

Packet Flow in an MPLS VPN Environment

Document ID: 10474

Contents

Introduction

Before You Begin

Conventions

Prerequisites

Components Used

Network Diagram

The Packet Flow Process

Related Information

Introduction

This document illustrates the packet flow through a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) cloud. It also introduces the concept of having multiple labels inside a packet.

VPN, when used with MPLS, allows several sites to transparently interconnect through a service provider's network. One service provider network can support several different IP VPNs. Each of these appears to its users as a private network, separate from all other networks. Within a VPN, each site can send IP packets to any other site in the same VPN.

Each VPN is associated with one or more VPN routing or forwarding instances (VRFs). A VRF consists of an IP routing table, a derived Cisco express forwarding (CEF) table and a set of interfaces that use this forwarding table.

The router maintains a separate routing and CEF table for each VRF. This prevents information being sent outside the VPN and allows the same subnet to be used in several VPNs without causing duplicate IP address problems.

The router using Border Gateway Protocol (BGP) distributes the VPN routing information using the BGP extended communities.

For more information regarding the propagation of updates through a VPN, refer to these documents:

- [VPN Route Target Communities.](#)
- [BGP Distribution of VPN Routing Information.](#)
- [MPLS Forwarding.](#)
- [Link to a Sample Configuration.](#)

The MPLS VPN feature was introduced in Cisco IOS[®] Software Release 12.0(5)T.

Before You Begin

Conventions

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

Prerequisites

There are no specific prerequisites for this document.


Components Used

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

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Network Diagram

In order to understand how VPN MPLS works, let's take the following sample configuration:



In this configuration:

- Rapid and Pound are the Customer Edge (CE) devices not running MPLS. They are associated with the VPN VRF101. For simplicity, we are only using one VRF here.
- Farm and Medina are the Provider Edge devices (PEs).
- Miles and Yard are LightStream 1010 routers. They constitute the MPLS backbone.

The Packet Flow Process

The output below shows what happens when Rapid sends packets to Pound inside the VPN VRF101:

```
rapid#ping 11.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 11.5.5.5, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

rapid#show ip route 11.5.5.5
Routing entry for 11.5.5.4/30
  Known via "rip", distance 120, metric 1
  Redistributing via rip
  Last update from 150.150.0.1 on FastEthernet0/1, 00:00:16 ago
  Routing Descriptor Blocks:
    * 150.150.0.1, from 150.150.0.1, 00:00:16 ago, via FastEthernet0/1
      Route metric is 1, traffic share count is 1
```

Farm learns the address 11.5.5.5 from Medina through BGP advertisements:

```
Farm#show ip bgp vpnv4 vrf vrf101 11.5.5.5
  BGP routing table entry for 1:101:11.5.5.4/30, version 56
  Paths: (1 available, best #1, table vrf101)
  Not advertised to any peer
  Local
    125.2.2.2 (metric 4) from 125.2.2.2 (125.2.2.2)
    Origin incomplete, metric 1, localpref 100, valid, internal, best
    Extended Community: RT:1:101

Farm#show ip route vrf vrf101 11.5.5.5
Routing entry for 11.5.5.4/30
  Known via "bgp 1", distance 200, metric 1, type internal
  Redistributing via rip
  Advertised by rip metric 0
  Last update from 125.2.2.2 01:29:20 ago
  Routing Descriptor Blocks:
  * 125.2.2.2 (Default-IP-Routing-Table), from 125.2.2.2, 01:29:20 ago
    Route metric is 1, traffic share count is 1
    AS Hops 0
```

Note: 125.2.2.2 is a loopback on Medina and is used to create the BGP pairing with Farm.

In order to forward the packet destined for 11.5.5.5 to Medina, Farm uses two labels. To see this, look at the CEF and the VPN label forwarding table on Farm:

```
Farm#show tag forwarding
-table vrf vrf101 11.5.5.5 detail
  Local   Outgoing   Prefix           Bytes tag  Outgoing   Next Hop
  tag     tag or VC  or Tunnel Id     switched   interface
  None    2/91      11.5.5.4/30      0          AT4/0.1    point2point
          MAC/Encaps=4/12, MTU=4466, Tag Stack{2/91(vcd=69) 40}
          00458847 0004500000028000

Farm#show ip cef vrf vrf101 11.5.5.5
11.5.5.4/30, version 25, cached adjacency to ATM4/0.1
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with AT4/0.1, point2point, tags imposed: {2/91(vcd=69) 40}
  via 125.2.2.2, 0 dependencies, recursive
  next hop 10.0.0.14, ATM4/0.1 via 125.2.2.2/32
  valid cached adjacency
  tag rewrite with AT4/0.1, point2point, tags imposed: {2/91(vcd=69) 40}
```

Two labels are applied to the packets that are leaving Farm and destined for 11.5.5.5. These can be represented like this:

2/91	40	Packet
------	----	--------

The label 40 is added to the packet and this is then segmented into cells with 2/91 as the VPI/VCI values. This means that the label is also called 2/91.

Note: On receiving a frame with several labels, the receiving device checks only the first one.

The labels are assigned as follows:

- 2/91 is assigned by Yard and corresponds to the address 125.2.2.2. This address is used to create the BGP pairing with Farm. Refer to MPLS VPN over ATM: with BGP or RIP on the Customer Site for more information. The label is used in the MPLS core to send frames from Farm to 125.2.2.2 on Medina.
- 40 is assigned to 11.5.5.5 by Medina. When a PE (Medina in this case) learns an IP prefix from a CE (Pound), the PE assigns a specific label to this route. The label depends on which VPN VRF the route has learned. It advertises the route and label to the other PEs using BGP enhanced communities.

Let's have a look at Medina:

```
Medina#show tag forwarding
-table vrf vrf101 11.5.5.5 detail
  Local   Outgoing   Prefix           Bytes tag   Outgoing   Next Hop
  tag     tag or VC   or Tunnel Id     switched    interface
  40      Untagged    11.5.5.4/30[V]   570         Et1/1      11.3.3.2
          MAC/Encaps=0/0, MTU=1500, Tag Stack{}
          VPN route: vrf101
          Per-packet load-sharing
```

Now that we know where the labels come from, we can see what happens to the packets destined for 11.5.5.5. Farm sends the segmented packet over the VC 2/91. Yard receives this. To look at what Yard does with these cells, use the following command:

```
Yard#show tag atm
-tdp bindings 125.2.2.2 32
  Destination: 125.2.2.2/32
  Transit ATM0/1/1 2/91 Active -> ATM4/0/0 1/82 Active
```

On receiving these cells on the VC 2/91 (cells which are destined for 125.2.2.2, also known as Medina), Yard switches these cells to Miles using the outgoing VC 1/82.

Note: Yard has not checked or modified label 40.

The same thing happens on Miles, switching the cells to Medina on the VC 1/33:

```
Miles#show tag atm
-tdp bindings 125.2.2.2 32
  Destination: 125.2.2.2/32
  Transit ATM0/1/3 1/82 Active -> ATM0/1/1 1/33 Active
```

The packet that arrives at Medina can be represented like this:

1/33	40	Packet
-------------	-----------	---------------

On receiving the cells on the VC 1/33, Medina checks the label 1/33 and sees that this label is local to the router. In doing so, Medina sees that the packet is destined for one of its own addresses:

```
Medina#show tag
-switching atm-tdp bindings local-tag 1 33
  Destination: 125.2.2.2/32
  Tailend Router ATM2/0.66 1/33 Active, VCD=406
```

Medina therefore removes the first label (1/33) and sees that the packet has another label (40). It then checks what this label corresponds to and switches the packet accordingly:

```
Medina#show tag
```

-switching forwarding-table tags 40

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
40	Untagged	11.5.5.4/30[V]	570	Et1/1	11.3.3.2

In this case, Medina sees that the packet is destined for a site connected by an ordinary IP link. It discards the label and forwards the IP packet on the interface ethernet 1/1.

Related Information

- [Tools and Resources](#)
- [Technical Support & 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: Jun 05, 2005

Document ID: 10474
