



BGP Support for Next-Hop Address Tracking

The BGP Support for Next-Hop Address Tracking feature introduces an event driven notification system to monitor the status of routes that are installed in the Routing Information Base (RIB) and to report next-hop changes that affect internal BGP (iBGP) or external BGP (eBGP) prefixes directly to the Border Gateway Protocol (BGP) process. This feature improves the overall BGP convergence time by allowing BGP to respond rapidly to next-hop changes for routes installed in the RIB.

Feature History for the BGP Support for Next-Hop Address Tracking feature

Release	Modification
12.0(29)S	This feature was introduced.
12.0(31)S	The default delay interval value was changed from 1 to 5 seconds.
12.3(14)T	This feature was integrated into Cisco IOS Release 12.3(14)T.

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 Support for Next-Hop Address Tracking, page 2](#)
- [Restrictions for BGP Support for Next-Hop Address Tracking, page 2](#)
- [Information About BGP Support for Next-Hop Address Tracking, page 2](#)
- [How to Configure BGP Next-Hop Address Tracking, page 2](#)
- [Configuration Examples for BGP Next-Hop Address Tracking, page 5](#)
- [Additional References, page 6](#)
- [Command Reference, page 7](#)



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

Copyright © 2004 Cisco Systems, Inc. All rights reserved.

Prerequisites for BGP Support for Next-Hop Address Tracking

This document assumes that BGP is enabled and peering has been established.

Restrictions for BGP Support for Next-Hop Address Tracking

This feature is not supported under the IPv6 address family.

Information About BGP Support for Next-Hop Address Tracking

- [Default BGP Scanner Behavior, page 2](#)
- [BGP Support for Next-Hop Address Tracking, page 2](#)

Default BGP Scanner Behavior

BGP monitors the next hop of installed routes to verify next-hop reachability and to select, install, and validate the BGP best path. By default, the BGP scanner is used to poll the RIB for this information every 60 seconds. During the 60 second time period between scan cycles, Interior Gateway Protocol (IGP) instability or other network failures can cause black holes and routing loops to temporarily form.

BGP Support for Next-Hop Address Tracking

The BGP Support for Next-Hop Address Tracking feature is enabled by default when a supporting Cisco IOS software image is installed. BGP next-hop address tracking is event driven. BGP prefixes are automatically tracked as peering sessions are established. Next-hop changes are rapidly reported to the BGP routing process as they are updated in the RIB. This optimization improves overall BGP convergence by reducing the response time to next-hop changes for routes installed in the RIB. When a bestpath calculation is run in between BGP scanner cycles, only next-hop changes are tracked and processed.

How to Configure BGP Next-Hop Address Tracking

The tasks in this section show how configure BGP next-hop address tracking. BGP next-hop address tracking significantly improves the response time of BGP to next-hop changes in the RIB. However, unstable Interior Gateway Protocol (IGP) peers can introduce instability to BGP neighbor sessions. We recommend that you aggressively dampen unstable IGP peering sessions to reduce the possible impact to BGP. For more details about route dampening, see the

- [Disabling BGP Next-Hop Address Tracking, page 3](#)
- [Adjusting the Delay Interval for BGP Next-Hop Address Tracking, page 4](#)
- [Configuration Examples for BGP Next-Hop Address Tracking, page 5](#)

Disabling BGP Next-Hop Address Tracking

Perform this task to disable BGP next-hop address tracking. BGP next-hop address tracking is enabled by default under the IPv4 and VPNv4 address families. Disabling next hop address tracking may be useful if you the network has unstable IGP peers and route dampening is not resolving the stability issues. To reenable BGP next-hop address tracking, use the **bgp nexthop** command with the **trigger** and **enable** keywords.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *as-number*
4. **address-family ipv4** [**mdt** | **multicast** | **tunnel** | **unicast** [**vrf** *vrf-name*] | **vrf** *vrf-name*] | **vpn4** [**unicast**]
5. **no bgp nexthop trigger enable**
6. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. • 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 64512	Enters router configuration mod to create or configure a BGP routing process.
Step 4	address-family ipv4 [mdt multicast tunnel unicast [vrf <i>vrf-name</i>] vrf <i>vrf-name</i>] vpn4 [unicast] Example: Router(config-router-af)# address-family ipv4 unicast	Enter address family configuration mode to configure BGP peers to accept address family-specific configurations. • The example creates an IPv4 unicast address family session.

	Command or Action	Purpose
Step 5	<pre>no bgp nexthop trigger enable</pre> <p>Example: Router(config-router-af)# no bgp nexthop trigger enable</p>	Disables BGP next-hop address tracking. <ul style="list-style-type: none"> • Next-hop address tracking is enabled by default for IPv4 and VPNv4 address family sessions. • The example disables next-hop address tracking.
Step 6	<pre>end</pre> <p>Example: Router(config-router-af)# end</p>	Exits address-family configuration mode, and enters Privileged EXEC mode.

Adjusting the Delay Interval for BGP Next-Hop Address Tracking

Perform this task to adjust the delay interval between routing table walks for BGP next-hop address tracking.

Delay Interval Tuning to Match the Interior Gateway Protocol

You can increase the performance of this feature by tuning the delay interval between full routing table walks to match the tuning parameters for the Interior Gateway protocol (IGP). The default delay interval is 5 seconds. This value is optimal for a fast-tuned IGP. In the case of an IGP that converges more slowly, you can change the delay interval to 20 seconds or more, depending on the IGP convergence time.

Aggressive IGP Route Dampening

BGP next-hop address tracking significantly improves the response time of BGP to next-hop changes in the RIB. However, unstable Interior Gateway Protocol (IGP) peers can introduce instability to BGP neighbor sessions. We recommend that you aggressively dampen unstable IGP peering sessions to reduce the possible impact to BGP.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *as-number*
4. **address-family ipv4** [**mdt** | **multicast** | **tunnel** | **unicast** [**vrf vrf-name**] | **vrf vrf-name**] | **vpn4** [**unicast**]
5. **bgp nexthop trigger delay** *delay-timer*
6. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>enable</code> Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none">Enter your password if prompted.
Step 2	<code>configure terminal</code> Example: Router# configure terminal	Enters global configuration mode.
Step 3	<code>router bgp as-number</code> Example: Router(config)# router bgp 64512	Enters router configuration mode to create or configure a BGP routing process.
Step 4	<code>address-family ipv4 [mdt multicast tunnel unicast [vrf vrf-name] vrf vrf-name] vpnv4 [unicast]</code> 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 creates an IPv4 unicast address family session.
Step 5	<code>bgp nexthop trigger delay delay-timer</code> Example: Router(config-router-af)# bgp nexthop trigger delay 20	Configures the delay interval between routing table walks for next-hop address tracking. <ul style="list-style-type: none">The time period determines how long BGP will wait before starting a full routing table walk after notification is received.The value for the <i>delay-timer</i> argument is a number from 1 to 100 seconds. The default value is 5 second.The example configures a delay interval of 20 seconds.
Step 6	<code>end</code> Example: Router(config-router-af)# end	Exits address-family configuration mode, and enters privileged EXEC mode.

Configuration Examples for BGP Next-Hop Address Tracking

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

- [Enabling and Disabling BGP Next-Hop Address Tracking: Example, page 5](#)
- [Adjusting the Delay Interval for BGP Next-Hop Address Tracking: Example, page 6](#)

Enabling and Disabling BGP Next-Hop Address Tracking: Example

In the following example, next-hop address tracking is disabled under the IPv4 address family session:

```
router bgp 50000
 address-family ipv4 unicast
  no bgp nexthop trigger enable
```

Adjusting the Delay Interval for BGP Next-Hop Address Tracking: Example

In the following example, the delay interval for next-hop tracking is configured to occur every 20 seconds under the IPv4 address family session:

```
router bgp 50000
 address-family ipv4 unicast
  bgp nexthop trigger delay 20
```

Where to Go Next

The BGP Support for Fast Peering Session Deactivation feature can also be configured to improve the response time of BGP to adjacency changes, improving overall BGP convergence. For information about BGP fast session deactivation, see the [BGP Support for Fast Peering Session Deactivation](#) feature.

Additional References

The following sections provide references related to the BGP Support for Next-Hop Address Tracking feature.

Related Documents

Related Topic	Document Title
BGP commands	<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
BGP Support for Fast Peering Session Deactivation	<ul style="list-style-type: none"> BGP Support for Fast Peering Session Deactivation
IP Event Dampening	<ul style="list-style-type: none"> IP Event Dampening

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 the following new commands:

- **bgp nexthop trigger enable**
- **bgp nexthop trigger delay**

bgp nexthop trigger enable

To enable Border Gateway Protocol (BGP) next-hop address tracking, use the **bgp nexthop trigger enable** command in address family or router configuration mode. To disable BGP next-hop address tracking, use the **no** form of this command.

bgp nexthop trigger enable

no bgp nexthop trigger enable

Syntax Description This command has no keywords or arguments.

Defaults BGP next-hop address tracking is enabled by default under IPv4 and VPNv4 address-family sessions.

Command Modes Address family configuration
Router Configuration

Command History	Release	Modification
	12.0(29)S	This command was introduced.
	12.3(14)T	This command was integrated into Cisco IOS Release 12.3(14)T.

Usage Guidelines The **bgp nexthop trigger enable** command is used to enable the BGP Support for Next-Hop Address Tracking feature. BGP next-hop address tracking is event driven. BGP prefixes are automatically tracked as peering sessions are established. Next-hop changes are rapidly reported to BGP as they are updated in the RIB. This optimization improves overall BGP convergence by reducing the response time to next-hop changes for routes installed in the RIB. When a bestpath calculation is run in between BGP scanner cycles, only the changes are processed and tracked.



Note

This feature improves BGP response time significantly. However, unstable IGP peers can introduce instability to the BGP session. We recommend that you aggressively dampen unstable IGP peering sessions to mitigate the possible impact to BGP.



Note

This feature is not supported under the IPv6 address family.

Examples In the following example, next-hop address tracking is disabled under the IPv4 address family session:

```
Router(config)# router bgp 50000
Router(config-router)# address-family ipv4 unicast
Router(config-router-af)# no bgp nexthop trigger enable
Router(config-router-af)# end
```

Related Commands	Command	Description
	bgp nexthop trigger delay	Configures the delay interval between routing table walks for BGP next-hop address tracking.
	neighbor fall-over	Configures BGP to use fast session deactivation for the specified peering session.

bgp nexthop trigger delay

To adjust the delay interval between routing table walks for Border Gateway Protocol (BGP) next-hop address tracking, use the **bgp nexthop trigger delay** command in address family or router configuration mode. To set the delay interval between routing table walks to the default value, use the **no** form of this command.

bgp nexthop trigger delay *seconds*

no bgp nexthop trigger delay

Syntax Description	<i>seconds</i>	The number of seconds for the delay. The range of values that can be entered is from 1 to 100 seconds.
---------------------------	----------------	--

Defaults The routing table is walked every 5 seconds.

Command Modes Address family configuration
Router Configuration

Command History	Release	Modification
	12.0(29)S	This command was introduced.
	12.0(31)S	The default delay interval was changed from 1 to 5 seconds.
	12.3(14)T	This command was integrated into Cisco IOS Release 12.3(14)T.

Usage Guidelines The **bgp nexthop trigger delay** command is used to adjust the delay interval between routing table walks for BGP next-hop address tracking. You can increase the performance of this feature by tuning the delay interval between full routing table walks to match the tuning parameters for the Interior Gateway protocol (IGP). The default delay interval is 5 seconds. This value is optimal for a fast-tuned IGP. In the case of an IGP that converges more slowly, you can change the delay interval to 20 seconds or more, depending on the IGP convergence time.



Note

This feature improves BGP response time significantly. However, unstable IGP peers can introduce instability to the BGP session. We recommend that you aggressively dampen unstable IGP peering sessions to mitigate the possible impact to BGP.

Examples In the following example, the delay interval is set to occur every 20 seconds under an IPv4 address family session:

```
Router(config)# router bgp 50000
Router(config-router)# address-family ipv4 unicast
Router(config-router-af)# bgp nexthop trigger delay 20
Router(config-router-af)# end
```

Related Commands	Command	Description
	bgp nexthop trigger enable	Enables or disables BGP next-hop address tracking
	neighbor fall-over	Configures BGP to use fast session deactivation for the specified peering session.

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.