

Inter-AS MPLS VPN Configuration with VPNv4 eBGP Sessions Between ASBRs

Document ID: 5811

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Configure

- Network Diagram
- Configurations
- Configuration Notes

Verifying and Understanding Inter-AS MPLS VPN

- PEsaro Router
- Pomerol Router
- Pulligny Router

The next-hop-self Command on ASBRs

Troubleshoot

Related Information

Introduction

This document shows a configuration example for a basic inter-autonomous system (inter-AS) Multiprotocol Label Switching (MPLS) VPN.

Prerequisites

Requirements

This document assumes a working knowledge of MPLS and MPLS VPN. Please see the Related Information section of this document for links to other MPLS VPN documents.

Components Used

The information in this document is based on the following software versions:

- Cisco IOS® Software Releases 12.2 and 12.2T

Note: Because this configuration uses a mixture of codes, Tag Distribution Protocol (TDP) is used rather than Label Distribution Protocol (LDP). In a pure LDP setup, Tag Forwarding Information Base (TFIB) would be replaced by Label Forwarding Information Base (LFIB), the **show tag forwarding** command would become **show mpls forwarding** command, and so forth.

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.

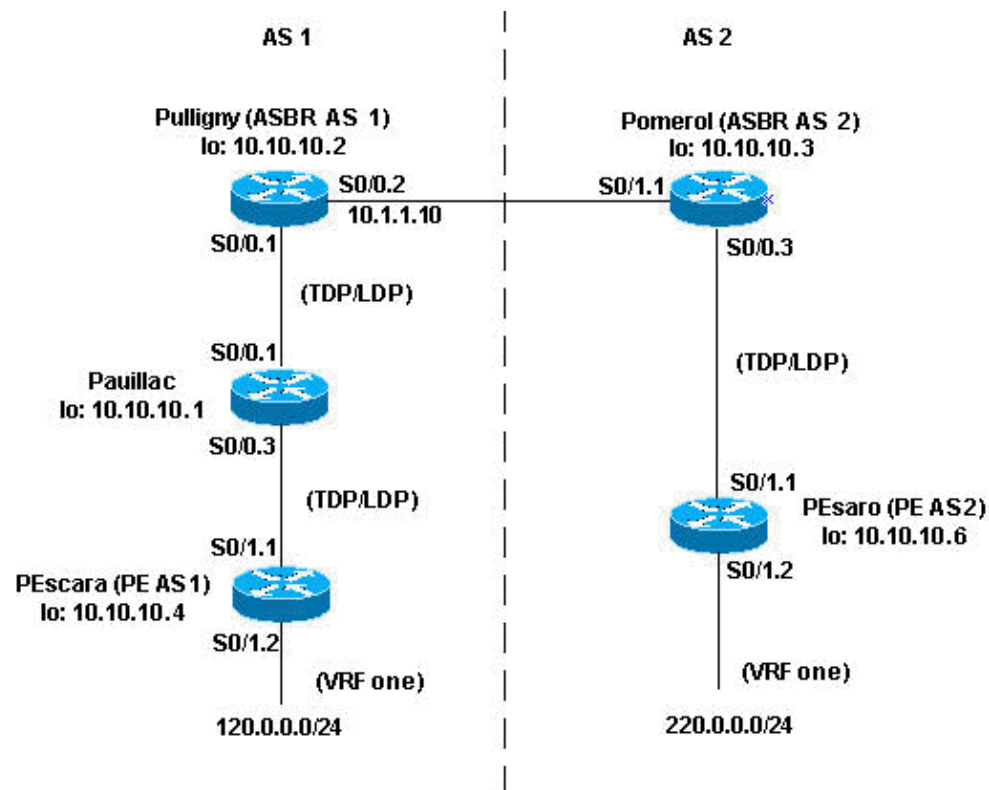
Configure

In this section, you are presented with the information to configure the features described in this document.

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

Network Diagram

This document uses the network setup shown in the diagram below.



Configurations

This document uses the configurations shown below for the routers in the network diagram above.

- Pulligny
- PEscara
- Pauillac
- Pomerol
- PEsaro

Pulligny
<pre>version 12.2 ! hostname Pulligny</pre>

```
!  
ip cef  
  
!--- Cisco Express Forwarding (CEF) must be enabled for MPLS.  
  
!  
interface Loopback0  
ip address 10.10.10.2 255.255.255.255  
ip router isis  
!  
interface Serial0/0  
no ip address  
encapsulation frame-relay  
!  
interface Serial0/0.1 point-to-point  
description to Paulliac  
ip address 10.1.1.2 255.255.255.252  
ip router isis  
tag-switching ip  
  
!--- TDP or LDP needs to be enabled inside each AS.  
  
frame-relay interface-dlci 913  
!  
interface Serial0/0.2 point-to-point  
description to Pomerol  
ip address 10.1.1.10 255.255.255.252  
frame-relay interface-dlci 912  
  
!--- No TDP, LDP, or routing protocols enabled on the inter-AS link.  
  
!  
router isis  
redistribute connected metric 20  
  
!--- To get the host route created by the Border Gateway Protocol (BGP)  
!--- into the Interior Gateway Protocol (IGP).  
  
net 47.0000.5555.5555.5555.00  
metric-style wide  
!  
router bgp 1  
no bgp default ipv4-unicast  
  
!--- BGP is not used for IPv4 unicast.  
  
no bgp default route-target filter  
  
!--- Needed to accept VPNv4 prefixes.  
  
neighbor 10.1.1.9 remote-as 2  
  
!--- External BGP (eBGP) to Pomerol.  
  
neighbor 10.10.10.4 remote-as 1  
neighbor 10.10.10.4 update-source Loopback0  
  
!--- Internal BGP (iBGP) to PEsacara.  
  
!  
address-family vpnv4  
neighbor 10.1.1.9 activate  
neighbor 10.1.1.9 send-community both  
neighbor 10.10.10.4 activate  
neighbor 10.10.10.4 send-community extended  
exit-address-family
```

```
!--- Address family to allow BGP to carry VPN-IPv4 protocols.

!
ip classless
!
end
```

PEscara

```
hostname PEscara
!
ip vrf one
  rd 1:1
  route-target export 1:1
  route-target import 1:1

!--- Create a virtual routing and forwarding (VRF) instance called one
!--- with a route distinguished value of 1:1. Routes with route target
!--- values of 1:1 will be imported into the VRF and exported out
!--- of the VRF.

ip cef

!--- CEF must be enabled for MPLS.

!
interface Loopback0
  ip address 10.10.10.4 255.255.255.255
  ip router isis
!
interface Serial0/1
  no ip address
  encapsulation frame-relay
  frame-relay lmi-type cisco
!
interface Serial0/1.1 point-to-point
  description to Pauillac
  ip address 10.1.1.14 255.255.255.252
  ip router isis
  tag-switching ip
  frame-relay interface-dlci 973
!
interface Serial0/1.2 point-to-point
  description to CEntrum
  ip vrf forwarding one

!--- Associates the interface with VRF one.

  ip address 120.0.0.1 255.255.255.0
  frame-relay interface-dlci 979
!
router isis
  net 47.0000.7777.7777.7777.00
  metric-style wide
!
router bgp 1
  no bgp default ipv4-unicast
  bgp log-neighbor-changes
  neighbor 10.10.10.2 remote-as 1
  neighbor 10.10.10.2 update-source Loopback0

!--- iBGP to Pulligny.

!
```

```

address-family ipv4 vrf one
redistribute connected
no auto-summary
no synchronization
exit-address-family

!--- Address family to allow BGP to carry IPv4 protocols for VRF one
!--- (each VRF that injects routes into BGP must be configured under
!--- the BGP process using its own address family).

!
address-family vpnv4
neighbor 10.10.10.2 activate
neighbor 10.10.10.2 send-community both
exit-address-family

!--- Address family to allow BGP to carry VPN-IPv4 protocols.

!
ip classless
!
end

```

Paulliac

```

!--- The Paulliac router is a Provider router (p) and its configuration
!--- is quite simple. It is configured for Intermediate System-to-
!--- Intermediate System (IS-IS) and MPLS/Tag switching.

!
hostname Paulliac
!
ip cef
!
interface Loopback0
ip address 10.10.10.1 255.255.255.255
ip router isis
!
interface Serial0/0
no ip address
encapsulation frame-relay
!
interface Serial0/0.1 point-to-point
description to Pulligny
ip address 10.1.1.1 255.255.255.252
ip router isis
tag-switching ip
frame-relay interface-dlci 931
!
interface Serial0/0.3 point-to-point
description to PEscara
ip address 10.1.1.13 255.255.255.252
ip router isis
tag-switching ip
frame-relay interface-dlci 937
!
router isis
net 47.0000.3333.3333.3333.00
metric-style wide
!
ip classless
!
end

```

Pomerol

*!--- The Pomerol router is an autonomous system border router (ASBR)
!--- like Pulligny and its configuration is very similar to Pomerol's.*

```
hostname Pomerol
!
ip cef
!
interface Loopback0
 ip address 10.10.10.3 255.255.255.255
 ip router isis
!
interface Serial0/0
 no ip address
 encapsulation frame-relay
!
interface Serial0/0.1 point-to-point
 description to Pulligny
 ip address 10.1.1.9 255.255.255.252
 frame-relay interface-dlci 921
!
interface Serial0/0.3 point-to-point
 description to PEsaro
 ip address 10.1.1.21 255.255.255.252
 ip router isis
 tag-switching ip
 frame-relay interface-dlci 925
!
router isis
 redistribute connected metric 20
 net 47.0000.2222.2222.2222.00
 metric-style wide
!
router bgp 2
 no bgp default ipv4-unicast
 no bgp default route-target filter
 bgp log-neighbor-changes
 neighbor 10.1.1.10 remote-as 1
 neighbor 10.10.10.6 remote-as 2
 neighbor 10.10.10.6 update-source Loopback0
!
 address-family vpnv4
 neighbor 10.1.1.10 activate
 neighbor 10.1.1.10 send-community both
 neighbor 10.10.10.6 activate
 neighbor 10.10.10.6 send-community extended
 exit-address-family
!
ip classless
!
end
```

PEsaro

*!--- The PEsaro router is a PE router like the PEsara router, and its
!--- configuration is very similar.*

```
hostname PEsaro
!
ip vrf one
 rd 1:1
 route-target export 1:1
```

```

route-target import 1:1
ip cef
!
interface Loopback0
 ip address 10.10.10.6 255.255.255.255
 ip router isis
!
interface Serial0/1
 no ip address
 encapsulation frame-relay
!
interface Serial0/1.1 point-to-point
 description to Pomerol
 ip address 10.1.1.22 255.255.255.252
 ip router isis
 tag-switching ip
 frame-relay interface-dlci 952
!
interface Serial0/1.2 point-to-point
 description to CErvela
 ip vrf forwarding one
 ip address 220.0.0.1 255.255.255.0
 frame-relay interface-dlci 958
!
router isis
 net 47.0000.5555.5555.00
 metric-style wide
!
router bgp 2
 no synchronization
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 redistribute connected
 neighbor 10.10.10.3 remote-as 2
 neighbor 10.10.10.3 update-source Loopback0
 no auto-summary
!
 address-family ipv4 vrf one
 redistribute connected
 no auto-summary
 no synchronization
 exit-address-family
!
 address-family vpnv4
 neighbor 10.10.10.3 activate
 neighbor 10.10.10.3 send-community both
 no auto-summary
 exit-address-family
!
ip classless
!
end

```

Configuration Notes

- eBGP between both ASs cannot be multihop.
- LDP or TDP is not required between both ASs.
- A common IGP is not required between both ASs. Each AS can use its own distinct IGP.
- BGP creates a connected host route for the eBGP peer in the remote AS once the session that needs to be injected into the IGP of each AS comes up.
- Unless the ASBR is a PE for every VRF that you want to exchange, you will need to configure the **no bgp default route-target filter** command to make sure that the ASBR accepts the BGP VPNv4 prefixes from the other PE routers inside the AS. The default behavior is to deny incoming VPNv4

prefixes that are not imported into any local VRF.

Verifying and Understanding Inter-AS MPLS VPN

This section provides information which you can use 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.

For more information on the commands used below, please see How to Troubleshoot the MPLS VPN.

The following examples show the packet flow from 220.0.0.0/24 which is a connected VRF one prefix on PEsaro to 120.0.0.0/24 which is a connected VRF one prefix on PEsara.

PEsaro Router

Below are the routes for VFR one in the PEsaro router.

```
PEsaro# show ip route vrf one
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```
C    220.0.0.0/24 is directly connected, Serial0/1.2
     120.0.0.0/24 is subnetted, 1 subnets
B    120.0.0.0 [200/0] via 10.1.1.10, 00:26:49
```

```
PEsaro# show ip cef vrf one 120.0.0.1
120.0.0.0/24, version 10, cached adjacency to Serial0/1.1
0 packets, 0 bytes
tag information set
  local tag: VPN-route-head
  fast tag rewrite with Se0/1.1, point2point, tags imposed: {17 22}
via 10.1.1.10, 0 dependencies, recursive
  next hop 10.1.1.21, Serial0/1.1 via 10.1.1.10/32
  valid cached adjacency
  tag rewrite with Se0/1.1, point2point, tags imposed: {17 22}
```

In the above output, you can see that the next hop for 120.0.0.0/24 is 10.1.1.10. You can also see that the tags imposed are 17 and 22. The TFIB and BGP information (below) shows you the source of these tags: the BGP route for 120.0.0.0/24 has an outgoing tag of 22, and the next hop that is associated with 120.0.0.0/24 has an outgoing tag of 17.

```
PEsaro# show ip bgp vpnv4 vrf one tags
Network          Next Hop      In tag/Out tag
Route Distinguisher: 1:1 (one)
  120.0.0.0/24   10.1.1.10    notag/22

!--- VPN label.

  220.0.0.0      0.0.0.0      18/aggregate(one)
```

```
PEsaro# show tag forwarding 10.1.1.10
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
```

```

tag      tag or VC      or Tunnel Id      switched      interface
17       17              10.1.1.10/32     0             Se0/1.1      point2point

```

```

PEsaro# show ip route 10.1.1.10
Routing entry for 10.1.1.10/32
  Known via "isis", distance 115, metric 30, type level-2
  Redistributing via isis
  Last update from 10.1.1.21 on Serial0/1.1, 00:30:39 ago
  Routing Descriptor Blocks:
    * 10.1.1.21, from 10.10.10.3, via Serial0/1.1
      Route metric is 30, traffic share count is 1

```

10.1.1.10 is advertised by Pomerol. It is the BGP-created host route, which is important because it shows up as being the BGP next-hop for the remote VPNv4 prefixes, as shown in the output above. PEsaro sends a packet with two labels to Pomerol; the inner or IGP label is 17.

Pomerol Router

```

Pomerol# show tag forwarding
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC  or Tunnel Id    switched   interface
16     Pop tag    10.10.10.6/32   0          Se0/0.3   point2point
17     Pop tag    10.1.1.10/32    540        Se0/0.1   point2point
18     18         1:1:220.0.0.0/24 1704       Se0/0.3   point2point

```

```

Pomerol# show ip route 10.1.1.10
Routing entry for 10.1.1.10/32
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via isis
  Advertised by isis metric 20 metric-type internal level-2
  Routing Descriptor Blocks:
    * directly connected, via Serial0/0.1
      Route metric is 0, traffic share count is 1

```

Note: The /32 route for 10.1.1.10 is created as soon as eBGP comes up.

In the above output, you can see that Pomerol will pop the top label if it is 17; the eBGP VPN label will be the only label left on packets sent to Pulligny in AS1.

Pulligny Router

```

Pulligny# show tag forwarding
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC  or Tunnel Id    switched   interface
16     Pop tag    10.1.1.9/32     540        Se0/0.2   point2point
17     Pop tag    10.1.1.12/30    0          Se0/0.1   point2point
18     Pop tag    10.10.10.1/32   0          Se0/0.1   point2point
19     16         10.10.10.4/32   0          Se0/0.1   point2point
21     Pop tag    10.1.1.4/30     0          Se0/0.1   point2point
22     16         1:1:120.0.0.0/24 1080       Se0/0.1   point2point

```

The output above shows that Pulligny has local and outgoing tags for 120.0.0.0/24; consider, however, the following detailed TFIB entry:

```

Pulligny# show tag forwarding detail | begin 22
22     16         1:1:120.0.0.0/24 1080        Se0/0.1   point2point
      MAC/Encaps=4/12, MTU=1496, Tag Stack{16 23}
      E4118847 0001000000017000
      No output feature configured

```

The output above shows that Pulligny has a VPNv4 entry in its TFIB, which is necessary because there is only one label over the inter-AS link. Since Pomerol popped label 17 from the label stack, Pulligny receives a packet with label 22. It will pop label 22 and push label 23 and label 16. By looking at the TFIB and BGP information (shown below) you can see that 16 is the IGP label to get to PEscara, and 23 is the VPN label that PEscara advertises for 1:1:120.0.0.0/24.

```
Pulligny# show ip bgp vpnv4 all tagging
Network          Next Hop          In tag/Out tag
Route Distinguisher: 1:1
 120.0.0.0/24    10.10.10.4        22/23
 220.0.0.0       10.1.1.9          notag/18

Pulligny# show tagging for 10.10.10.4
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC  or Tunnel Id    switched   interface
19     16         10.10.10.4/32   0          Se0/0.1   point2point
```

Notice that Pulligny changes the VPN label that it gets from PEscara and advertises a different label to Pomerol. In doing so, it adds an entry for 1:1:120.0.0.0/24 to its TFIB. This behavior is the default because the label changes whenever the next-hop changes (such as in eBGP sessions between both ASs); therefore, the ASBR will have a TFIB entry for each VPN prefix inside its local AS. The output below shows that Pauillac has a simple label swap or pop because the PE is the next-hop (Penultimate Hop Popping [PHP]).

```
Pauillac# show tag forwarding
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC  or Tunnel Id    switched   interface
16     Pop tag    10.10.10.4/32   111049     Se0/0.3   point2point
17     Pop tag    10.10.10.2/32   94769      Se0/0.1   point2point
18     Pop tag    10.1.1.8/30     0          Se0/0.1   point2point
19     16         10.1.1.9/32     560        Se0/0.1   point2point
```

When PEscara receives label 23, it will recognize it as being part of VRF one and it will perform an IP lookup to forward it to the correct interface (as shown in the output below).

```
PEscara# show tag forwarding
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC  or Tunnel Id    switched   interface
16     Pop tag    10.1.1.0/30     0          Se0/1.1   point2point
17     Pop tag    10.1.1.4/30     0          Se0/1.1   point2point
18     Pop tag    10.10.10.1/32   0          Se0/1.1   point2point
19     17         10.10.10.2/32   0          Se0/1.1   point2point
21     18         10.1.1.8/30     0          Se0/1.1   point2point
22     19         10.1.1.9/32     0          Se0/1.1   point2point
23     Aggregate  120.0.0.0/24[V] 1040

PEscara# show tag forwarding detail | begin 23
23     Aggregate  120.0.0.0/24[V] 1040
      MAC/Encaps=0/0, MTU=0, Tag Stack{}
      VPN route: one
      Feature Quick flag set
```

The output below is an example of one packet being sent using the packet flow explained above.

```
Pomerol#
*Mar  1 19:49:59.103: TAG: Se0/0.3: recvd: CoS=0, TTL=255, Tag(s)=17/22
*Mar  1 19:49:59.103: TAG: Se0/0.1: xmit: CoS=0, TTL=254, Tag(s)=22

Pulligny#
*Mar  1 20:21:35: TAG: Se0/0.2: recvd: CoS=0, TTL=254, Tag(s)=22
*Mar  1 20:21:35: TAG: Se0/0.1: xmit: CoS=0, TTL=253, Tag(s)=16/23

Pauillac#
```

```
*Mar 1 19:56:39.627: TAG: Se0/0.1: recvd: CoS=0, TTL=253, Tag(s)=16/23
*Mar 1 19:56:39.627: TAG: Se0/0.3: xmit: CoS=0, TTL=252, Tag(s)=23
```

```
PEscara#
```

```
*Mar 1 19:57:00.463: TAG: Se0/1.1: recvd: CoS=0, TTL=252, Tag(s)=23
```

The next-hop-self Command on ASBRs

If you choose to configure the **next-hop-self** command towards the iBGP peers on the ASBRs, the ASBR advertises a different label (because a new label is advertised when the next-hop changes) to the iBGP peers, adding an entry for the remote AS VPNv4 routes in its TFIB. This creates an additional state in the ASBRs. You do not need to redistribute the BGP-created host route into the IGP. As shown below, the ASBR now has both entries for local and remote VPN prefixes.

```
Pomerol# show tag forwarding
Local  Outgoing  Prefix          Bytes tag  Outgoing  Next Hop
tag    tag or VC  or Tunnel Id    switched   interface
16     Pop tag    10.10.10.6/32   0          Se0/0.3   point2point
17     Pop tag    10.1.1.10/32    540        Se0/0.1   point2point
18     18        1:1:220.0.0.0/24 0          Se0/0.3   point2point
19     22        1:1:120.0.0.0/24 0          Se0/0.1   point2point
```

Troubleshoot

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

Related Information

- [Inter-Autonomous Systems for MPLS VPNs](#)
 - [MPLS Tech Notes](#)
 - [IGRP Technical Support](#)
 - [IP Switching Technical Support](#)
 - [MPLS Technical Support](#)
 - [Technical Support – Cisco Systems](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Aug 10, 2005

Document ID: 5811
