

Network Configuration Example

Interconnecting a Layer 2 Circuit with a Layer 3 VPN

Release

11.4



Published: 2011-11-08

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

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Network Configuration Example Interconnecting a Layer 2 Circuit with a Layer 3 VPN

Release 11.4

Copyright © 2011, Juniper Networks, Inc.

All rights reserved.

Revision History

October 2011—R1 Junos OS 11.4

The information in this document is current as of the date listed in the revision history.

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

Layer 2 Circuit Overview	1
Layer 3 VPN Overview	3
Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN	5
Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN	7

Layer 2 Circuit Overview

A Layer 2 circuit is a point-to-point Layer 2 connection transported using Multiprotocol Label Switching (MPLS) or other tunneling technology on the service provider's network. A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple virtual circuits (VCs) are transported over a single shared label-switched path (LSP) tunnel between two provider edge (PE) routers. In contrast, each CCC requires a separate dedicated LSP.

To establish a Layer 2 circuit, the Link Integrity Protocol (LIP) is used as the signaling protocol to advertise the ingress label to the remote PE routers. For this purpose, a targeted remote LDP neighbor session is established using the extended discovery mechanism described in LDP, and the session is brought up to the remote PE loopback IP address. Because LDP looks at the Layer 2 circuit configuration and initiates extended neighbor discovery for all the Layer 2 circuit neighbors (the remote PEs), no new configuration is necessary in LDP. Each Layer 2 circuit is represented by the logical interface connecting the local PE router to the local customer edge (CE) router. Note that LDP must be enabled on the lo0.0 interface for extended neighbor discovery to function correctly.

Packets are sent to remote CE routers over an egress VPN label advertised by the remote PE router, using a targeted LDP session. The VPN label is sent over an LDP LSP to the remote PE router connected to the remote CE router. Return traffic from the remote CE router destined to the local CE router is sent using an ingress VPN label advertised by the local PE router, which is also sent over the LDP LSP to the local PE router from the remote PE router.

Related Documentation

- [Layer 3 VPN Overview on page 3](#)
- Layer 2 VPN Overview
- Layer 2 VPN Applications
- Applications for Interconnecting a Layer 2 Circuit with a Layer 2 Circuit
- [Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 5](#)
- Example: Interconnecting a Layer 2 Circuit with a Layer 2 Circuit
- [Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 7](#)
- Example: Interconnecting a Layer 2 Circuit with a Layer 2 VPN

Layer 3 VPN Overview

Layer 3 VPNs are based on RFC 2547bis, *BGP/MPLS IP VPNs*. RFC 2547bis defines a mechanism by which service providers can use their IP backbones to provide VPN services to their customers. A Layer 3 VPN is a set of sites that share common routing information and whose connectivity is controlled by a collection of policies. The sites that make up a Layer 3 VPN are connected over a provider's existing Internet backbone. RFC 2547bis VPNs are also known as BGP/MPLS VPNs because BGP is used to distribute VPN routing information across the provider's backbone, and MPLS is used to forward VPN traffic across the backbone to remote VPN sites.

Customer networks, because they are private, can use either public addresses or private addresses, as defined in RFC 1918, *Address Allocation for Private Internets*. When customer networks that use private addresses connect to the Internet infrastructure, the private addresses might overlap with the same private addresses used by other network users. MPLS/BGP VPNs solve this problem by adding a *route distinguisher*, a VPN identifier prefix to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and within the Internet.

In addition, each VPN has its own VPN-specific routing table that contains the routing information for that VPN only. To separate VPN routes from routes in the Internet or those in other VPNs, the PE router creates a separate routing table for each VPN called a VPN routing and forwarding (VRF) table. The PE router creates one VRF table for each VPN that has a connection to a customer edge (CE) router. Any customer or site that belongs to the VPN can access only the routes in the VRF tables for that VPN. Every VRF table has one or more extended community attributes associated with it that identify the route as belonging to a specific collection of routers. One of these, the *route target* attribute, identifies a collection of sites (VRF tables) to which a PE router distributes routes. The PE router uses the route target to constrain the import of remote routes into its VRF tables.

When an ingress PE router receives routes advertised from a directly connected CE router, it checks the received route against the VRF export policy for that VPN.

- If the route matches, the route is converted to VPN-IPv4 format—that is, the route distinguisher is added to the route. The PE router then announces the route in VPN-IPv4 format to the remote PE routers. It also attaches a route target to each route learned from the directly connected sites. The route target attached to each route is based on the configured export target policy of the VRF table. The routes are then distributed using IBGP sessions, which are configured in the provider's core network.
- If the route from the CE router does not match, it is not exported to other PE routers, but it can still be used locally for routing, for example, if two CE routers in the same VPN are directly connected to the same PE router.

When an egress PE router receives a route, it checks it against the import policy on the IBGP session between the PE routers. If it passes, the router places the route into its `bgp.l3vpn.0` table. At the same time, the router checks the route against the VRF import policy for the VPN. If it matches, the route distinguisher is removed from the route and

the route is placed into the VRF table (the *routing-instance-name*.inet.0 table) in IPv4 format.

**Related
Documentation**

- [Layer 2 Circuit Overview on page 1](#)
- Layer 2 VPN Overview
- Interconnecting Layer 2 VPNs with Layer 3 VPNs Overview
- [Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 5](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 7](#)
- Example: Interconnecting a Layer 2 VPN with a Layer 3 VPN

Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN

MPLS-based Layer 2 services are growing in demand among enterprise and service providers. This creates new challenges related to interoperability between Layer 2 and Layer 3 services for service providers who want to provide end-to-end value-added services. There are various reasons to stitch different Layer 2 services to one another and to Layer 3 services. For example, to expand the service offerings and to expand geographically. The Junos OS has various features to address the needs of the service provider.

Interconnecting a Layer 2 Circuit with a Layer 3 VPN includes the following benefits:

- Interconnecting a Layer 2 Circuit with a Layer 3 VPN enables the sharing of a service provider's core network infrastructure between IP and Layer 2 circuit services, reducing the cost of providing those services. A Layer 2 MPLS circuit allows service providers to create a Layer 2 circuit service over an existing IP and MPLS backbone.
- Service providers do not have to invest in separate Layer 2 equipment to provide Layer 2 circuit service. A service provider can configure a provider edge router to run any Layer 3 protocol in addition to the Layer 2 protocols. Customers who prefer to maintain control over most of the administration of their own networks want Layer 2 circuit connections with their service provider instead of a Layer 3 VPN connection.

Related Documentation

- [Layer 2 Circuit Overview on page 1](#)
- [Layer 3 VPN Overview on page 3](#)
- [Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 7](#)

Example: Interconnecting a Layer 2 Circuit with a Layer 3 VPN

This example provides a step-by-step procedure and commands for configuring and verifying a Layer 2 circuit to Layer 3 VPN interconnection. It contains the following sections:

- [Requirements on page 7](#)
- [Overview and Topology on page 7](#)
- [Configuration on page 9](#)
- [Verifying the Layer 2 Circuit to Layer 3 VPN Interconnection on page 19](#)

Requirements


This example uses the following hardware and software components:

- Junos OS Release 9.3 or later
- 3 MX Series routers
- 1 M Series routers
- 1 T Series router
- 1 EX Series router
- 1 J Series router

Overview and Topology

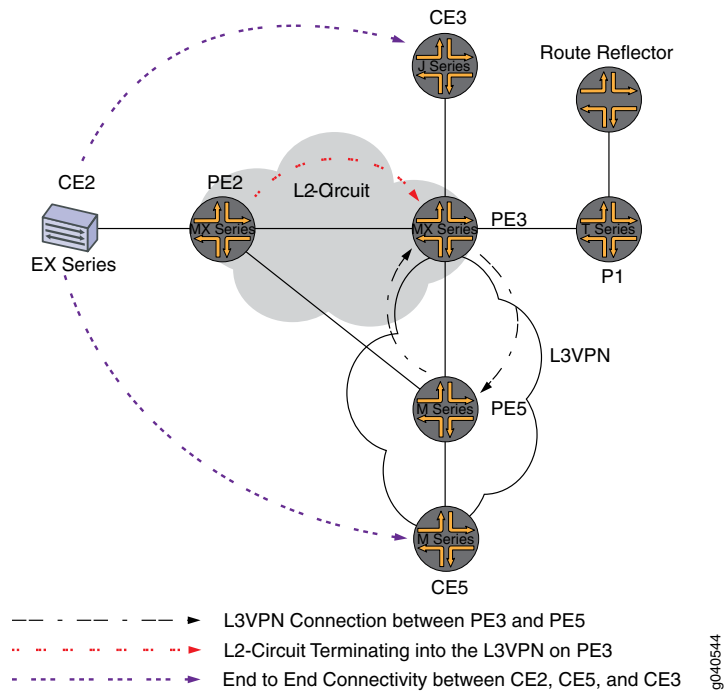
The physical topology of a Layer 2 circuit to Layer 3 VPN interconnection is shown in [Figure 1 on page 8](#).

Figure 1: Physical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection



The logical topology of a Layer 2 circuit to Layer 3 VPN interconnection is shown in [Figure 2 on page 8](#).

Figure 2: Logical Topology of a Layer 2 Circuit to Layer 3 VPN Interconnection



Configuration



NOTE: In any configuration session, it is good practice to verify periodically that the configuration can be committed using the `commit check` command.

In this example, the router being configured is identified using the following command prompts:

- **CE2** identifies the customer edge 2 (CE2) router
- **PE1** identifies the provider edge 1 (PE1) router
- **CE3** identifies the customer edge 3 (CE3) router
- **PE3** identifies the provider edge 3 (PE3) router
- **CE5** identifies the customer edge 5 (CE5) router
- **PE5** identifies the provider edge 5 (PE5) router

This example contains the following procedures:

- [Configuring PE Router Customer-facing and Loopback Interfaces on page 9](#)
- [Configuring Core-facing Interfaces on page 11](#)
- [Configuring Protocols on page 12](#)
- [Configuring Routing Instances and Layer 2 Circuits on page 15](#)
- [Configuring the Route Reflector on page 17](#)
- [Interconnecting the Layer 2 Circuit with the Layer 3 VPN on page 18](#)

Configuring PE Router Customer-facing and Loopback Interfaces

Step-by-Step Procedure

To begin building the interconnection, configure the interfaces on the PE routers. If your network contains provider (P) routers, configure the interfaces on the P routers also. This example shows the configuration for Router PE2, Router PE3, and Router PE5.

1. On Router PE2, configure the **ge-1/0/2** interface encapsulation. To configure the interface encapsulation, include the **encapsulation** statement and specify the **ethernet-ccc** option (**vlan-ccc** encapsulation is also supported). Configure the **ge-1/0/2.0** logical interface family for circuit cross-connect functionality. To configure the logical interface family, include the **family** statement and specify the **ccc** option. The encapsulation should be configured the same way for all routers in the Layer 2 circuit domain.

```
[edit interfaces]
ge-1/0/2 {
  encapsulation ethernet-ccc;
  unit 0 {
    family ccc;
  }
}
```

2. On Router PE2, configure the **lo0.0** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **2.2.2.2/32** as the loopback IPv4 address.

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address 2.2.2.2/32;
    }
  }
}
```

3. On Router PE3, configure the **ge-1/0/1** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **90.90.90.1/24** as the interface address for this device.

```
[edit interfaces]
ge-1/0/1 {
  unit 0 {
    family inet {
      address 90.90.90.1/24;
    }
  }
}
```

4. On Router PE3, configure the **lo0.0** loopback interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **3.3.3.3/32** as the loopback IPv4 address for this router.

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address 3.3.3.3/32;
    }
  }
}
```

5. On Router PE5, configure the **ge-2/0/0** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **80.80.80.1/24** as the interface address.

```
[edit interfaces]
ge-2/0/0 {
  unit 0 {
    family inet {
      address 80.80.80.1/24;
    }
  }
}
```

6. On Router PE5, configure the **lo0.0** interface. Include the **family** statement and specify the **inet** option. Include the **address** statement and specify **5.5.5.5/32** as the loopback IPv4 address for this router.

```
[edit interfaces]
lo0 {
```

```

    unit 0 {
      family inet {
        address 5.5.5.5/32;
      }
    }
  }
}

```

Configuring Core-facing Interfaces

Step-by-Step Procedure

This procedure describes how to configure the core-facing interfaces on the PE routers. This example does not include all the core-facing interfaces shown in the physical topology illustration. Enable the **mpls** and **inet** address families on the core-facing interfaces.

1. On Router PE2, configure the **xe-0/2/0** interface. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify **10.10.5.1/30** as the interface address. Include the **family** statement and specify the **mpls** address family.

```

[edit interfaces]
xe-0/2/0 {
  unit 0 {
    family inet {
      address 10.10.5.1/30;
    }
    family mpls;
  }
}

```

2. On Router PE3, configure the core-facing interfaces. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify the IPv4 addresses shown in the example as the interface addresses. Include the **family** statement and specify the **mpls** address family. In the example, the **xe-2/1/0** interface is connected to Router PE5, and the **xe-2/2/0** interface is connected to Router PE2.

```

[edit interfaces]
xe-2/0/0 {
  unit 0 {
    family inet {
      address 10.10.20.2/30;
    }
    family mpls;
  }
}
xe-2/1/0 {
  unit 0 {
    family inet {
      address 10.10.6.1/30;
    }
    family mpls;
  }
}
xe-2/2/0 {
  unit 0 {
    family inet {

```

```
        address 10.10.5.2/30;
    }
    family mpls;
}
}
xe-2/3/0 {
    unit 0 {
        family inet {
            address 10.10.1.2/30;
        }
        family mpls;
    }
}
```

3. On Router PE5, configure the **xe-0/1/0** interface. Include the **family** statement and specify the **inet** address family. Include the **address** statement and specify **10.10.6.2/30** as the interface address. Include the **family** statement and specify the **mpls** address family.

```
[edit interfaces]
xe-0/1/0 {
    unit 0 {
        family inet {
            address 10.10.6.2/30;
        }
        family mpls;
    }
}
```

Configuring Protocols

Step-by-Step Procedure

This procedure describes how to configure the protocols used in this example. If your network contains P routers, configure the interfaces on the P routers also.

1. On Router PE3, enable OSPF as the IGP. Enable the MPLS, LDP, and BGP protocols on all interfaces except **fxp0.0**. LDP is used as the signaling protocol for the Layer 2 circuit to Router PE2. The following configuration snippet shows the protocol configuration for Router PE3:

```
[edit]
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path to-RR {
            to 7.7.7.7;
        }
        label-switched-path to-PE2 {
            to 2.2.2.2;
        }
        label-switched-path to-PE5 {
```



```

        to 5.5.5.5;
    }
    label-switched-path to-PE4 {
        to 4.4.4.4;
    }
    label-switched-path to-PE1 {
        to 1.1.1.1;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group RR {
        type internal;
        local-address 3.3.3.3;
        family inet-vpn {
            unicast;
        }
        family l2vpn {
            signaling;
        }
        neighbor 7.7.7.7;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
}

```

2. On Router PE2, configure the MPLS, OSPF, and LDP protocols.

```

[edit ]
protocols {
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {

```

```
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
}
```

3. On Router PE5, enable OSPF as the IGP. Enable the MPLS, RSVP, and BGP protocols on all interfaces except **fxp0.0**. Enable core-facing interfaces with the **mpls** and **inet** address families.

```
[edit]
protocols {
    rsvp {
        interface all {
            link-protection;
        }
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path to-RR {
            to 7.7.7.7;
        }
        label-switched-path to-PE2 {
            to 2.2.2.2;
        }
        label-switched-path to-PE3 {
            to 3.3.3.3;
        }
        label-switched-path to-PE4 {
            to 4.4.4.4;
        }
        label-switched-path to-PE1 {
            to 1.1.1.1;
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
bgp {
    group to-rr {
        type internal;
        local-address 5.5.5.5;
        family inet-vpn {
            unicast;
        }
    }
}
```

```

        family l2vpn {
            signaling;
        }
        neighbor 7.7.7.7;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
}
}

```

Configuring Routing Instances and Layer 2 Circuits

Step-by-Step Procedure

This procedure describes how to configure the Layer 2 circuit and the Layer 3 VPN.

1. On Router PE2, configure the Layer 2 circuit. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE3 as the neighbor. Include the interface statement and specify **ge-1/0/2.0** as the logical interface that is participating in the Layer 2 circuit. Include the **virtual-circuit-id** statement and specify **100** as the identifier. Include the **no-control-word** statement for equipment that does not support the control word.

```

[edit ]
protocols {
    l2circuit {
        neighbor 3.3.3.3 {
            interface ge-1/0/2.0 {
                virtual-circuit-id 100;
                no-control-word;
            }
        }
    }
}
}

```

2. On Router PE3, configure the Layer 2 circuit to Router PE2. Include the **l2circuit** statement. Include the **neighbor** statement and specify the loopback IPv4 address of Router PE2 as the neighbor. Include the interface statement and specify **lt-1/1/10.0** as the logical tunnel interface that is participating in the Layer 2 circuit. Include the **virtual-circuit-id** statement and specify **100** as the identifier. Include the **no-control-word** statement.

```

[edit ]
protocols {
    l2circuit {
        neighbor 2.2.2.2 {
            interface lt-1/1/10.0 {
                virtual-circuit-id 100;
                no-control-word;
            }
        }
    }
}

```

```
    }  
  }  
}
```

3. On Router PE3, configure the Layer 3 VPN (**L3VPN**) routing instance to Router PE5 at the **[edit routing-instances]** hierarchy level. Also configure the BGP peer group at the **[edit routing-instances L3VPN protocols]** hierarchy level.

```
[edit ]  
routing-instances {  
  L3VPN {  
    instance-type vrf;  
    interface ge-1/0/1.0;  
    interface lt-1/1/10.1;  
    route-distinguisher 65000:33;  
    vrf-target target:65000:2;  
    vrf-table-label;  
    protocols {  
      bgp {  
        export direct;  
        group ce3 {  
          neighbor 90.90.90.2 {  
            peer-as 100;  
          }  
        }  
      }  
    }  
  }  
}
```

4. On Router PE5, configure the Layer 3 VPN routing instance (**L3VPN**) at the **[edit routing-instances]** hierarchy level. Also configure the BGP peer group at the **[edit routing-instances L3VPN protocols]** hierarchy level.

```
[edit ]  
routing-instances {  
  L3VPN {  
    instance-type vrf;  
    interface ge-2/0/0.0;  
    route-distinguisher 65000:5;  
    vrf-target target:65000:2;  
    vrf-table-label;  
    protocols {  
      bgp {  
        group ce5 {  
          neighbor 80.80.80.2 {  
            peer-as 200;  
          }  
        }  
      }  
    }  
  }  
}
```

Configuring the Route Reflector

Step-by-Step Procedure Although a route reflector is not required to interconnect a Layer 2 circuit with a Layer 3 VPN, this examples uses a route reflector. This procedure shows the relevant portion of the route reflector configuration.

1. Configure the route reflector with RSVP, MPLS, BGP and OSPF. The route reflector is a BGP peer with the PE routers. Notice that the BGP peer group configuration includes the **family** statement and specifies the **inet-vpn** option. The **inet-vpn** option enables BGP to advertise network layer reachability information (NLRI) for the Layer 3 VPN routes. The configuration also includes the **family** statement and specifies the **l2vpn** option. The **l2vpn** option enables BGP to advertise NLRI for the Layer 2 circuit. Layer 2 circuits use the same internal BGP infrastructure as Layer 2 VPNs.

```
[edit ]
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path to-pe3 {
      to 3.3.3.3;
    }
    label-switched-path to-pe5 {
      to 5.5.5.5;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group RR {
      type internal;
      local-address 7.7.7.7;
      family inet {
        unicast;
      }
      family inet-vpn {
        unicast;
      }
      family l2vpn {
        signaling;
      }
      cluster 7.7.7.7;
      neighbor 1.1.1.1;
      neighbor 2.2.2.2;
      neighbor 4.4.4.4;
      neighbor 5.5.5.5;
      neighbor 3.3.3.3;
    }
  }
}
```

```

    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface all;
            interface fxp0.0 {
                disable;
            }
        }
    }
}

```

Interconnecting the Layer 2 Circuit with the Layer 3 VPN

Step-by-Step Procedure

Before you can configure the logical tunnel interface in an MX Series router, you must create the tunnel services interface to be used for tunnel services.

1. Create the tunnel service interface on Router PE3. Include the **bandwidth** statement at the **[edit chassis fpc slot-number pic slot-number tunnel-services]** hierarchy level and specify the amount of bandwidth to reserve for tunnel services in gigabits per second.

```

[edit chassis]
fpc 1 {
    pic 1 {
        tunnel-services {
            bandwidth 1g;
        }
    }
}

```

2. On Router PE3, configure the **lt-1/1/10** logical tunnel interface unit 0.

Router PE3 is the router that is *stitching* the Layer 2 circuit to the Layer 3 VPN using the logical tunnel interface. The configuration of the peer unit interfaces is what makes the interconnection.

Include the **encapsulation** statement and specify the **ethernet-ccc** option. Include the **peer-unit** statement and specify the logical interface unit 1 as the peer tunnel interface. Include the **family** statement and specify the **ccc** option.

Configure the **lt-1/1/10** logical interface unit 1 with **ethernet** encapsulation. Include the **peer-unit** statement and specify the logical interface unit 0 as the peer tunnel interface. Include the **family** statement and specify the **inet** option. Also include the **address** statement and specify **70.70.70.1/24** as the IPv4 address of the interface.



NOTE: The peering logical interfaces must belong to the same logical tunnel interface derived from the Tunnel Services PIC.

```

[edit interfaces]
lt-1/1/10 {
    unit 0 {

```

```

        encapsulation ethernet-ccc;
        peer-unit 1;
        family ccc;
    }
    unit 1 {
        encapsulation ethernet;
        peer-unit 0;
        family inet {
            address 70.70.70.1/24;
        }
    }
}

```

3. On each router, commit the configuration.

```

user@host> commit check
configuration check succeeds
user@host> commit

```

Verifying the Layer 2 Circuit to Layer 3 VPN Interconnection

To verify that the interconnection is working properly, perform these tasks:

- [Verifying That the Layer 2 Circuit Connection to Router PE3 is Up on page 19](#)
- [Verifying LDP Neighbors and Targeted LDP LSPs on Router PE2 on page 20](#)
- [Verifying the Layer 2 Circuit Routes on Router PE2 on page 20](#)
- [Verifying That the Layer 2 Circuit Connection to Router PE2 is Up on page 21](#)
- [Verifying LDP Neighbors and Targeted LDP LSPs on Router PE3 on page 22](#)
- [Verifying a BGP Peer Session with the Route Reflector on Router PE3 on page 22](#)
- [Verifying the Layer 3 VPN Routes on Router PE3 on page 23](#)
- [Verifying the Layer 2 Circuit Routes on Router PE3 on page 23](#)
- [Verifying the MPLS Routes on Router PE3 on page 24](#)
- [Verifying Traffic Flow Between Router CE2 and Router CE3 on page 25](#)
- [Verifying Traffic Flow Between Router CE2 and Router CE5 on page 25](#)

Verifying That the Layer 2 Circuit Connection to Router PE3 is Up

Purpose To verify that the Layer 2 circuit connection from Router PE2 to Router PE3 is **Up**. To also document the incoming and outgoing LDP labels and the circuit ID used by this Layer 2 circuit connection.

Action Verify that the Layer 2 circuit connection is up, using the **show l2circuit connections** command.

```
user@PE2> show l2circuit connections
```

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface h/w not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational

VM -- vlan id mismatch CF -- Call admission control failure
 OL -- no outgoing label IB -- TDM incompatible bitrate
 NC -- intf encaps not CCC/TCC TM -- TDM misconfiguration
 BK -- Backup Connection ST -- Standby Connection
 CB -- rcvd cell-bundle size bad SP -- Static Pseudowire
 LD -- local site signaled down RS -- remote site standby
 RD -- remote site signaled down XX -- unknown

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 3.3.3.3

Interface	Type	St	Time last up	# Up trans
ge-1/0/2.0(vc 100)	rmt	Up	Jan 7 02:14:13 2010	1

Remote PE: 3.3.3.3, Negotiated control-word: No
 Incoming label: 301488, Outgoing label: 315264
 Negotiated PW status TLV: No
 Local interface: ge-1/0/2.0, Status: Up, Encapsulation: ETHERNET

Meaning The output shows that the Layer 2 circuit connection from Router PE2 to Router PE3 is **Up** and the connection is using the **ge-1/0/2.0** interface. Note that the outgoing label is **315264** and the incoming label is **301488**, the virtual circuit (VC) identifier is **100** and the encapsulation is **ETHERNET**.

Verifying LDP Neighbors and Targeted LDP LSPs on Router PE2

Purpose To verify that Router PE2 has a targeted LDP LSP to Router PE3 and that Router PE2 and Router PE3 are LDP neighbors.

Action Verify that Router PE2 has a targeted LDP LSP to Router PE3 and that Router PE2 and Router PE3 are LDP neighbors, using the **show ldp neighbor** command.

```

user@PE2> show ldp neighbor
Address          Interface      Label space ID      Hold time
3.3.3.3          lo0.0          3.3.3.3:0           38
  
```

Meaning The output shows that Router PE2 has an LDP neighbor with the IPv4 address of **3.3.3.3**. Address 3.3.3.3 is the lo0.0 interface address of Router PE3. Notice that Router PE2 uses the local **lo0.0** interface for the LSP.

Verifying that the routers are LDP neighbors also verifies that the targeted LSP is established.

Verifying the Layer 2 Circuit Routes on Router PE2

Purpose To verify that Router PE2 has a route for the Layer 2 circuit and that the route uses the LDP MPLS label to Router PE3.

Action Verify that Router PE2 has a route for the Layer 2 circuit and that the route uses the LDP MPLS label to Router PE3, using the **show route table mpls.0** command.

```

user@PE2> show route table mpls.0
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
  
```



```

0          *[MPLS/0] 1w3d 05:24:11, metric 1
           Receive
1          *[MPLS/0] 1w3d 05:24:11, metric 1
           Receive
2          *[MPLS/0] 1w3d 05:24:11, metric 1
           Receive
300560     *[LDP/9] 16:12:23, metric 1
           > to 10.10.2.1 via xe-0/1/0.0, Pop
300560(S=0) *[LDP/9] 16:12:23, metric 1
           > to 10.10.2.1 via xe-0/1/0.0, Pop
301008     *[LDP/9] 16:12:23, metric 1
           > to 10.10.4.2 via xe-0/3/0.0, Swap 299856
301488     *[L2CKT/7] 11:07:28
           > via ge-1/0/2.0, Pop
301536     *[LDP/9] 16:12:23, metric 1
           > to 10.10.4.2 via xe-0/3/0.0, Pop
301536(S=0) *[LDP/9] 16:12:23, metric 1
           > to 10.10.4.2 via xe-0/3/0.0, Pop
301712     *[LDP/9] 12:41:22, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Swap 315184
301728     *[LDP/9] 12:41:22, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Pop
301728(S=0) *[LDP/9] 12:41:22, metric 1
           > to 10.10.5.2 via xe-0/2/0.0, Pop
ge-1/0/2.0 *[L2CKT/7] 11:07:28, metric2 1
           > to 10.10.5.2 via xe-0/2/0.0, Push 315264

```

Meaning The output shows that Router PE2 pushes the **315264** outgoing label on the **L2CKT** route going out interface **ge-1/0/2.0**. The output also shows that Router PE2 pops the **301488** incoming label on the **L2CKT** coming from interface **ge-1/0/2.0**

Verifying That the Layer 2 Circuit Connection to Router PE2 is Up

Purpose To verify that the Layer 2 circuit connection from Router PE3 to Router PE2 is **Up**, To also document the incoming and outgoing LDP labels and the circuit ID used by this Layer 2 circuit connection.

Action Verify that the Layer 2 circuit connection is up, using the **show l2circuit connections** command.

user@PE3> show l2circuit connections

Layer-2 Circuit Connections:

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface h/w not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational
VM -- vlan id mismatch	CF -- Call admission control failure
OL -- no outgoing label	IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC	TM -- TDM misconfiguration
BK -- Backup Connection	ST -- Standby Connection
CB -- rcvd cell-bundle size bad	XX -- unknown

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 2.2.2.2

```

Interface                Type St      Time last up          # Up trans
lt-1/1/10.0(vc 100)      rmt  Up      Jan  7 02:15:03 2010      1
Remote PE: 2.2.2.2, Negotiated control-word: No
Incoming label: 315264, Outgoing label: 301488
Local interface: lt-1/1/10.0, Status: Up, Encapsulation: ETHERNET

```

Meaning The output shows that the Layer 2 circuit connection from Router PE3 to Router PE2 is **Up** and the connection is using the logical tunnel (lt) interface. Note that the incoming label is **315264** and the outgoing label is **301488**, the virtual circuit (VC) identifier is **100**, and that the encapsulation is **ETHERNET**.

Verifying LDP Neighbors and Targeted LDP LSPs on Router PE3

Purpose To verify that Router PE3 has a targeted LDP LSP to Router PE2 and that Router PE3 and Router PE2 are LDP neighbors.

Action Verify that Router PE2 has a targeted LDP LSP to Router PE3 and that Router PE2 and Router PE3 are LDP neighbors, using the **show ldp neighbor** command.

```

user@PE2> show ldp neighbor
Address          Interface      Label space ID      Hold time
2.2.2.2          lo0.0          2.2.2.2:0           43
4.4.4.4          lo0.0          4.4.4.4:0           33

```

Meaning The output shows that Router PE3 has an LDP neighbor with the IPv4 address of **2.2.2.2**. Address 2.2.2.2 is the lo0.0 interface address of Router PE2. The output also shows that the interface used on Router PE3 for the LSP is **lo0.0**. Verifying that the routers are LDP neighbors also verifies that the targeted LSP is established.

Verifying a BGP Peer Session with the Route Reflector on Router PE3

Purpose To verify that Router PE3 has a peer session established with the route reflector.

Action Verify that Router PE3 has a peer session established with the route reflector, using the **show bgp summary** command.

```
user@PE2> show bgp summary
```

```

Groups: 2 Peers: 2 Down peers: 0
Table      Tot Paths  Act Paths Suppressed  History  Damp State   Pending
bgp.13vpn.0      1          1          0          0          0          0
Peer          AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
7.7.7.7        65000     1597     1612      0        1    12:03:21 Establ
  bgp.12vpn.0: 0/0/0/0
  bgp.13vpn.0: 1/1/1/0
  L3VPN.inet.0: 1/1/1/0

```

Meaning The output shows that Router PE3 has a peer session with the router with the IPv4 address of **7.7.7.7**. Address 7.7.7.7 is the lo0.0 interface address of the route reflector. The output also shows that the peer session state is **Establ**, meaning that the session is established.

Verifying the Layer 3 VPN Routes on Router PE3

Purpose To verify that Router PE3 has Layer 3 VPN routes to Router CE2, Router CE3, and Router CE5.

Action Verify that Router PE3 has routes to Router CE2, Router CE3, and Router CE5 in the Layer 3 VPN route table, using the **show route table L3VPN.inet.0** command. In this example, **L3VPN** is the name configured for the routing instance.

```
user@PE3> show route table L3VPN.inet.0
L3VPN.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

70.70.70.0/24      * [Direct/0] 11:13:59
                  > via lt-1/1/10.1
70.70.70.1/32     * [Local/0] 11:13:59
                  Local via lt-1/1/10.1
80.80.80.0/24     * [BGP/170] 11:00:41, localpref 100, from 7.7.7.7
                  AS path: I
                  > to 10.10.6.2 via xe-2/1/0.0, Push 16
90.90.90.0/24     * [Direct/0] 11:54:41
                  > via ge-1/0/1.0
90.90.90.1/32     * [Local/0] 11:54:41
                  Local via ge-1/0/1.0
```

Meaning The output shows that Router PE3 has a route to the IPv4 subnetwork address of **70.70.70.0**. Address 70.70.70.2 is the interface address of Router CE2. The output shows that Router PE3 has a route to the IPv4 subnetwork address of **80.80.80.0**. Address 80.80.80.2 is the interface address of Router CE5. The output shows that Router PE3 has a route to the IPv4 subnetwork address of **90.90.90.0**. Address 90.90.90.2 is the interface address of Router CE3.

Verifying the Layer 2 Circuit Routes on Router PE3

Purpose To verify that Router PE3 has a route to Router PE2 in the Layer 2 circuit route table.

Action Verify that Router PE3 has a route to Router PE2 in the Layer 2 circuit route table, using the **show route table l2circuit.0** command.

```
user@PE3> show route table l2circuit.0
2.2.2.2:NoCtrlWord:5:100:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop type: Indirect
    Next-hop reference count: 1
    Next hop type: Router
    Next hop: 10.10.5.1 via xe-2/2/0.0, selected
    Protocol next hop: 2.2.2.2
    Indirect next hop: 8cae0a0 -
    State: <Active Int>
    Local AS: 65000
    Age: 11:16:50 Metric2: 1
    Task: l2 circuit
    Announcement bits (1): 0-LDP
```

```
AS path: I
VC Label 315264, MTU 1500
```

Meaning The output shows that Router PE3 has a route to the IPv4 address of 2.2.2.2. Address 2.2.2.2 is the lo0.0 interface address of Router PE2. Note that the VC label is 315264. This label is the same as the incoming MPLS label displayed using the **show l2circuit connections** command.

Verifying the MPLS Routes on Router PE3

Purpose To verify that Router PE3 has a route to Router PE2 in the MPLS route table.

Action Verify Router PE3 has a route to Router PE2 in the MPLS route table, using the **show route table mpls.0** command.

```
user@PE3> show route table mpls.0
mpls.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 1w3d 05:29:02, metric 1
            Receive
1          *[MPLS/0] 1w3d 05:29:02, metric 1
            Receive
2          *[MPLS/0] 1w3d 05:29:02, metric 1
            Receive
16         *[VPN/0] 12:22:45
            to table L3VPN.inet.0, Pop
315184     *[LDP/9] 12:45:14, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, Pop
315184(S=0) *[LDP/9] 12:45:14, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, Pop
315200     *[LDP/9] 00:03:53, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, Swap 625297
            to 10.10.6.2 via xe-2/1/0.0, Swap 299856
315216     *[LDP/9] 12:45:14, metric 1
            > to 10.10.6.2 via xe-2/1/0.0, Pop
315216(S=0) *[LDP/9] 12:45:14, metric 1
            > to 10.10.6.2 via xe-2/1/0.0, Pop
315232     *[LDP/9] 12:45:06, metric 1
            > to 10.10.1.1 via xe-2/3/0.0, Pop
315232(S=0) *[LDP/9] 12:45:06, metric 1
            > to 10.10.1.1 via xe-2/3/0.0, Pop
315248     *[LDP/9] 12:45:14, metric 1
            > to 10.10.5.1 via xe-2/2/0.0, Pop
315248(S=0) *[LDP/9] 12:45:14, metric 1
            > to 10.10.5.1 via xe-2/2/0.0, Pop
315264     *[L2CKT/7] 11:11:20
            > via lt-1/1/10.0, Pop
315312     *[RSVP/7] 11:26:01, metric 1
            > to 10.10.6.2 via xe-2/1/0.0, label-switched-path to-pe5
315312(S=0) *[RSVP/7] 11:26:01, metric 1
            > to 10.10.6.2 via xe-2/1/0.0, label-switched-path to-pe5
315328     *[RSVP/7] 11:26:01, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, label-switched-path to-RR
315360     *[RSVP/7] 11:26:01, metric 1
            > to 10.10.20.1 via xe-2/0/0.0, label-switched-path to-RR
316208     *[RSVP/7] 00:03:32, metric 1
            > to 10.10.6.2 via xe-2/1/0.0, label-switched-path
```

```

Bypass->10.10.9.1
316208(S=0)      *[RSVP/7] 00:03:32, metric 1
                  > to 10.10.6.2 via xe-2/1/0.0, label-switched-path
Bypass->10.10.9.1
lt-1/1/10.0      *[L2CKT/7] 11:11:20, metric2 1
                  > to 10.10.5.1 via xe-2/2/0.0, Push 301488

```

Meaning The output shows that Router PE3 has a route for the Layer 2 circuit and that the route uses the LDP MPLS label to Router PE2. Notice that the **301488** label is the same as the outgoing label displayed on Router PE2 using the **show l2circuit connections** command.

Verifying Traffic Flow Between Router CE2 and Router CE3

Purpose To verify that the CE routers can send and receive traffic across the interconnection.

Action Verify that Router CE2 can send traffic to and receive traffic from Router CE3 across the interconnection, using the **ping** command.

```

user@CE2>ping 90.90.90.2
PING 90.90.90.2 (90.90.90.2): 56 data bytes
64 bytes from 90.90.90.2: icmp_seq=0 ttl=63 time=0.708 ms
64 bytes from 90.90.90.2: icmp_seq=1 ttl=63 time=0.610 ms

```

Meaning The output shows that Router CE2 can send an ICMP request to and receive a response from Router CE3 across the interconnection.

Verifying Traffic Flow Between Router CE2 and Router CE5

Purpose To verify that the CE routers can send and receive traffic across the interconnection.

Action Verify that Router CE2 can send traffic to and receive traffic from Router CE5 across the interconnection, using the **ping** command.

```

user@CE2>ping 80.80.80.2
PING 80.80.80.2 (80.80.80.2): 56 data bytes
64 bytes from 80.80.80.2: icmp_seq=0 ttl=62 time=0.995 ms
64 bytes from 80.80.80.2: icmp_seq=1 ttl=62 time=1.005 ms

```

Meaning The output shows that Router CE2 can send an ICMP request to and receive a response from Router CE5 across the interconnection.

Related Documentation

- [Layer 2 Circuit Overview on page 1](#)
- [Layer 3 VPN Overview on page 3](#)
- [Applications for Interconnecting a Layer 2 Circuit with a Layer 3 VPN on page 5](#)

