show ip route vrf** command in EXEC mode.

show ip route vrf *vrf-name* [**connected**] [*protocol* [*as-number*] [*tag*] [*output-modifiers*]] [*ip-prefix*] [**list** *number* [*output-modifiers*]] [**profile**] [**static** [*output-modifiers*]] [**summary** [*output-modifiers*]] [**supernets-only** [*output-modifiers*]]

Syntax Description

vrf-name	The name assigned to the VRF.
connected	(Optional) Displays all connected routes in a VRF.
protocol	(Optional) To specify a routing protocol, use one of the following keywords: **bgp**, **egp**, **eigrp**, **hello**, **igrp**, **isis**, **ospf**, or **rip**.
as-number	(Optional) The autonomous system number.
tag	(Optional) Cisco IOS routing area label.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
ip-prefix	(Optional) Specifies a network to display.
list *number*	(Optional) Specifies the IP access list to display.
profile	(Optional) Displays the IP routing table profile.
static	(Optional) Displays static routes.
summary	(Optional) Displays a summary of routes.
supernets-only	(Optional) Displays supernet entries only.

Defaults

No default behavior or values.

Command Modes

EXEC

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	The *ip-prefix* argument was added. The output of the **show ip route vrf** *vrf-name ip-prefix* command was enhanced to display information on the multipaths to the specified network.
12.2(14)S	The *ip-prefix* argument was added. The output of the **show ip route vrf** *vrf-name ip-prefix* command was enhanced to display information on the multipaths to the specified network.

Usage Guidelines

This command displays specified information from the IP routing table of a VRF.

Examples

This example shows the IP routing table associated with the VRF called vrf1:

```
Router# show ip route vrf vrf1

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR
       T - traffic engineered route

Gateway of last resort is not set

B    51.0.0.0/8 [200/0] via 13.13.13.13, 00:24:19
C    50.0.0.0/8 is directly connected, Ethernet1/3
B    11.0.0.0/8 [20/0] via 50.0.0.1, 02:10:22
B    12.0.0.0/8 [200/0] via 13.13.13.13, 00:24:20
```

This example shows BGP entries in the IP routing table associated with the VRF named vrf1:

```
Router# show ip route vrf vrf1 bgp

B  51.0.0.0/8 [200/0] via 13.13.13.13, 03:44:14
B  11.0.0.0/8 [20/0] via 51.0.0.1, 03:44:12
B  12.0.0.0/8 [200/0] via 13.13.13.13, 03:43:14
```

This example shows the IP routing table associated with a VRF named PATH and network 10.22.22.0:
Router# **show ip route vrf PATH 10.22.22.0**

```
Routing entry for 10.22.22.0/24
  Known via "bgp 1", distance 200, metric 0
  Tag 22, type internal
  Last update from 10.22.5.10 00:01:07 ago
  Routing Descriptor Blocks:
  * 1.22.7.8 (Default-IP-Routing-Table), from 1.11.3.4, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    1.22.1.9 (Default-IP-Routing-Table), from 1.11.1.2, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    1.22.6.10 (Default-IP-Routing-Table), from 1.11.6.7, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    1.22.4.10 (Default-IP-Routing-Table), from 1.11.4.5, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
    1.22.5.10 (Default-IP-Routing-Table), from 1.11.5.6, 00:01:07 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
```

Table 9 describes the significant fields shown when using the **show ip route vrf** *vrf-name ip-prefix* command.

Table 9 show ip route vrf Field Descriptions

Field	Description
Routing entry for 10.22.22.0/24	Network number.
Known via ...	Indicates how the route was derived.
distance	Administrative distance of the information source.
metric	The metric to reach the destination network.
Tag	Integer that is used to implement the route.
type	Indicates that the route is a L1 type or L2 type route.
Last update from *10.22.5.10*	Indicates the IP address of a router that is the next hop to the remote network and the router interface on which the last update arrived.
hh:*mm*:*ss* ago	Specifies the last time the route was updated (in hours:minutes:seconds).
Routing Descriptor Blocks:	Displays the next hop IP address followed by the information source.
ip-address, from *ip-address*, *hh*:*mm*:*ss* ago	Indicates the next hop address, the address of the gateway that sent the update, and the time that has elapsed since this update was received (in hours:minutes:seconds).
Route metric	This value is the best metric for this routing descriptor block.
traffic share count	Number of uses for this routing descriptor block.
AS Hops	Hops to the destination or to the router where the route first enters iBGP.

Related Commands

Command	Description
show ip cache	Displays the CEF forwarding table associated with a VRF.
show ip vrf	Displays the set of defined VRFs and associated interfaces.