



# BGP 4 MIB Support for Per-Peer Received Routes

This document describes BGP 4 MIB support for per-peer received routes. This feature introduces a table in the CISCO-BGP4-MIB that provides the capability to query (by using Simple Network Management Protocol [SNMP] commands) for routes that are learned from individual Border Gateway Protocol (BGP) peers.

- [Finding Feature Information, page 1](#)
- [Restrictions on BGP 4 MIB Support for Per-Peer Received Routes, page 1](#)
- [Information About BGP 4 MIB Support for Per-Peer Received Routes, page 2](#)
- [Additional References for BGP Restart Neighbor Session After Max-Prefix Limit Reached, page 6](#)
- [Feature Information for BGP 4 MIB Support for Per-Peer Received Routes, page 7](#)
- [Glossary, page 7](#)

## Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table at the end of this module.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to [www.cisco.com/go/cfn](http://www.cisco.com/go/cfn). An account on Cisco.com is not required.

## Restrictions on BGP 4 MIB Support for Per-Peer Received Routes

BGP 4 MIB Support for per-Peer Received Routes supports only routes that are contained in IPv4 AFIs and unicast SAFIs in the local BGP RIB table. The BGP 4 MIB Support for per-Peer Received Routes enhancement is supported only by BGP Version 4.

# Information About BGP 4 MIB Support for Per-Peer Received Routes

## Overview of BGP 4 MIB Support for Per-Peer Received Routes

The BGP 4 MIB support for per-peer received routes feature introduces a table in the CISCO-BGP4-MIB that provides the capability to query (by using SNMP commands) for routes that are learned from individual BGP peers.

Before this new MIB table was introduced, a network operator could obtain the routes learned by a local BGP-speaking router by querying the local BGP speaker with an SNMP command (for example, the `snmpwalk` command). The network operator used the SNMP command to query the `bgp4PathAttrTable` of the CISCO-BGP4-MIB. The routes that were returned from a `bgp4PathAttrTable` query were indexed in the following order:

- Prefix
- Prefix length
- Peer address

Because the `bgp4PathAttrTable` indexes the prefixes first, obtaining routes learned from individual BGP peers will require the network operator to "walk through" the complete `bgp4PathAttrTable` and filter out routes from the interested peer. A BGP Routing Information Base (RIB) could contain 10,000 or more routes, which makes a manual "walk" operation impossible and automated walk operations very inefficient.

BGP 4 MIB Support for per-Peer Received Routes introduces a Cisco-specific enterprise extension to the CISCO-BGP4-MIB that defines a new table called the `cbgpRouterTable`. The `cbgpRouterTable` provides the same information as the `bgp4PathAttrTable` with the following two differences:

- Routes are indexed in the following order:
  - Peer address
  - Prefix
  - Prefix length

The search criteria for SNMP queries of local routes are improved because peer addresses are indexed before prefixes. A search for routes that are learned from individual peers is improved with this enhancement because peer addresses are indexed before prefixes. A network operator will no longer need to search through potentially thousands of routes to obtain the learned routes of a local BGP RIB table.

- Support is added for multiprotocol BGP, Address Family Identifier (AFI), and Subsequent Address Family Identifier (SAFI) information. This information is added in the form of indexes to the `cbgpRouterTable`. The CISCO-BGP4-MIB can be queried for any combination of AFIs and SAFIs that are supported by the local BGP speaker.

**Note**

The MIB will be populated only if the router is configured to run a BGP process. The present implementation of BGP 4 MIB Support for Per-Peer Received Routes will show only routes contained in IPv4 AFI and unicast SAFI BGP local RIB tables. Support for showing routes contained in other local RIB tables will be added in the future.

## BGP 4 Per-Peer Received Routes Table Elements and Objects

The following sections describe new table elements, AFI and SAFI tables and objects, and network address prefixes in the Network Layer Reachability Information (NLRI) fields that have been introduced by the BGP 4 MIB Support for Per-Peer Received Routes enhancement.

### MIB Tables and Objects

The table below describes the MIB indexes of the `cbgpRouterTable`.

For a complete description of the MIB, see the CISCO-BGP4-MIB file CISCO-BGP4-MIB.my, available through Cisco.com at the following URL:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

**Table 1: MIB Indexes of the `cbgpRouterTable`**

MIB Indexes	Description
<code>cbgpRouteAfi</code>	Represents the AFI of the network layer protocol that is associated with the route.
<code>cbgpRouteSafi</code>	Represents the SAFI of the route. It gives additional information about the type of the route. The AFI and SAFI are used together to determine which local RIB (Loc-RIB) contains a particular route.
<code>cbgpRoutePeerType</code>	Represents the type of network layer address that is stored in the <code>cbgpRoutePeer</code> object.
<code>cbgpRoutePeer</code>	Represents the network layer address of the peer from which the route information has been learned.
<code>cbgpRouteAddrPrefix</code>	Represents the network address prefix that is carried in a BGP update message.  See the table below for information about the types of network layer addresses that can be stored in specific types of AFI and SAFI objects.

MIB Indexes	Description
cbgpRouteAddrPrefixLen	Represents the length in bits of the network address prefix in the NLRI field.  See the table below for a description of the 13 possible entries.

## AFIs and SAFIs

The table below lists the AFI and SAFI values that can be assigned to or held by the `cbgpRouteAfi` and `cbgpRouteSafi` indexes, respectively. The table below also displays the network address prefix type that can be held by specific combinations of AFIs and SAFIs. The type of network address prefix that can be carried in a BGP update message depends on the combination of AFIs and SAFIs.

**Table 2: AFIs and SAFIs**

AFI	SAFI	Type
ipv4(1)	unicast(1)	IPv4 address
ipv4(1)	multicast(2)	IPv4 address
ipv4(1)	vpn(128)	VPN-IPv4 address
ipv6(2)	unicast(1)	IPv6 address



**Note**

A VPN-IPv4 address is a 12-byte quantity that begins with an 8-byte Route Distinguisher (RD) and ends with a 4-byte IPv4 address. Any bits beyond the length specified by `cbgpRouteAddrPrefixLen` are represented as zeros.

## Network Address Prefix Descriptions for the NLRI Field

The table below describes the length in bits of the network address prefix in the NLRI field of the `cbgpRouteTable`. Each entry in the table provides information about the route that is selected by any of the six indexes in the table below.

**Table 3: Network Address Prefix Descriptions for the NLRI Field**

Table or Object (or Index)	Description
<code>cbgpRouteOrigin</code>	The ultimate origin of the route information.
<code>cbgpRouteASPathSegment</code>	The sequence of autonomous system path segments.

Table or Object (or Index)	Description
cbgpRouteNextHop	The network layer address of the autonomous system border router that traffic should pass through to get to the destination network.
cbgpRouteMedPresent	Indicates that the MULTI_EXIT_DISC attribute for the route is either present or absent.
cbgpRouteMultiExitDisc	Metric that is used to discriminate between multiple exit points to an adjacent autonomous system. The value of this object is irrelevant if the value of the cbgpRouteMedPresent object is "false(2)."
cbgpRouteLocalPrefPresent	Indicates that the LOCAL_PREF attribute for the route is either present or absent.
cbgpRouteLocalPref	Determines the degree of preference for an advertised route by an originating BGP speaker. The value of this object is irrelevant if the value of the cbgpRouteLocalPrefPresent object is "false(2)."
cbgpRouteAtomicAggregate	Determines if the system has selected a less specific route without selecting a more specific route.
cbgpRouteAggregatorAS	The autonomous system number of the last BGP speaker that performed route aggregation. A value of 0 indicates the absence of this attribute.
cbgpRouteAggregatorAddrType	Represents the type of network layer address that is stored in the cbgpRouteAggregatorAddr object.
cbgpRouteAggregatorAddr	The network layer address of the last BGP 4 speaker that performed route aggregation. A value of all zeros indicates the absence of this attribute.
cbgpRouteBest	An indication of whether this route was chosen as the best BGP 4 route.
cbgpRouteUnknownAttr	One or more path attributes not understood by the local BGP speaker. A size of 0 indicates that this attribute is absent.

## Benefits of BGP 4 MIB Support for Per-Peer Received Routes

- Improved SNMP Query Capabilities--The search criteria for SNMP queries for routes that are advertised by individual peers are improved because the peer address is indexed before the prefix. A network

operator will no longer need to search through potentially thousands of routes to obtain the learned routes of a local BGP RIB table.

- Improved AFI and SAFI Support--Support is added for multiprotocol BGP. AFI and SAFI are added as indexes to the table. The CISCO-BGP4-MIB can be queried for any combination of AFIs and SAFIs that are supported by the local BGP speaker.

## Additional References for BGP Restart Neighbor Session After Max-Prefix Limit Reached

### Related Documents

Related Topic	Document Title
Cisco IOS commands	<a href="#">Cisco IOS Master Command List, All Releases</a>
BGP commands	<a href="#">Cisco IOS IP Routing: BGP Command Reference</a>

### Standards and RFCs

Standard/RFC	Title
RFC 2918	<i>Route Refresh Capability for BGP-4</i>
RFC 4486	<i>Subcodes for BGP Cease Notification Message</i>

### Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	<a href="http://www.cisco.com/cisco/web/support/index.html">http://www.cisco.com/cisco/web/support/index.html</a>

# Feature Information for BGP 4 MIB Support for Per-Peer Received Routes

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to [www.cisco.com/go/cfn](http://www.cisco.com/go/cfn). An account on Cisco.com is not required.

**Table 4: Feature Information for BGP 4 MIB Support for Per-Peer Received Routes**

Feature Name	Releases	Feature Information
BGP 4 MIB support for per-peer received routes	Cisco IOS XE Release 2.1	This feature was introduced on the Cisco ASR 1000 Series Aggregation Services Routers.
BGP received routes MIB	Cisco IOS XE Release 2.1	This feature was introduced on the Cisco ASR 1000 Series Routers.

## Glossary

**AFI**--Address Family Identifier. Carries the identity of the network layer protocol that is associated with the network address.

**BGP**--Border Gateway Protocol. An interdomain routing protocol that exchanges reachability information with other BGP systems. It is defined by RFC 1163, A Border Gateway Protocol (BGP). The current implementation of BGP is BGP Version 4 (BGP4). BGP4 is the predominant interdomain routing protocol that is used on the Internet. It supports CIDR and uses route aggregation mechanisms to reduce the size of routing tables.

**MBGP**--multiprotocol BGP. An enhanced version of BGP that carries routing information for multiple network layer protocols and IP multicast routes. It is defined in RFC 2858, Multiprotocol Extensions for BGP-4.

**MIB**--Management Information Base. A group of managed objects that are contained within a virtual information store or database. MIB objects are stored so that values can be assigned to object identifiers and to assist managed agents by defining which MIB objects should be implemented. The value of a MIB object can be changed or retrieved using SNMP or CMIP commands, usually through a GUI network management system. MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches.

**NLRI**--Network Layer Reachability Information. Carries route attributes that describe a route and how to connect to a destination. This information is carried in BGP update messages. A BGP update message can carry one or more NLRI prefixes.

**RIB**--Routing Information Base (RIB). A central repository of routes that contains Layer 3 reachability information and destination IP addresses or prefixes. The RIB is also known as the routing table.

**SAFI**--Subsequent Address Family Identifier. Provides additional information about the type of the Network Layer Reachability Information that is carried in the attribute.

**SNMP**--Simple Network Management Protocol. A network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices and to manage configurations, statistics collection, performance, and security.

**snmpwalk** --The **snmpwalk** command is an SNMP application that is used to communicate with a network entity MIB using SNMP.

**VPN**--Virtual Private Network. Enables IP traffic to travel securely over a public TCP/IP network by encrypting all traffic from one network to another. A VPN uses a tunnel to encrypt all information at the IP level.