

Network Configuration Example

Implementing Interprovider Layer 3 VPN Option B



Published: 2014-01-10

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Network Configuration Example Implementing Interprovider Layer 3 VPN Option B
NCE0002
Copyright © 2014, Juniper Networks, Inc.
All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

Introduction	1
Interprovider Layer 3 VPN Option B Overview	1
Applications	1
Implementation	2
Example: Configuring Interprovider Layer 3 VPN Option B	3

Introduction

This document describes one of four recommended interprovider and carrier-of-carriers solutions for situations in which the customer of a VPN service provider might be another service provider rather than an end customer.

Interprovider Layer 3 VPN Option B Overview

A customer service provider depends on a virtual private network (VPN) service provider (SP) to deliver a VPN transport service between the customer service provider's points of presence (POPs) or regional networks.

If the customer service provider's sites have different autonomous system (AS) numbers, then the VPN transit service provider supports carrier-of-carriers VPN service for the interprovider VPN service. This functionality might be used by a VPN customer who has connections to several different Internet service providers (ISPs), or different connections to the same ISP in different geographic regions, each of which has a different AS number.

Applications

A customer might require VPN services for different sites, yet the same SP is not available for all of those sites.

RFC 4364 suggests several methods to resolve this problem, including:

- Interprovider VRF-to-VRF connections at the AS boundary routers (ASBR) (not very scalable). This option is presented in *Implementing Interprovider Layer 3 VPN Option A*.
- Interprovider external border gateway protocol (EBGP) redistribution of labeled VPN-IPv4 routes from AS to neighboring AS (somewhat scalable). This option is presented in *Implementing Interprovider Layer 3 VPN Option B*.
- Interprovider multihop EBGP redistribution of labeled VPN-IPv4 routes between source and destination ASs, with EBGP redistribution of labeled IPv4 routes from AS to neighboring AS (very scalable). This option is presented in *Implementing Interprovider Layer 3 VPN Option C*.

Solutions might include elements of both the interprovider VPN solutions and the carrier-of-carriers solution. For example, a transit carrier might supply a service provider whose sites have different AS numbers, which makes the solution topology look like an interprovider solution (because of the different AS numbers). However, it is the same service for the transit carrier, so it really is a carrier-of-carriers service. This type of service solution is referred to as carrier-of-carriers VPN service for the interprovider VPN service.

In contrast, if the customer service provider's sites have the same AS number, then the VPN transit service provider delivers a carrier-of-carriers VPN service.

In addition to resolving the initial problem described above, carrier-of-carriers or interprovider VPN solutions may be used to solve other problems such as scalability and merging two service providers.

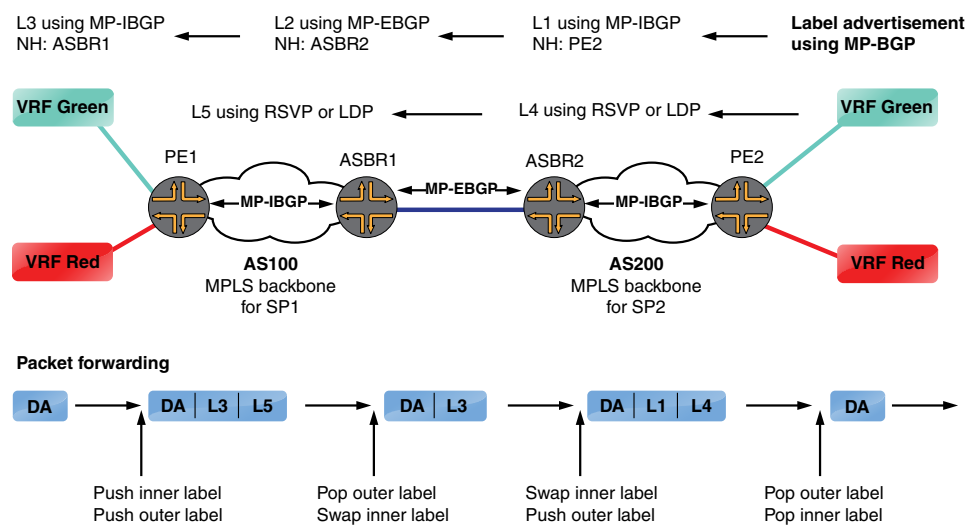
Implementation

In this solution, ASBR routers keep all VPN-IPv4 routes in the routing information base (RIB), and the labels associated with the prefixes are kept in the forwarding information base (FIB). Because the RIB and FIB tables can take too much of the respective allocated memory, this solution is not very scalable for an interprovider VPN.

If a transit SP is used between SP1 and SP2, the transit SP also has to keep all VPN-IPv4 routes in the RIB and the corresponding labels in the FIB. The ASBRs at the transit SP have the same functionality as ASBRs at SP1 or SP2 in this solution.

The provider edge (PE) routers within an AS use multiprotocol internal BGP (MP-IBGP) to distribute labeled VPN-IPv4 routes to an AS boundary router or to a route reflector of which the AS boundary router is a client. The AS boundary router uses MP-EBGP to distribute the labeled VPN-IPv4 routes to its peer AS boundary router in the neighboring AS. The peer AS boundary router then uses multiprotocol internal BGP (MP-IBGP) to distribute labeled VPN-IPv4 routes to PE routers, or to a route reflector of which the PE routers are a client. The logical topology of the network is shown in [Figure 1 on page 2](#).

Figure 1: Logical Topology of Interprovider Layer 3 VPN Option B



Note: The assumption is that explicit null signaling is used in all actions in this figure.

DA = Destination address
NH = Next hop
L1, L2, etc. = Label 1, Label 2

g040500

Related Documentation

- [Example: Configuring Interprovider Layer 3 VPN Option B on page 3](#)

Example: Configuring Interprovider Layer 3 VPN Option B

Interprovider Layer 3 VPN Option B provides interprovider EBGp redistribution of labeled VPN-IPv4 routes from AS to neighboring AS. This solution is considered to be more scalable than Option A, but not as scalable as Option C.

This example provides a step-by-step procedure to configure interprovider layer 3 VPN option B, which is one of the recommended implementations of an MPLS VPN when that service is required by a customer that has more than one AS, but not all of the customer's ASs can be serviced by the same service provider. It is organized in the following sections:

- [Requirements on page 3](#)
- [Configuration Overview and Topology on page 3](#)
- [Configuration on page 5](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.5 or later.
- Eight M Series, T Series, TX Series, or MX Series Juniper Networks routers.



NOTE: This configuration example has been tested using the software release listed and is assumed to work on all later releases.

Configuration Overview and Topology

Interprovider layer 3 VPN option B is a somewhat scalable solution to the problem of providing VPN services to a customer that has different sites, not all of which can use the same service provider. *RFC 4364*, section 10, refers to this method as interprovider EBGp redistribution of labeled VPN-IPv4 routes from AS to neighboring AS.

In the topology shown in Figure 1, the following events occur:

- The PE routers use IBGP to redistribute labeled VPN-IPv4 routes either to an ASBR, or to a route reflector of which an ASBR is a client.
- The ASBR then uses EBGp to redistribute those labeled VPN-IPv4 routes to an ASBR in another AS, which distributes them to the PE routers in that AS, or to another ASBR for distribution.
- Labeled VPN-IPv4 routes are distributed between ASBR routers on each site. There is no need to define a separate VPN routing and forwarding instance (VRF) for each common VPN that resides on two different SPs.
- Router PE2 distributes VPN-IPv4 routes to Router ASBR2 using MP-IBGP.
- Router ASBR2 distributes these labeled VPN-IPv4 routes to Router ASBR1, using the MP-EBGP session between them.

- Router ASBR1 redistributes those routes to Router PE1, using MP-IBGP. Each time a label is advertised, routers change the next-hop information and labels.
- An MPLS path is established between Router PE1 and Router PE2. This path enables changing of the next-hop attribute for the routes that are learned from the neighbor SP router and map the incoming label for the given routes to the outgoing label advertised to PE routers in the internal network.
- The ingress PE router inserts two labels onto the IP packet coming from the end customer. The inner label is for the VPN-IPv4 routes learned from internal ASBRs and the outer label is for the route to the internal ASBR, obtained through resource reservation protocol (RSVP) or label distribution protocol (LDP).
- When a packet arrives at the ASBR, it removes the outer label (when explicit-null signaling is used; otherwise, penultimate hop-popping (PHP) pops the label) and swaps the inner label with the label obtained from the neighbor ASBR through MP-EBGP label and prefix advertisements.
- The second ASBR swaps the VPN-IPv4 label and pushes another label to reach the PE router in its own AS.
- The remaining process is the same as for a regular VPN.




NOTE: In this solution, ASBR routers keep all VPN-IPv4 routes in the routing information base (RIB), and the labels associated with the prefixes are kept in the forwarding information base (FIB). Because the RIB and FIB tables can take occupy much of the respective allocated memory, this solution is not very scalable for an interprovider VPN.

If a transit SP is used between SP1 and SP2, the transit SP also has to keep all VPN-IPv4 routes in the RIB and the corresponding labels in the FIB. The ASBRs at the transit SP have the same functionality as ASBRs in the SP1 or SP2 networks in this solution.

The topology of the network is shown in [Figure 2 on page 5](#).

Figure 2: Physical Topology of Interprovider Layer 3 VPN Option B



Configuration



NOTE: The procedure presented here is written with the assumption that the reader is already familiar with MPLS MVPN configuration. This example focuses on explaining the unique configuration required for carrier-of-carriers solutions for VPN services to different sites.

To configure layer 3 VPN option B, perform the following tasks:

- [Configuring Router CE1 on page 5](#)
- [Configuring Router PE1 on page 6](#)
- [Configuring Router P1 on page 9](#)
- [Configuring Router ASBR1 on page 10](#)
- [Configuring Router ASBR2 on page 11](#)
- [Configuring Router P2 on page 13](#)
- [Configuring Router PE2 on page 14](#)
- [Configuring Router CE2 on page 17](#)
- [Verifying the VPN Operation on page 18](#)

Configuring Router CE1

Step-by-Step Procedure

1. On Router CE1, configure the IP address and protocol family on the Fast Ethernet interface for the link between Router CE1 and Router PE1. Specify the **inet** address family type.

```
[edit interfaces fe-0/0/1.0]
family inet {
  address 18.18.18.1/30;
}
```

2. On Router CE1, configure the IP address and protocol family on the loopback interface. Specify the **inet** address family type.

```
[edit interfaces lo0]
unit 0 {
  family inet {
    address 1.1.1.1/32;
  }
}
```

3. On Router CE1, configure an IGP. Include the logical interface for the link between Router CE1 and Router PE1 and the logical loopback interface of Router CE1. The IGP can be a static route, RIP, OSPF, ISIS, or EBGp. In this example we configure OSPF.

```
[edit protocols]
ospf {
  area 0.0.0.2 {
    interface fe-0/0/1.0;
    interface lo0.0;
  }
}
```

Configuring Router PE1

Step-by-Step Procedure

1. On Router PE1, configure IPv4 addresses on the SONET, Fast Ethernet, and logical loopback interfaces. Specify the **inet** address family on all of the interfaces. Specify the **mpls** address family on the SONET and Fast Ethernet interfaces.

```
[edit interfaces]
so-0/2/0 {
  unit 0 {
    family inet {
      address 19.19.19.1/30;
    }
    family mpls;
  }
}
fe-1/2/3 {
  unit 0 {
    family inet {
      address 18.18.18.2/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 2.2.2.2/32;
    }
  }
}
```

```
}  
}
```

2. On Router PE1, configure the routing instance for VPN2. Specify the **vrf** instance type and specify the customer-facing Fast Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Apply the VRF import and export policies to enable the sending and receiving of route targets. Configure the OSPF protocol within the VRF. Specify the customer-facing Fast Ethernet interface and specify the export policy to export BGP routes into OSPF.

```
[edit routing-instances]  
vpn2CE1 {  
  instance-type vrf;  
  interface fe-1/2/3.0;  
  route-distinguisher 1:100;  
  vrf-import vpnimport;  
  vrf-export vpnexport;  
  protocols {  
    ospf {  
      export bgp-to-ospf;  
      area 0.0.0.2 {  
        interface fe-1/2/3.0;  
      }  
    }  
  }  
}
```

3. On Router PE1, configure the RSVP and MPLS protocols to support the label-switched path (LSP). Configure the LSP to Router ASBR1 and specify the IP address of the logical loopback interface on Router ASBR1. Configure a BGP group. Specify the group type as **internal**. Specify the local address as the logical loopback interface on Router PE1. Specify the neighbor address as the logical loopback interface on Router ASBR1. Specify the **inet-vpn** address family and **unicast** traffic type to enable BGP to carry IPv4 network layer reachability information (NLRI) for VPN routes. Configure the OSPF protocol. Specify the core-facing SONET interface and specify the logical loopback interface on Router PE1.

```
[edit protocols]  
rsvp {  
  interface so-0/2/0.0;  
  interface lo0.0;  
}  
mpls {  
  label-switched-path To-ASBR1 {  
    to 4.4.4.4;  
  }  
  interface so-0/2/0.0;  
  interface lo0.0;  
}  
bgp {  
  group To_ASBR1 {  
    type internal;  
    local-address 2.2.2.2;  
    neighbor 4.4.4.4 {  
      family inet-vpn {
```

```
        unicast;
    }
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface so-0/2/0.0;
        interface lo0.0;
    }
}
```

4. On Router PE1, configure the BGP local autonomous system number.

```
[edit routing-options]
autonomous-system 100;
```

5. On Router PE1, configure a policy to export the BGP routes into OSPF.

```
[edit policy-options]
policy-statement bgp-to-ospf {
    term 1 {
        from protocol bgp;
        then accept;
    }
    term 2 {
        then reject;
    }
}
```

6. On Router PE1, configure a policy to add the VRF route target to the routes being advertised for this VPN.

```
[edit policy-options]
policy-statement vpnexport {
    term 1 {
        from protocol ospf;
        then {
            community add test_comm;
            accept;
        }
    }
    term 2 {
        then reject;
    }
}
```

7. On Router PE1, configure a policy to import routes from BGP that have the **test_comm** community attached.

```
[edit policy-options]
policy-statement vpnimport {
    term 1 {
        from {
            protocol bgp;
            community test_comm;
        }
    }
}
```

```

        then accept;
    }
    term 2 {
        then reject;
    }
}

```

8. On Router PE1, define the **test_comm** BGP community with a route target.

```

[edit policy-options]
community test_comm members target:1:100;

```

Configuring Router P1

Step-by-Step Procedure

1. On Router P1, configure IP addresses for the SONET and Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** address families. Configure the IP addresses for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```

[edit interfaces]
so-0/2/1 {
    unit 0 {
        family inet {
            address 19.19.19.2/30;
        }
        family mpls;
    }
}
ge-1/3/0 {
    unit 0 {
        family inet {
            address 20.20.20.1/30;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 3.3.3.3/32;
        }
    }
}

```

2. On Router P1, configure the RSVP and MPLS protocols to support the LSP. Specify the SONET and Gigabit Ethernet interfaces.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```

[edit protocols]
rsvp {
    interface so-0/2/1.0;
    interface ge-1/3/0.0;
    interface lo0.0;
}

```

```
}
mpls {
  interface lo0.0;
  interface ge-1/3/0.0;
  interface so-0/2/1.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-1/3/0.0;
    interface so-0/2/1.0;
    interface lo0.0;
  }
}
```

Configuring Router ASBR1

Step-by-Step Procedure

1. On Router ASBR1, configure IP addresses for the Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** addresses families. Configure the IP addresses for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 20.20.20.2/30;
    }
    family mpls;
  }
}
ge-0/1/1 {
  unit 0 {
    family inet {
      address 21.21.21.1/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 4.4.4.4/32;
    }
  }
}
```

2. On Router ASBR1, configure the RSVP and MPLS protocols to support the LSP. Specify the Gigabit Ethernet interfaces and the **lo0.0** logical loopback interface.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```
[edit protocols]
rsvp {
```

```

interface ge-0/0/0.0;
interface lo0.0;
}
mpls {
  label-switched-path To_PE1 {
    to 2.2.2.2;
  }
  interface lo0.0;
  interface ge-0/0/0.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0;
  }
}
}

```

3. On Router ASBR1, create the **To-PE1** internal BGP peer group. Specify the local IP peer address as the local **lo0.0** address. Specify the neighbor IP peer address as the **lo0.0** interface address of Router PE1.

```

[edit protocols]
bgp {
  group To-PE1 {
    type internal;
    local-address 4.4.4.4;
    neighbor 2.2.2.2 {
      family inet-vpn {
        unicast;
      }
    }
  }
}

```

4. On Router ASBR1, create the **To-ASBR2** external BGP peer group. Enable the router to use BGP to advertise NLRI for unicast routes. Specify the neighbor IP peer address as the Gigabit Ethernet interface address of Router ASBR2.

```

[edit protocols]
bgp {
  group To-ASBR2 {
    type external;
    family inet-vpn {
      unicast;
    }
    neighbor 21.21.21.2 {
      peer-as 200;
    }
  }
}

```

Configuring Router ASBR2

- | | |
|-------------------------------|---|
| Step-by-Step Procedure | <ol style="list-style-type: none"> 1. On Router ASBR2, configure IP addresses for the Gigabit Ethernet interfaces. Enable the interfaces to process the inet and mpls address families. Configure the IP address |
|-------------------------------|---|

for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```
[edit interfaces]
ge-0/1/1 {
  unit 0 {
    family inet {
      address 21.21.21.2/30;
    }
    family mpls;
  }
}
ge-0/2/3 {
  unit 0 {
    family inet {
      address 22.22.22.1/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 5.5.5.5/32;
    }
  }
}
```

2. On Router ASBR2, configure the RSVP and MPLS protocols to support the LSP. Specify the Gigabit Ethernet interfaces.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```
[edit protocols]
rsvp {
  interface ge-0/2/3.0;
  interface lo0.0;
}
mpls {
  label-switched-path To_PE2 {
    to 7.7.7.7;
  }
  interface lo0.0;
  interface ge-0/2/3.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/2/3.0;
    interface lo0.0;
  }
}
```


-
3. On Router ASBR2, create the **To-PE2** internal BGP peer group. Specify the local IP peer address as the local **lo0.0** address. Specify the neighbor IP peer address as the **lo0.0** interface address of Router PE2.

```
[edit protocols]
bgp {
  group To-PE2 {
    type internal;
    local-address 5.5.5.5;
    neighbor 7.7.7.7 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
```

4. On Router ASBR2, create the **To-ASBR1** external BGP peer group. Enable the router to use BGP to advertise NLRI for unicast routes. Specify the neighbor IP peer address as the Gigabit Ethernet interface on Router ASBR1.

```
[edit protocols]
bgp {
  group To-ASBR1 {
    type external;
    family inet-vpn {
      unicast;
    }
    neighbor 21.21.21.1 {
      peer-as 100;
    }
  }
}
```

Configuring Router P2

Step-by-Step Procedure

1. On Router P2, configure IP addresses for the SONET and Gigabit Ethernet interfaces. Enable the interfaces to process the **inet** and **mpls** addresses families. Configure the IP addresses for the **lo0.0** loopback interface and enable the interface to process the **inet** address family.

```
[edit interfaces]
so-0/0/0 {
  unit 0 {
    family inet {
      address 23.23.23.1/30;
    }
    family mpls;
  }
}
ge-0/2/2 {
  unit 0 {
    family inet {
      address 22.22.22.2/30;
    }
  }
}
```

```
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 6.6.6.6/32;
        }
    }
}
```

2. On Router P2, configure the RSVP and MPLS protocols to support the LSP. Specify the SONET and Gigabit Ethernet interfaces.

Configure the OSPF protocol. Specify the SONET and Gigabit Ethernet interfaces and specify the logical loopback interface. Enable OSPF to support traffic engineering extensions.

```
[edit protocols]
rsvp {
    interface so-0/0/0.0;
    interface ge-0/2/2.0;
    interface lo0.0;
}
mpls {
    interface lo0.0;
    interface ge-0/2/2.0;
    interface so-0/0/0.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/2/2.0;
        interface so-0/0/0.0;
        interface lo0.0;
    }
}
```

Configuring Router PE2

Step-by-Step Procedure

1. On Router PE2, configure IPv4 addresses on the SONET, Fast Ethernet, and logical loopback interfaces. Specify the **inet** address family on all of the interfaces. Specify the **mpls** address family on the SONET and Fast Ethernet interfaces.

```
[edit interfaces]
so-0/0/1 {
    unit 0 {
        family inet {
            address 23.23.23.2/30;
        }
        family mpls;
    }
}
fe-0/3/1 {
    unit 0 {
        family inet {
```

```

        address 24.24.24.1/30;
    }
    family mpls;
}
}
lo0 {
    unit 0 {
        family inet {
            address 7.7.7.7/32;
        }
    }
}
}

```

2. On Router PE2, configure the routing instance for VPN2. Specify the **vrf** instance type and specify the customer-facing Fast Ethernet interface. Configure a route distinguisher to create a unique VPN-IPv4 address prefix. Apply the VRF import and export policies to enable the sending and receiving of route targets. Configure the BGP peer group within the VRF. Specify AS **20** as the peer AS and specify the IP address of the Fast Ethernet interface on Router CE1 as the neighbor address.

```

[edit routing-instances]
vpn2CE2 {
    instance-type vrf;
    interface fe-0/3/1.0;
    route-distinguisher 1:100;
    vrf-import vpnimport;
    vrf-export vpnexport;
    protocols {
        bgp {
            group To_CE2 {
                peer-as 20;
                neighbor 24.24.24.2;
            }
        }
    }
}

```

3. On Router PE2, configure the RSVP and MPLS protocols to support the LSP. Configure the LSP to ASBR2 and specify the IP address of the logical loopback interface on Router ASBR2. Configure a BGP group. Specify the group type as **internal**. Specify the local address as the logical loopback interface on Router PE2. Specify the neighbor address as the logical loopback interface on the Router ASBR2. Specify the **inet-vpn** address family and **unicast** traffic type to enable BGP to carry IPv4 NLRI for VPN routes. Configure the OSPF protocol. Specify the core-facing SONET interface and the logical loopback interface on Router PE2.

```

[edit protocols]
rsvp {
    interface so-0/0/1.0;
    interface lo0.0;
}
mpls {
    label-switched-path To-ASBR2 {
        to 5.5.5.5;
    }
}

```

```
interface so-0/0/1.0;
interface lo0.0;
}
bgp {
  group To_ASBR2 {
    type internal;
    local-address 7.7.7.7;
    neighbor 5.5.5.5 {
      family inet-vpn {
        unicast;
      }
    }
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface so-0/0/1.0;
    interface lo0.0;
  }
}
```

4. On Router PE2, configure the BGP local autonomous system number.

```
[edit routing-options]
autonomous-system 200;
```

5. On Router PE2, configure a policy to add the VRF route target to the routes being advertised for this VPN.

```
[edit policy-options]
policy-statement vpnexport {
  term 1 {
    from protocol bgp;
    then {
      community add test_comm;
      accept;
    }
  }
  term 2 {
    then reject;
  }
}
```

6. On Router PE2, configure a policy to import routes from BGP that have the **test_comm** community attached.

```
[edit policy-options]
policy-statement vpnimport {
  term 1 {
    from {
      protocol bgp;
      community test_comm;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}
```

```
}  
}
```

7. On Router PE1, define the **test_comm** BGP community with a route target.

```
[edit policy-options]  
community test_comm members target:1:100;
```

Configuring Router CE2

Step-by-Step Procedure

1. On Router CE2, configure the IP address and protocol family on the Fast Ethernet interface for the link between Router CE2 and Router PE2. Specify the **inet** address family type.

```
[edit interfaces]  
fe-3/0/0 {  
  unit 0 {  
    family inet {  
      address 24.24.24.2/30;  
    }  
  }  
}
```

2. On Router CE2, configure the IP address and protocol family on the loopback interface. Specify the **inet** address family type.

```
[edit interfaces]  
lo0 {  
  unit 0 {  
    family inet {  
      address 8.8.8.8/32;  
    }  
  }  
}
```

3. On Router CE2, configure an IGP. The IGP can be a static route, RIP, OSPF, ISIS, or EBGP. In this example, we configure EBGP. Specify AS **200** as the peer AS and specify the BGP neighbor IP address as the Fast Ethernet interface of Router PE2. Include the **export** statement.

```
[edit protocols]  
bgp {  
  group To_PE2 {  
    neighbor 24.24.24.1 {  
      export myroutes;  
      peer-as 200;  
    }  
  }  
}
```

Verifying the VPN Operation

Step-by-Step Procedure

1. Commit the configuration on each router.



NOTE: The MPLS labels shown in this example will be different than the labels used in your configuration.

2. On Router PE1, display the routes for the **vpn2CE1** routing instance using the **show ospf route** command. Verify that the 1.1.1.1 route is learned from OSPF.

```
user@PE1> show ospf route instance vpn2CE1
Topology default Route Table:
```

Prefix	Path	Route	NH	Metric	NextHop	NextHop
	Type	Type	Type		Interface	addr/label
1.1.1.1	Intra	Router	IP	1	fe-1/2/3.0	18.18.18.1
1.1.1.1/32	Intra	Network	IP	1	fe-1/2/3.0	18.18.18.1

3. On Router PE1, use the **show route advertising-protocol** command to verify that Router PE1 advertises the 1.1.1.1 route to Router ASBR1 using MP-BGP with the VPN MPLS label.

```
user@PE1> show route advertising-protocol bgp 4.4.4.4 extensive
vpn2CE1.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 1.1.1.1/32 (1 entry, 1 announced)
  BGP group To_ASBR1 type Internal
    Route Distinguisher: 1:100
    VPN Label: 299952
    Nexthop: Self
    Flags: Nexthop Change
    MED: 1
    Localpref: 100
    AS path: [100] I
    Communities: target:1:100 rte-type:0.0.0.2:1:0
```

4. On Router ASBR1, use the **show route receive-protocol** command to verify that the router receives and accepts the 1.1.1.1 route and places it in the **bgp.l3vpn.0** routing table.

```
user@ASBR1> show route receive-protocol bgp 2.2.2.2 extensive
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 1:100:1.1.1.1/32 (1 entry, 1 announced)
  Route Distinguisher: 1:100
  VPN Label: 299952
  Nexthop: 2.2.2.2
  MED: 1
  Localpref: 100
  AS path: I
  Communities: target:1:100 rte-type:0.0.0.2:1:0
```

-
5. On Router ASBR1, use the **show route advertising-protocol** command to verify that Router ASBR1 advertises the 1.1.1.1 route to Router ASBR2.

```
user@ASBR1> show route advertising-protocol bgp 21.21.21.2 extensive
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 1:100:1.1.1/32 (1 entry, 1 announced)
  BGP group To-ASBR2 type External
    Route Distinguisher: 1:100
    VPN Label: 299984
    Nexthop: Self
    Flags: Nexthop Change
    AS path: [100] I
    Communities: target:1:100 rte-type:0.0.0.2:1:0
```

6. On Router ASBR2, use the **show route receive-protocol** command to verify that the router receives and accepts the 1.1.1.1 route and places it in the **bgp.l3vpn.0** routing table.

```
user@ASBR2> show route receive-protocol bgp 21.21.21.1 extensive
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 1:100:1.1.1/32 (1 entry, 1 announced)
  Accepted
  Route Distinguisher: 1:100
  VPN Label: 299984
  Nexthop: 21.21.21.1
  AS path: 100 I
  Communities: target:1:100 rte-type:0.0.0.2:1:0
```

7. On Router ASBR2, use the **show route advertising-protocol** command to verify that Router ASBR2 advertises the 1.1.1.1 route to Router PE2 in the **To-PE2** routing instance.

```
user@ASBR2> show route advertising-protocol bgp 7.7.7.7 extensive
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 1:100:1.1.1/32 (1 entry, 1 announced)
  BGP group To-PE2 type Internal
    Route Distinguisher: 1:100
    VPN Label: 300048
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [200] 100 I
    Communities: target:1:100 rte-type:0.0.0.2:1:0
```

8. On Router PE2, use the **show route receive-protocol** command to verify that the router receives and accepts the 1.1.1.1 route and places it in the **To_CE2.inet.0** routing table.

```
user@PE2> show route receive-protocol bgp 5.5.5.5 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

__juniper_private1__.inet.0: 14 destinations, 14 routes (8 active, 0 holddown,
6 hidden)
```

```
__juniper_private2__.inet.0: 1 destinations, 1 routes (0 active, 0 holddown,
1 hidden)
```

```
To_CE2.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
```

```
* 1.1.1/32 (1 entry, 1 announced)
```

```
Accepted
```

```
Route Distinguisher: 1:100
```

```
VPN Label: 300048
```

```
Nexthop: 5.5.5.5
```

```
Localpref: 100
```

```
AS path: 100 I
```

```
AS path: Recorded
```

```
Communities: target:1:100 rte-type:0.0.0.2:1:0
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```
bgp.l3vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
* 1:100:1.1.1/32 (1 entry, 0 announced)
```

```
Accepted
```

```
Route Distinguisher: 1:100
```

```
VPN Label: 300048
```

```
Nexthop: 5.5.5.5
```

```
Localpref: 100
```

```
AS path: 100 I
```

```
AS path: Recorded
```

```
Communities: target:1:100 rte-type:0.0.0.2:1:0
```

```
__juniper_private1__.inet6.0: 4 destinations, 4 routes (4 active, 0 holddown,
0 hidden)
```

9. On Router PE2, use the **show route advertising-protocol** command to verify that Router PE2 advertises the 1.1.1 route to Router CE2 through the **To_CE2** peer group.

```
user@PE2> show route advertising-protocol bgp 24.24.24.2 extensive
```

```
To_CE2.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
```

```
* 1.1.1/32 (1 entry, 1 announced)
```

```
BGP group To_CE2 type External
```

```
Nexthop: Self
```

```
AS path: [200] 100 I
```

```
Communities: target:1:100 rte-type:0.0.0.2:1:0
```

10. On Router CE2, use the **show route** command to verify that Router CE2 receives the 1.1.1 route from Router PE2.

```
user@CE2> show route 1.1.1
```

```
inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
1.1.1.1/32          *[BGP/170] 00:25:36, localpref 100
                    AS path: 200 100 I
                    > to 24.24.24.1 via fe-3/0/0.0
```

11. On Router CE2, use the **ping** command and specify **8.8.8.8** as the source of the ping packets to verify connectivity with Router CE1.

```
user@CE2> ping 1.1.1.1 source 8.8.8.8
```

```
PING 1.1.1.1 (1.1.1.1): 56 data bytes
```

```
64 bytes from 1.1.1.1: icmp_seq=0 ttl=58 time=4.786 ms
```

```
64 bytes from 1.1.1.1: icmp_seq=1 ttl=58 time=10.210 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=58 time=10.588 ms
```

12. On Router PE2, use the **show route** command to verify that the traffic is sent with an inner label of **300048** and a top label of **299776**.

```
user@PE2> show route 1.1.1.1 detail
To_CE2.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
1.1.1.1/32 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
                Route Distinguisher: 1:100
                Next hop type: Indirect
                Next-hop reference count: 3
                Source: 5.5.5.5
                Next hop type: Router, Next hop index: 653
                Next hop: via so-0/0/1.0 weight 0x1, selected
                Label-switched-path To-PE2
                Label operation: Push 300048, Push 299776(top)
                Protocol next hop: 5.5.5.5
                Push 300048
                Indirect next hop: 8c61138 262143
                State: <Secondary Active Int Ext>
                Local AS: 200 Peer AS: 200
                Age: 27:48 Metric2: 2
                Task: bgp_200.5.5.5+60185
                Announcement bits (3): 0-RT 1-KRT 2-BGP RT Background
                AS path: 100 I
                AS path: Recorded
                Communities: target:1:100 rte-type:0.0.0.2:1:0
                Accepted
                VPN Label: 300048
                Localpref: 100
                Router ID: 5.5.5.5
                Primary Routing Table bgp.l3vpn.0
```

13. On Router ASBR2, use the **show route table** command to verify that Router ASBR2 receives the traffic after the top label is popped by Router P2, that label **300048** is swapped with label **299984**, and that the packet is sent toward Router ASBR1 through interface **ge-0/1/1.0**.

```
user@ASBR2> show route table mpls.0 detail
300048 (1 entry, 1 announced)
    *VPN      Preference: 170
                Next hop type: Router, Next hop index: 648
                Next-hop reference count: 2
                Source: 21.21.21.1
                Next hop: 21.21.21.1 via ge-0/1/1.0, selected
                Label operation: Swap 299984
                State: <Active Int Ext>
                Local AS: 200
                Age: 30:39
                Task: BGP RT Background
                Announcement bits (1): 0-KRT
                AS path: 100 I
                Ref Cnt: 1
                Communities: target:1:100 rte-type:0.0.0.2:1:0
```

14. On Router ASBR1, use the **show route table** command to verify that Router ASBR1 receives the traffic with label **299984**, swaps the label with **299952**, and pushes a new top label of **299792**.

```
user@ASBR1> show route table mpls.0 detail
299984 (1 entry, 1 announced)
  *VPN      Preference: 170
            Next hop type: Indirect
            Next-hop reference count: 2
            Source: 2.2.2.2
            Next hop type: Router, Next hop index: 538
            Next hop: 20.20.20.1 via ge-0/0/0.0 weight 0x1, selected
            Label-switched-path To_PE1
            Label operation: Swap 299952, Push 299792(top)
            Protocol next hop: 2.2.2.2
            Swap 299952
            Indirect next hop: 8af70a0 262142
            State: <Active Int Ext>
            Local AS: 100
            Age: 34:09      Metric2: 2
            Task: BGP RT Background
            Announcement bits (1): 0-KRT
            AS path: I
            Ref Cnt: 1
            Communities: target:1:100 rte-type:0.0.0.2:1:0
```

15. On Router PE1, use the **show route table** command to verify that Router PE1 receives the traffic with label **299952**, and then pops the inner label.

```
user@PE1> show route route table mpls.0 detail
299952 (1 entry, 1 announced)
  *VPN      Preference: 170
            Next hop type: Router, Next hop index: 536
            Next-hop reference count: 2
            Next hop: 18.18.18.1 via fe-1/2/3.0, selected
            Label operation: Pop
            State: Active Int Ext
            Local AS: 100
            Age: 40:26
            Task: BGP RT Background
            Announcement bits (1): 0-KRT
            AS path: I
            Ref Cnt: 1
            Communities: rte-type:0.0.0.2:1:0
```

Related Documentation • [Interprovider Layer 3 VPN Option B Overview on page 1](#)