Route Leaking in MPLS/VPN Networks

Document ID: 47807

Contents

Introduction
Prerequisites
Requirements
Components Used
Conventions
Configure
Route Leaking from a Global Routing Table into a VRF and Route Leaking from a VRF into a Global Routing Table
Route Leaking Between Different VRFs
Troubleshoot
Related Information

Introduction

This document provides sample configurations for route leaking in an MPLS/VPN environment.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

For more information on document conventions, refer to the Cisco Technical Tips Conventions .

Configure

This sections contains these two configuration examples:

- Route leaking from a global routing table into a VPN routing/forwarding instance (VRF) and route leaking from a VRF into a global routing table
- Route leaking between different VRFs

Note: To find additional information about the commands in this document, use the Command Lookup Tool (registered customers only) .

Route Leaking from a Global Routing Table into a VRF and Route Leaking from a VRF into a Global Routing Table

This configuration describes route leaking from a global routing table into a VRF and route leaking from a VRF into a global routing table.

Network Diagram

This configuration uses this network setup:



Configuration

In this example, a Network Management System (NMS) station located in a VRF is accessed from the global routing table. The provider edge (PE) routers and provider (P) routers have to export the netflow information to an NMS station (10.0.2.2) in a VRF. 10.0.2.2 is reachable via a VRF interface on PE−4.

To access 10.0.2.0/30 from the global table, a static route to 10.0.2.0/30 that points out of the VRF interface is introduced on the PE−4. This static route is then redistributed via Interior Gateway Protocol (IGP) to all PE and P routers. This ensures that all PE and P routers can reach 10.0.2.0/30 via PE−4.

A static VRF route is also added. The static VRF route points to the subnet in the global network that sends the traffic to this NMS station. Without this addition, the PE−4 drops traffic, from the NMS station, that is received on the VRF interface; and the PE−4 sends the `ICMP: host unreachable rcv` message to the NMS station.

This section uses this configuration:

- PE−4

PE−4

```
!
ip cef
!
ip vrf vpn2
rd 200:1
route-target export 200:1
route-target import 200:1
!
interface Serial1/0
ip address 10.1.2.5 255.255.255.252
no ip directed-broadcast
!
interface Serial2/0
ip vrf forwarding vpn2
ip address 10.0.2.1 255.255.255.0
no ip directed-broadcast
!
ip classless
ip route 10.0.2.0 255.255.255.252 Serial2/0
ip route vrf vpn2 10.1.2.4 255.255.255.252 Serial1/0
```

```
!
```

The static routes can now be redistributed into any IGP to be announced network−wide. The same applies if the VRF interface is a LAN interface (for example, Ethernet). The exact configuration command for that is:

```
ip route 10.0.2.0 255.255.255.252 Ethernet2/0 10.0.2.2
```

Note: The IP address configured after the interface name is only used by Address Resolution Protocol (ARP), to know what address to resolve.

Note: For 4500 series switches, you must configure static ARP entries in the VRF tables for the respective next hop addresses.

Note: By default, Cisco IOS® software accepts static VRF routes as configured. This might compromise security because it might introduce route leaking between different VRFs. You can use the **no ip route static inter−vrf** command to prevent the installation of such static VRF routes. Refer to MPLS Virtual Private Networks (VPNs) for more information on the **no ip route static inter−vrf** command.

Verify

This section provides information to confirm that your configuration is working properly.

Certain **show** commands are supported by the Output Interpreter Tool (registered customers only) , which allows you to view an analysis of **show** command output.

- **show ip route 10.0.2.0** Displays a specified IP address routing entry.
- **show ip route vrf vpn2 10.1.2.4** Displays a specified IP address VRF routing entry.

```
PE−4# show ip route 10.0.2.0

Routing entry for 10.0.2.0/30
Known via "static", distance 1, metric 0 (connected)
Routing Descriptor Blocks:
* directly connected, via Serial2/0
Route metric is 0, traffic share count is 1

PE−4# show ip route vrf vpn2 10.1.2.4

Routing entry for 10.1.2.4/30
Known via "static", distance 1, metric 0 (connected)
Redistributing via bgp 1
Advertised by bgp 1
Routing Descriptor Blocks:
* directly connected, via Serial1/0
Route metric is 0, traffic share count is 1
```

Route Leaking Between Different VRFs

This configuration describes route leaking between different VRFs.

Network Diagram

This configuration uses this network diagram:



Configuration

You can not configure two static routes to advertise each prefix between the VRFs, because this method is not supported packets will not be routed by the router. To achieve route leaking between VRFs, you must use the import functionality of route−target and enable Border Gateway Protocol (BGP) on the router. No BGP neighbor is required.

This section uses this configuration:

- PE−4

PE−4

```
!
ip vrf vpn1
 rd 100:1
 route-target export 100:1
 route-target import 100:1
 route-target import 200:1
!
ip vrf vpn2
 rd 200:1
 route-target export 200:1
 route-target import 200:1
 route-target import 100:1
!
interface Serial1/0
 ip vrf forwarding vpn1
 ip address 10.1.2.5 255.255.255.252
 no ip directed-broadcast
!
interface Serial2/0
 ip vrf forwarding vpn2
 ip address 10.0.2.1 255.255.255.0
 no ip directed-broadcast
router bgp 1
!
address-family ipv4 vrf vpn2
 redistribute connected
 !
address-family ipv4 vrf vpn1
 redistribute connected
!
```

Verify

This section provides information to troubleshoot your configuration.

Certain **show** commands are supported by the Output Interpreter Tool (registered customers only) , which allows you to view an analysis of **show** command output.

- **show ip bgp vpnv4 all** Displays all VPNv4 prefixes that are learned via BGP.

```
PE−4# show ip bgp vpnv4 all

BGP table version is 13, local router ID is 7.0.0.4
Status codes: s suppressed, d damped, h history, * valid,
> best, i − internal, r RIB−failure, S Stale
Origin codes: i − IGP, e − EGP, ? − incomplete

Network Next Hop Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf vpn1)
*> 10.0.2.0/24 0.0.0.0 0 32768 ?
*> 10.1.2.4/30 0.0.0.0 0 32768 ?
Route Distinguisher: 200:1 (default for vrf vpn2)
*> 10.0.2.0/24 0.0.0.0 0 32768 ?
*> 10.1.2.4/30 0.0.0.0 0 32768 ?
```

Note: The other way of leaking routes between VRFs is to connect together two Ethernet interfaces on the PE−4 router and associate each Ethernet interface with one of the VRFs. You also must configure static ARP entries in the VRF tables for the respective next hop addresses. However, this is not a recommended solution for route leaking between VRFs; the previously described BGP technique is the recommended solution.

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Related Information

- **MPLS Support Page**
- **Technical Support and Documentation − Cisco Systems**

Contacts & Feedback | Help | Site Map
© 2014 − 2015 Cisco Systems, Inc. All rights reserved. Terms & Conditions | Privacy Statement | Cookie Policy | Trademarks of Cisco Systems, Inc.

Updated: Aug 10, 2005 Document ID: 47807