



MPLS Virtual Private Network Enhancements

This document describes the Multiprotocol Label Switching (MPLS) virtual private network (VPN) enhancements available in Cisco IOS Release 12.0(7)T. It also includes information about the new Border Gateway Protocol (BGP) commands available for enhanced VPN traffic management.

This document includes the following sections:

- [Feature Overview, page 1](#)
- [Supported Platforms, page 4](#)
- [Supported Standards, MIBs, and RFCs, page 4](#)
- [Configuration Tasks, page 4](#)
- [Configuration Examples, page 6](#)
- [Command Reference, page 6](#)

Feature Overview

The MPLS VPN enhancements provide increased BGP functionality enabling you to manage and route traffic within a VPN. With these MPLS VPN enhancements you can

- Configure BGP hub and spoke connections
- Configure faster convergence for BGP VPN routing/forwarding instance (VRF) routes
- Limit the number of BGP VRF routes
- Identify BGP site-of-origin routers
- Distribute BGP open shortest path first (OSPF) routing information

[Table 1](#) lists the MPLS VPN enhancements and the associated BGP commands.



Table 1 *MPLS VPN Enhancements in Release 12.0(7)T*

Name of Feature	Command	This command enables you to...
Configuring Faster Convergence for BGP VPN Routing/Forwarding Instance (VRF) Routes	bgp scan-time import	Configure scanning intervals of BGP routers to decrease import processing time of routing information.
Limiting VRF Routes	maximum routes	Limit the number of routes in a VRF to prevent a PE router from importing too many routes.
Configuring Border Gateway Protocol (BGP) Hub and Spoke Connections	neighbor allowas-in	Configure provider edge (PE) routers to allow customer edge (CE) routers to readvertise all prefixes that contain duplicate autonomous system numbers (ASNs) to neighboring PE routers.
Reusing ASNs in an MPLS VPN Environment	neighbor as-override	Configure a PE router to reuse the same ASN on all sites within an MPLS VPN by overriding private ASNs.

MPLS VPN Overview

Using MPLS VPNs in a Cisco IOS network provide the capability to deploy and administer scalable Layer 3 VPN backbone services including applications, data hosting network commerce, and telephony services, to business customers. A VPN is a secure IP-based network that shares resources on one or more physical networks. A VPN contains geographically dispersed sites that can communicate securely over a shared backbone.

A one-to-one relationship does not necessarily exist between customer sites and VPNs; a given site can be a member of multiple VPNs. However, a site can associate with only one VRF. Each VPN is associated with one or more VPN routing/forwarding instances (VRFs). A VRF includes routing and forwarding tables and rules that define the VPN membership of customer devices attached to CE routers. A VRF consists of the following:

- IP routing table
- Cisco Express Forwarding (CEF) table
- Set of interfaces that use the CEF forwarding table
- Set of rules and routing protocol parameters to control the information in the routing tables

VPN routing information is stored in the IP routing table and the CEF table for each VRF. A separate set of routing and CEF tables is maintained for each VRF. These tables prevent information from being forwarded outside a VPN and also prevent packets that are outside a VPN from being forwarded to a router within the VPN.

Distribution of VPN Routing Information

The distribution of VPN routing information is controlled through the use of VPN route target communities, implemented by border gateway protocol (BGP) extended communities. Distribution of VPN routing information works as follows:

- When a VPN route learned from a CE router is injected into BGP, a list of VPN route target extended community attributes is associated with it. Typically the list of route target community values is set from an export list of route targets associated with the VRF from which the route was learned.
- An import list of route target extended communities is associated with each VRF. The import list defines route target extended community attributes that a route must have in order for the route to be imported into the VRF. For example, if the import list for a particular VRF includes route target communities A, B, and C, then any VPN route that carries any of those route target extended communities—A, B, *or* C—is imported into the VRF.

BGP Operation

A service provider edge (PE) router can learn an IP prefix from a customer edge (CE) router by static configuration, through a BGP session with the CE router, or through the routing information protocol (RIP) exchange with the CE router. The IP prefix is a member of the IPv4 address family. After it learns the IP prefix, the PE converts it into a VPN-IPv4 prefix by combining it with an 8-byte route distinguisher (RD). The generated prefix is a member of the VPN-IPv4 address family. It uniquely identifies the customer address, even if the customer site is using globally nonunique (unregistered private) IP addresses.

The route distinguisher used to generate the VPN-IPv4 prefix is specified by a configuration command associated with the VRF on the PE router.

BGP distributes reachability information for VPN-IPv4 prefixes for each VPN. BGP communication takes place at two levels: within IP domains, known as autonomous systems (interior BGP or IBGP) and between autonomous systems (external BGP or EBGP). PE-PE or PE-RR (route reflector) sessions are IBGP sessions, and PE-CE sessions are EBGP sessions.

BGP propagates reachability information for VPN-IPv4 prefixes among PE routers by means of the BGP multiprotocol extensions (see RFC 2283, *Multiprotocol Extensions for BGP-4*), which define support for address families other than IPv4. It does this in a way that ensures the routes for a given VPN are learned only by other members of that VPN, enabling members of the VPN to communicate with each other.

Benefits

Configuring BGP Hub and Spoke Connections—Configuring PE routers in a hub and spoke configuration allows a CE router to readvertise all prefixes containing duplicate autonomous system numbers (ASNs) to neighboring PE routers. Using duplicate ASNs in a hub and spoke configuration provides faster convergence of routing information within geographically dispersed locations.

Configuring Faster Convergence for BGP VRF Routes—Configuring scanning intervals of BGP routers decreases import processing time of VPNv4 routing information, thereby providing faster convergence of routing information. Routing tables are updated with routing information about VPNv4 routes learned from PE routers or route reflectors.

Limiting VPN VRFs—Limiting the number of routes in a VRF prevents a PE router from importing too many routes, thus diminishing the router's performance. This enhancement can also be used to enforce the maximum number of members that can join a VPN from a particular site. A threshold is set in the VRF routing table to limit the number of VRF routes imported.

Reuse ASNs in an MPLS VPN Environment—Configuring a PE router to reuse an existing ASN allows customers to configure BGP routes with the same ASNs in multiple geographically dispersed sites, providing better scalability between sites.

Distributing BGP OSPF Routing Information—Setting a separate router ID for each interface or subinterface on a PE router attached to multiple CE routers within a VPN provides increased flexibility through OSPF when routers exchange routing information between sites.

Related Documents

For more information about the MPLS VPN functionality including BGP distribution of routing information, see the *MPLS Virtual Private Network Feature Module*, Cisco IOS Release 12.0(5)T on CCO at

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/index.htm>.

Supported Platforms

The MPLS VPN enhancements support the following platforms:

- Cisco 3640 series
- Cisco 4500 series
- Cisco 7200 series
- Cisco 7500 series

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.

For descriptions of supported MIBs and how to use MIBs, see the Cisco MIB web site on CCO at <http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>.

RFCs

No new or modified RFCs are supported by this feature.

Configuration Tasks

Perform the following tasks to configure and verify VPNs:

- [Defining VPNs](#)
- [Configuring BGP Routing Sessions](#)
- [Verifying VPN Operation](#)

Defining VPNs

To define VPNs, perform the following steps on the PE router:

	Command	Purpose
Step 1	Router(config)# ip vrf <i>vrf-name</i>	Enter VRF configuration mode and assign a VRF name.
Step 2	Router(config-vrf)# rd <i>route-distinguisher</i>	Create routing and forwarding tables.
Step 3	Router(config-vrf)# route-target { import export both } <i>route-target-ext-community</i>	Create a list of import and/or export route target communities for the specified VRF.
Step 4	Router(config-vrf)# import map <i>route-map</i>	(Optional) Associate the specified route map with the VRF.
Step 5	Router(config-vrf)# ip vrf forwarding <i>vrf-name</i>	Associate a VRF with an interface or subinterface.

Configuring BGP Routing Sessions

To configure BGP routing sessions in a provider network, perform the following steps on the PE routers:

	Command	Purpose
Step 1	Router(config)# router bgp <i>autonomous-system</i>	Configure the BGP routing process with the autonomous system number passed along to other BGP routers.
Step 2	Router(config-router)# neighbor { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>number</i>	Specify a neighbor's IP address or BGP peer group identifying it to the local autonomous system.
Step 3	Router(config-router)# neighbor <i>ip-address</i> activate	Activate the advertisement of the IPv4 address family.

Verifying VPN Operation

To verify VPN operation by displaying routing information on the PE routers, you may issue any of the following show commands in any order:

Command	Purpose
Router# Router# show ip vrf	Display the set of defined VRFs and interfaces.
Router# Router# show ip vrf [{ brief detail interfaces }] <i>vrf-name</i>	Displays information about defined VRFs and associated interfaces.
Router# show ip bgp vpnv4 all [tags]	Displays information about all BGP VPN-IPv4 prefixes.
Router# show ip interface <i>interface-number</i>	Displays the VRF table associated with an interface.

Configuration Examples

For MPLS VPN configuration examples, see the *MPLS Virtual Private Network Feature Module*, Cisco IOS Release 12.0(5)T.

Command Reference

This section documents the new VPN commands for Cisco IOS Release 12.0(7). All other commands used with MPLS VPNs are documented in Cisco IOS Release 12.0(5)T.

- [bgp scan-time](#)
- [maximum routes](#)
- [neighbor allowas-in](#)
- [neighbor as-override](#)

bgp scan-time

To configure scanning intervals of BGP routers for next hop validation or to decrease import processing time of Virtual Private Network version 4 (VPNv4) routing information, use the **bgp scan-time** command in address family or router configuration mode. To return the scanning interval of a router to its default scanning interval of 15 seconds, use the **no** form of this command.

bgp scan-time [**import**] *scanner-interval*

no bgp scan-time [**import**] *scanner-interval*

Syntax Description		
<code>import</code>	(Optional) Configures import processing of VPNv4 unicast routing information from BGP routers into routing tables.	
<code>scanner-interval</code>	Specifies the scanning interval of BGP routing information. Valid values used for selecting the desired scanning interval are from 5 to 60 seconds. The default is 15 seconds.	

Defaults The default scanning interval is 15 seconds.

Command Modes Address family configuration
Router configuration

Command History	Release	Modification
	12.07(T)	This command was introduced.

Usage Guidelines The **import** keyword is supported in address family VPNv4 unicast mode only. Entering the **no** form of this command does not disable scanning, but removes it from the output of the **show running-config** command.

Examples In the following router configuration example, the scanning interval for next hop validation of IPv4 unicast routes for BGP routing tables is set to 20 seconds:

```
router bgp 100
 no synchronization
 bgp scan-time 20
```

In the following address family configuration example, the scanning interval for next hop validation of address family VPNv4 unicast routes for BGP routing tables is set to 45 seconds:

```
router bgp 150
 address-family vpn4 unicast
 bgp scan-time 45
```

In the following address family configuration example, the scanning interval for importing address family VPNv4 routes into IP routing tables is set to 30 seconds:

```
router bgp 150
 address-family vpnv4 unicast
  bgp scan-time import 30
```

Related Commands

Command	Description
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.

maximum routes

To limit the maximum number of routes in a Virtual Private Network routing/forwarding instance (VRF) to prevent a provider edge (PE) router from importing too many routes, use the **maximum routes** command in VRF configuration mode. To remove the limit on the maximum number of routes allowed, use the **no** form of this command.

maximum routes *limit* {*warn-threshold* | **warn-only**}

no maximum routes

Syntax Description		
	<i>limit</i>	Specifies the maximum number of routes allowed in a VRF. You may select from 1 to 4,294,967,295 routes to be allowed in a VRF.
	<i>warn threshold</i>	Generates a warning when the number of routes set by the <i>warn-threshold</i> argument is reached and rejects routes that exceed the maximum number set in the <i>limit</i> argument. The warning threshold is a percentage of the maximum number of routes specified in the <i>limit</i> argument, from 1 to 100.
	warn-only	Issues a SYSLOG error message when the maximum number of routes allowed for a VRF exceeds the limit threshold. However, additional routes are still allowed.

Defaults No default behavior or values.

Command Modes VRF configuration mode

Command History	Release	Modification
	12.0(7)T	This command was introduced.

Usage Guidelines You can use the **maximum routes** command to monitor and limit the number of routes in a VRF on a PE router.

To limit the number of routes allowed in the VRF, use the **maximum routes** *limit* command with the *warn-threshold* argument. The *warn-threshold* argument generates a warning and does not allow the addition of routes to the VRF when the maximum number set by the *limit* argument is reached. The software generates a warning message everytime a route is added to a VRF when the VRF route count is above the warning threshold. The software also generates a route rejection notification when the maximum threshold is reached and everytime a route is rejected after the limit is reached.

To set a number of routes at which you receive a notification, but which does not limit the number of routes that can be imported into the VRF, use the **maximum routes** *limit* command with the **warn-only** keyword.

To use the **maximum routes** command, you must enter the VRF configuration submode.

Examples

The following example shows how to set a limit threshold of VRF routes to 1000. When the number of routes for the VRF reaches 1000, the router issues a SYSLOG error message, but continues to accept new VRF routes.

```
Router(config)# ip vrf vrf1

Router(config-vrf)# rd 100:1

Router(config-vrf)# route-target import 100:1

Router(config-vrf)# maximum routes 1000 warn-only
```

The following example shows how to set the maximum number of VRF routes allowed to 1000 and set the warning threshold at 80 percent of the maximum. When the number of routes for the VRF reaches 800, the router issues a warning message. When the number of routes for the VRF reaches 1000, the router issues a SYSLOG error message and rejects any new routes.

```
Router(config)# ip vrf vrf2

Router(config-vrf)# rd 200:1

Router(config-vrf)# route-target import 200:1

Router(config-vrf)# maximum routes 1000 80
```

Related Commands

Command	Description
rd	Creates VRF routing and forwarding tables and specifies the default route distinguisher for a VPN.
route-target	Configures a VRF route target community for importing and exporting extended community attributes.
import map	Configures an import route map for a specified VRF for more control over routes imported into the VRF.

neighbor allowas-in

To configure PE routers to allow readvertisement of all prefixes containing duplicate ASNs, use the **neighbor allowas-in** router configuration command. To disable the readvertisement of a PE router's ASN, use the **no** form of this command.

neighbor allowas-in *number*

no neighbor allowas-in *number*

Syntax Description	<i>number</i>	Specifies the number of times to allow the advertisement of a PE router's ASN. Valid values are from 1 to 10 times.
---------------------------	---------------	---

Defaults	No default behavior or values.
-----------------	--------------------------------

Command Modes	Router configuration
----------------------	----------------------

Command History	Release	Modification
	12.0(7)T	This command was introduced.

Usage Guidelines In a hub and spoke configuration, a PE router readvertises all prefixes containing duplicate autonomous system numbers. Use the **neighbor allowas-in** command to configure two VRFs on each PE router to receive and readvertise prefixes:

1. One VRF receives prefixes with ASNs from all PE routers and then advertises them to neighboring PE routers.
2. The other VRF receives prefixes with ASNs from the CE router and readvertises them to all PE routers in the hub and spoke configuration.

You control the number of times an ASN is advertised by specifying a number from 1 to 10.

Examples In the following example, the PE router with ASN 100 is configured to allow prefixes from the VRF address family VPN IPv4 vrf1. The neighboring PE router with the IP address 192.168.255.255 is set to be readvertised to other PE routers with the same ASN 6 times:

```
router bgp 100
 address-family ipv4 vrf vrf1
 neighbor 192.168.255.255 allowas-in 6
```

Related Commands	Command	Description
	address-family	Enters the address family submode used to configure routing protocols including BGP, OSPF, RIP, and static routing.

neighbor as-override

To configure a PE router to override a site's ASN with a provider's ASN, use the **neighbor as-override** router configuration command. To remove VPN IPv4 prefixes from a specified router, use the **no** form of this command.

neighbor ip-address as-override

no neighbor ip-address as-override

Syntax Description	<i>ip-address</i>	Specifies the router's IP address to override with the ASN provided.
---------------------------	-------------------	--

Defaults	No default behavior or values.
-----------------	--------------------------------

Command Modes	Router configuration
----------------------	----------------------

Command History	Release	Modification
	12.0(7)T	This command was introduced.

Usage Guidelines	This command is used in conjunction with the site-of-origin feature, identifying the site where a route originated from, and preventing routing loops between routers within a VPN.
-------------------------	---

Examples	In the following example, the router's ASN of 100 overrides the neighboring routers IP address 192.168.255.255.
-----------------	---

```
router bgp 100
 neighbor 192.168.255.255 remote-as 100
 neighbor 192.168.255.255 update-source loopback0
 address-family ipv4 vrf vpn1
 neighbor 192.168.255.255 activate
 neighbor 192.168.255.255 as-override
```

Related Commands	Command	Description
	neighbor activate	Enables the exchange of information with a BGP neighboring router.
	neighbor remote-as	Allows a neighboring router's IP address to be included in the BGP routing table.
	neighbor update-source	Allows internal BGP sessions to use any operational interface for TCP/IP connections.
	route-map	Redistributes routes from one routing protocol to another.

