Network Configuration Example

Configuring VPLS Multihoming Using Autodiscovery

Release
13.1

Published: 2013-02-11

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

Network Configuration Example Configuring VPLS Multihoming Using Autodiscovery
Release 13.1
NCE0072

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

This document describes the uses for and benefits of using autodiscovery for VPLS multihoming. It also provides a step-by-step configuration example for VPLS multihoming, which includes steps to validate that the configuration is working.

Advantages of Using Autodiscovery for VPLS Multihoming

Virtual private LAN service (VPLS) provides a multipoint-to-multipoint Ethernet service that can span one or more metropolitan areas and provides connectivity between multiple sites as if these sites were attached to the same Ethernet LAN.

VPLS uses an IP and MPLS service provider infrastructure. From a service provider's point of view, use of IP and MPLS routing protocols and procedures instead of the Spanning Tree Protocol (STP), and MPLS labels instead of VLAN IDs, significantly improves the scalability of the VPLS service.

It is frequently a requirement for a service provider to supply its customers with redundant connectivity to one or more sites. This capability is called multihoming.

VPLS multihoming enables you to connect a customer site to two or more PE routers to provide redundant connectivity. A redundant PE router can provide network service to the customer site as soon as a failure is detected. VPLS multihoming helps to maintain VPLS service and traffic forwarding to and from the multihomed site in the event of the following types of network failures:

- PE router to CE device link failure
- PE router failure
- MPLS-reachability failure between the local PE router and a remote PE router

The Junos® operating system (Junos OS) supports both forwarding equivalency class (FEC) 128 and FEC 129. FEC 128 requires manually configured pseudowires. FEC 129 uses VPLS autodiscovery to convey endpoint information. After PE routers are autodiscovered, pseudowires are created automatically.

VPLS multihoming with support for FEC 129 enables you to interoperate with other vendor's auto-discovery for LDP-signaled VPLS. This interoperability allows you to select the vendor that offers the best value.

Related Documentation

- VPLS Multihoming Overview on page 2
- Example: Configuring VPLS Multihoming (FEC 129) on page 3

VPLS Multihoming Overview

Virtual private LAN service (VPLS) multihoming enables you to connect a customer site to two or more PE routers to provide redundant connectivity. A redundant PE router can provide network service to the customer site as soon as a failure is detected. VPLS multihoming helps to maintain VPLS service and traffic forwarding to and from the multihomed site in the event of the following types of network failures:

- PE router to CE device link failure
- PE router failure
- MPLS-reachability failure between the local PE router and a remote PE router

Figure 1: CE Device Multihomed to Two PE Routers

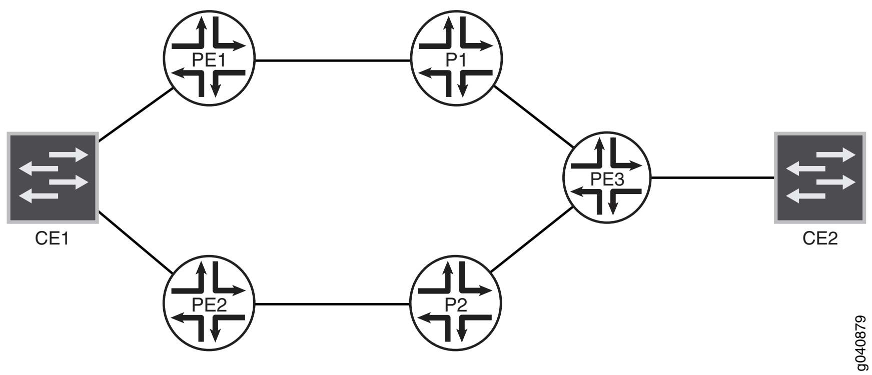

Figure 1 on page 2 illustrates how a CE device could be multihomed to two PE routers. Device CE1 is multihomed to Routers PE1 and PE2. Device CE2 has two potential paths to reach Device CE1, but only one path is active at any one time. If Router PE1 were the designated VPLS edge (VE) device (also called a designated forwarder), BGP would signal a pseudowire from Router PE3 to Router PE1. If a failure occurred over this path, Router PE2 would be made the designated VE device, and BGP would re-signal the pseudowire from Router PE3 to Router PE2.

Multihomed PE routers advertise network layer reachability information (NLRI) for the multihomed site to the other PE routers in the VPLS network. The NLRI includes the site ID for the multihomed PE routers. For all of the PE routers multihomed to the same CE device, you need to configure the same site ID. The remote VPLS PE routers use the site ID to determine where to forward traffic addressed to the customer site. To avoid route collisions, the site ID shared by the multihomed PE routers must be different than the site IDs configured on the remote PE routers in the VPLS network.

Although you configure the same site ID for each of the PE routers multihomed to the same CE device, you can configure unique values for other parameters, such as the route distinguisher. These values help to determine which multihomed PE router is selected as the designated VE device to be used to reach the customer site.

BEST PRACTICE: We recommend that you configure unique route distinguishers for each multihomed PE router. Configuring unique route distinguishers helps with faster convergence when the connection to a primary multihomed PE router goes down. If you configure unique route distinguishers, the other PE routers in the VPLS network must maintain additional state for the multihomed PE routers.

Remote PE routers in the VPLS network need to determine which of the multihomed PE routers should forward traffic to reach the CE device. To make this determination, remote PE routers use the VPLS path-selection process to select one of the multihomed PE routers based on its NLRI advertisement. Because remote PE routers pick only one of the NLRI advertisements, it establishes a pseudowire to only one of the multihomed PE routers, the PE router that originated the winning advertisement. This prevents multiple paths from being created between sites in the network, preventing the formation of Layer 2 loops. If the selected PE router fails, all PE routers in the network automatically switch to the backup PE router and establish new pseudowires to it.

BEST PRACTICE: To prevent the formation of Layer 2 loops between the CE devices and the multihomed PE routers, we recommend that you employ the Spanning Tree Protocol (STP) on your CE devices. Layer 2 loops can form due to incorrect configuration. Temporary Layer 2 loops can also form during convergence after a change in the network topology.

The PE routers run the BGP path selection procedure on locally originated and received Layer 2 route advertisements to establish that the routes are suitable for advertisement to other peers, such as BGP route reflectors. If a PE router in a VPLS network is also a route reflector, the path selection process for the multihomed site has no effect on the path selection process performed by this PE router for the purpose of reflecting Layer 2 routes. Layer 2 prefixes that have different route distinguishers are considered to have different NLRIs for route reflection. The VPLS path selection process enables the route reflector to reflect all routes that have different route distinguishers to the route reflector clients, even though only one of these routes is used to create the VPLS pseudowire to the multihomed site.

Related Documentation

- Configuring VPLS Multihoming (FEC 128)
- Advantages of Using Autodiscovery for VPLS Multihoming on page 1

Example: Configuring VPLS Multihoming (FEC 129)

This example shows how to configure virtual private LAN service (VPLS) multihoming. Multihoming allows a customer site to connect to multiple provider edge (PE) routers. A VPLS site multihomed to two or more PE routers provides redundant connectivity in the event of a PE router-to-CE device link failure or the failure of a PE router. The example demonstrates BGP-based multihoming support for FEC 129 VPLS (also known as LDP VPLS with BGP-based autodiscovery).

Requirements

This example has the following hardware and software requirements:

- One or more CE devices to represent a VPLS site.
- Two or more PE devices.
- Junos OS Release 12.3 or later running on the PE devices that are connected to the multihomed VPLS site.

Overview

BGP-based VPLS autodiscovery (FEC 129) enables each VPLS PE router to discover the other PE routers that are in the same VPLS domain. VPLS autodiscovery also automatically detects when PE routers are added or removed from the VPLS domain. You do not need to manually configure the VPLS and maintain the configuration when a PE router is added or deleted. VPLS autodiscovery uses BGP to discover the VPLS members and to set up and tear down pseudowires in the VPLS.

BGP multihoming enables you to connect a customer site to two or more PE routers to provide redundant connectivity while preventing the formation of Layer 2 loops in the service provider's network. The redundant connectivity maintains the VPLS service and traffic forwarding to and from the multihomed site in the event of a PE router-to-CE device link failure, the failure of a PE router, or an MPLS reachability failure between the local PE router and a remote PE router. A redundant PE router can begin providing service to the customer site as soon as the failure is detected.

When a CE device connects to multiple PE routers, each of these routers advertises reachability for the multihomed site—routes that have the same site ID in the Layer 2 network layer reachability information (NLRI). The other PE routers in the network use a BGP path selection process to select only one of the advertising routers to which they send traffic destined for the CE device. This path selection process eliminates Layer 2 loops in the VPLS network.

Autodiscovery is not specifically related to multihoming. Autodiscovery is not required for multihoming to work. They are two separate features. That said, the meaning of FEC 129 is that VPLS does autodiscovery. So when you configure multihoming for FEC 129, you must also, by definition, configure autodiscovery (with the **auto-discovery-only** statement).

There are two places in the configuration where you can configure VPLS multihoming. One is for FEC 128, and the other is for FEC 129:

- For FEC 128—**routing-instances *instance-name* protocols vpls site *site-name* multi-homing**
- For FEC 129—**routing-instances *instance-name* protocols vpls multi-homing**

The following statements are used for configuring multihoming for FEC 129:

```
[edit routing-instances instance-name protocols vpls]
multi-homing {
  peer-active;
  site site-name {
    active-interface interface-name {
      any;
      primary interface-name;
    }
    identifier identifier;
    interface interface-name {
      preference preference-value;
    }
    peer-active;
    preference (preference-value | backup | primary);
  }
}
```

This example shows Device CE1 multihomed to Router PE1 and Router PE2. In addition, Device CE2 is single-homed to Router PE1. Device PE3 is the remote PE router, connected to Device CE3. Multihoming is not enabled on Device PE3. "CLI Quick Configuration" on page 5 shows the configuration for all of the devices in Figure 2 on page 5. The section "Configuring Device PE1" on page 9 has step-by-step instructions for configuring Device PE1.

Figure 2: Topology for FEC 129 Multihoming

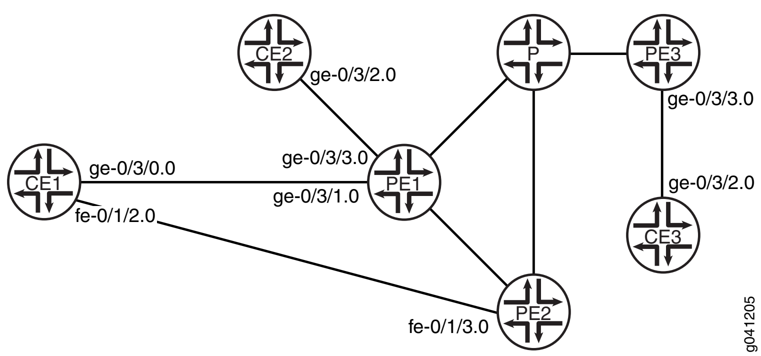

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device PE1

```
set interfaces ge-0/3/3 encapsulation ethernet-vpls
```

```
set interfaces ge-0/3/3 unit 0 description PE1-to-CE2
set interfaces ge-0/3/3 unit 0 family vpls
set interfaces ge-0/3/1 encapsulation ethernet-vpls
set interfaces ge-0/3/1 unit 0 description PE1-to-CE1
set interfaces ge-0/3/1 unit 0 family vpls
set interfaces ge-1/2/0 unit 1 description PE1-to-P
set interfaces ge-1/2/0 unit 1 family inet address 10.1.1.1/30
set interfaces ge-1/2/0 unit 1 family mpls
set interfaces ge-1/2/1 unit 5 description PE1-to-PE2
set interfaces ge-1/2/1 unit 5 family inet address 10.1.1.5/30
set interfaces ge-1/2/1 unit 5 family mpls
set interfaces lo0 unit 2 family inet address 1.1.1.2/32
set protocols mpls interface ge-1/2/0.1
set protocols mpls interface ge-1/2/1.5
set protocols bgp local-address 1.1.1.2
set protocols bgp group pe-pe type internal
set protocols bgp group pe-pe family l2vpn auto-discovery-only
set protocols bgp group pe-pe family l2vpn signaling
set protocols bgp group pe-pe neighbor 1.1.1.3
set protocols bgp group pe-pe neighbor 1.1.1.4
set protocols bgp group pe-pe neighbor 1.1.1.5
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-1/2/0.1
set protocols ospf area 0.0.0.0 interface ge-1/2/1.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ldp interface ge-1/2/0.1
set protocols ldp interface ge-1/2/1.5
set protocols ldp interface lo0.2
set routing-instances green instance-type vpls
set routing-instances green interface ge-0/3/1.0
set routing-instances green interface ge-0/3/3.0
set routing-instances green route-distinguisher 1.1.1.2:1
set routing-instances green l2vpn-id l2vpn-id:100:100
set routing-instances green vrf-target target:100:100
set routing-instances green protocols vpls no-tunnel-services
set routing-instances green protocols vpls oam ping-interval 600
set routing-instances green protocols vpls oam bfd-liveness-detection minimum-interval
  200
set routing-instances green protocols vpls multi-homing site test identifier 1
set routing-instances green protocols vpls multi-homing site test interface ge-0/3/1.0
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 100
```

Device PE2

```
set interfaces fe-0/1/3 encapsulation ethernet-vpls
set interfaces fe-0/1/3 unit 0 description PE2-to-CE1
set interfaces fe-0/1/3 unit 0 family vpls
set interfaces ge-1/2/0 unit 6 description PE2-to-PE1
set interfaces ge-1/2/0 unit 6 family inet address 10.1.1.6/30
set interfaces ge-1/2/0 unit 6 family mpls
set interfaces ge-1/2/2 unit 10 description PE2-to-P
set interfaces ge-1/2/2 unit 10 family inet address 10.1.1.10/30
set interfaces ge-1/2/2 unit 10 family mpls
set interfaces lo0 unit 4 family inet address 1.1.1.4/32
set protocols mpls interface ge-1/2/0.6
set protocols mpls interface ge-1/2/2.10
set protocols bgp local-address 1.1.1.4
```



```
set protocols bgp group pe-pe type internal
set protocols bgp group pe-pe family l2vpn auto-discovery-only
set protocols bgp group pe-pe family l2vpn signaling
set protocols bgp group pe-pe neighbor 1.1.1.2
set protocols bgp group pe-pe neighbor 1.1.1.3
set protocols bgp group pe-pe neighbor 1.1.1.5
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-1/2/0.6
set protocols ospf area 0.0.0.0 interface ge-1/2/2.10
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ldp interface ge-1/2/0.6
set protocols ldp interface ge-1/2/2.10
set protocols ldp interface lo0.4
set routing-instances green instance-type vpls
set routing-instances green interface fe-0/1/3.0
set routing-instances green route-distinguisher 1.1.1.4:1
set routing-instances green l2vpn-id l2vpn-id:100:100
set routing-instances green vrf-target target:100:100
set routing-instances green protocols vpls no-tunnel-services
set routing-instances green protocols vpls oam ping-interval 600
set routing-instances green protocols vpls oam bfd-liveness-detection minimum-interval
  200
set routing-instances green protocols vpls multi-homing site test identifier 1
set routing-instances green protocols vpls multi-homing site test interface fe-0/1/3.0
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 100
```

Device PE3

```
set interfaces ge-0/3/3 unit 0
set interfaces ge-1/2/0 unit 14 description PE3-to-P
set interfaces ge-1/2/0 unit 14 family inet address 10.1.1.14/30
set interfaces ge-1/2/0 unit 14 family mpls
set interfaces lo0 unit 5 family inet address 1.1.1.5/32
set protocols rsvp interface ge-1/2/0.14
set protocols mpls interface ge-1/2/0.14
set protocols bgp local-address 1.1.1.5
set protocols bgp group pe-pe type internal
set protocols bgp group pe-pe family l2vpn auto-discovery-only
set protocols bgp group pe-pe family l2vpn signaling
set protocols bgp group pe-pe neighbor 1.1.1.2
set protocols bgp group pe-pe neighbor 1.1.1.3
set protocols bgp group pe-pe neighbor 1.1.1.4
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-1/2/0.14
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ldp interface ge-1/2/0.14
set protocols ldp interface lo0.5
set routing-instances green instance-type vpls
set routing-instances green interface ge-0/3/3.0
set routing-instances green route-distinguisher 1.1.1.5:100
set routing-instances green l2vpn-id l2vpn-id:100:100
set routing-instances green vrf-target target:100:100
set routing-instances green protocols vpls no-tunnel-services
set routing-instances green protocols vpls oam ping-interval 600
set routing-instances green protocols vpls oam bfd-liveness-detection minimum-interval
  200
set routing-instances green protocols vpls oam ping-interval 600
```

```
set routing-instances green protocols vpls oam bfd-liveness-detection minimum-interval
  200
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 100
```

Device CE1

```
set interfaces ge-0/3/0 unit 0 description CE1-to-PE1
set interfaces ge-0/3/0 unit 0 family inet address 20.1.1.5/24
set interfaces fe-0/1/2 unit 0 description CE1-to-PE2
set interfaces fe-0/1/2 unit 0 family inet address 20.1.1.1/24
```

Device CE2

```
set interfaces ge-0/3/2 unit 0 description CE2-to-PE1
set interfaces ge-0/3/2 unit 0 family inet address 20.1.1.6/24
```

Device CE3

```
set interfaces ge-0/3/2 unit 0 description CE3-to-PE3
set interfaces ge-0/3/2 unit 0 family inet address 20.1.1.7/24
```

Device P

```
set interfaces ge-1/2/0 unit 2 description P-to-PE1
set interfaces ge-1/2/0 unit 2 family inet address 10.1.1.2/30
set interfaces ge-1/2/0 unit 2 family mpls
set interfaces ge-3/2/0 unit 9 description P-to-PE2
set interfaces ge-3/2/0 unit 9 family inet address 10.1.1.9/30
set interfaces ge-3/2/0 unit 9 family mpls
set interfaces ge-4/2/0 unit 13 description P-to-PE3
set interfaces ge-4/2/0 unit 13 encapsulation ethernet
set interfaces ge-4/2/0 unit 13 peer-unit 14
set interfaces ge-4/2/0 unit 13 family inet address 10.1.1.13/30
set interfaces ge-4/2/0 unit 13 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface ge-1/2/0.2
set protocols mpls interface ge-3/2/0.9
set protocols mpls interface ge-4/2/0.13
set protocols bgp local-address 1.1.1.3
set protocols bgp group pe-pe type internal
set protocols bgp group pe-pe family l2vpn signaling
set protocols bgp group pe-pe neighbor 1.1.1.2
set protocols bgp group pe-pe neighbor 1.1.1.4
set protocols bgp group pe-pe neighbor 1.1.1.5
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-1/2/0.2
set protocols ospf area 0.0.0.0 interface ge-3/2/0.9
set protocols ospf area 0.0.0.0 interface ge-4/2/0.13
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ldp interface ge-1/2/0.2
set protocols ldp interface ge-3/2/0.9
set protocols ldp interface ge-4/2/0.13
set protocols ldp interface lo0.3
set routing-options router-id 1.1.1.3
set routing-options autonomous-system 100
```


Configuring Device PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see Using the CLI Editor in Configuration Mode in the CLI User Guide.

To configure Device PE1:

1. Configure the interfaces.

   Configure **family mpls** on the provider-facing interfaces. Configure **family vpls** on the customer-facing interfaces.

   ```
   [edit interfaces]
   user@PE1# set ge-0/3/3 encapsulation ethernet-vpls
   user@PE1# set ge-0/3/3 unit 0 description PE1-to-CE2
   user@PE1# set ge-0/3/3 unit 0 family vpls

   user@PE1# set ge-0/3/1 encapsulation ethernet-vpls
   user@PE1# set ge-0/3/1 unit 0 description PE1-to-CE1
   user@PE1# set ge-0/3/1 unit 0 family vpls

   user@PE1# set ge-1/2/0 unit 1 description PE1-to-P
   user@PE1# set ge-1/2/0 unit 1 family inet address 10.1.1.1/30
   user@PE1# set ge-1/2/0 unit 1 family mpls

   user@PE1# set ge-1/2/1 unit 5 description PE1-to-PE2
   user@PE1# set ge-1/2/1 unit 5 family inet address 10.1.1.5/30
   user@PE1# set ge-1/2/1 unit 5 family mpls

   user@PE1# set lo0 unit 2 family inet address 1.1.1.2/32
   ```

2. Configure the interior gateway protocol (IGP) and signaling protocols on the provider-facing interfaces.

   The **traffic-engineering** statement enables OSPF to advertise the label-switched path (LSP) metric in summary link-state advertisements (LSAs).

   ```
   [edit protocols]
   user@PE1# set ldp interface ge-1/2/0.1
   user@PE1# set ldp interface ge-1/2/1.5
   user@PE1# set ldp interface lo0.2

   user@PE1# set mpls interface ge-1/2/0.1
   user@PE1# set mpls interface ge-1/2/1.5

   user@PE1# set ospf traffic-engineering
   user@PE1# set ospf area 0.0.0.0 interface ge-1/2/0.1
   user@PE1# set ospf area 0.0.0.0 interface ge-1/2/1.5
   user@PE1# set ospf area 0.0.0.0 interface lo0.2 passive
   ```

3. Configure BGP.

The **auto-discovery-only** statement notifies the routing process (rpd) to expect autodiscovery-related NLRI messages so that information can be deciphered and used by LDP and VPLS. The **auto-discovery-only** statement must be configured on all PE routers in a VPLS. If you configure route reflection, the **auto-discovery-only** statement is also required on provider (P) routers that act as the route reflector in supporting FEC 129-related updates.

For interoperation scenarios in which a PE router must support both types of NLRI (FEC 128 and FEC 129), this example also includes the **signaling** statement.

```
[edit protocols bgp]
user@PE1# set local-address 1.1.1.2
user@PE1# set group pe-pe type internal
user@PE1# set group pe-pe family l2vpn auto-discovery-only
user@PE1# set group pe-pe family l2vpn signaling
user@PE1# set group pe-pe neighbor 1.1.1.3
user@PE1# set group pe-pe neighbor 1.1.1.4
user@PE1# set group pe-pe neighbor 1.1.1.5
```

4. Configure the routing instance.

   Both CE-facing interfaces are included in the routing instance. Only the multihomed interface is included in the multihoming site.

   As a convention, the route distinguisher is composed of Device PE1's loopback interface address and the multihoming site identifier.

```
[edit routing-instances green]
user@PE1# set instance-type vpls
user@PE1# set interface ge-0/3/1.0
user@PE1# set interface ge-0/3/3.0
user@PE1# set route-distinguisher 1.1.1.2:1
user@PE1# set l2vpn-id l2vpn-id:100:100
user@PE1# set vrf-target target:100:100
user@PE1# set protocols vpls no-tunnel-services
user@PE1# set protocols vpls multi-homing site test identifier 1
user@PE1# set protocols vpls multi-homing site test interface ge-0/3/1.0
```

5. (Optional) Configure bidirectional forwarding detection (BFD) for FEC 129 VPLS.

```
[edit routing-instances green]
user@PE1# set protocols vpls oam ping-interval 600
user@PE1# set protocols vpls oam bfd-liveness-detection minimum-interval 200
```

6. Configure the autonomous system (AS) number and router ID.

```
[edit routing-options]
user@PE1# set router-id 1.1.1.2
user@PE1# set autonomous-system 100
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
ge-0/3/1 {
```

```
  encapsulation ethernet-vpls;
  unit 0 {
    description PE1-to-CE1;
    family vpls;
  }
}
ge-0/3/3 {
  encapsulation ethernet-vpls;
  unit 0 {
    description PE1-to-CE2;
    family vpls;
  }
}
ge-1/2/0 {
  unit 1 {
    description PE1-to-P;
    family inet {
      address 10.1.1.1/30;
    }
    family mpls;
  }
}
ge-1/2/1 {
  unit 5 {
    description PE1-to-PE2;
    family inet {
      address 10.1.1.5/30;
    }
    family mpls;
  }
}
lo0 {
  unit 2 {
    family inet {
      address 1.1.1.2/32;
    }
  }
}
```

```
user@PE1# show protocols
mpls {
  interface ge-1/2/0.1;
  interface ge-1/2/1.5;
}
bgp {
  local-address 1.1.1.2;
  group pe-pe {
    type internal;
    family l2vpn {
      auto-discovery-only;
      signaling;
    }
    neighbor 1.1.1.3;
    neighbor 1.1.1.4;
    neighbor 1.1.1.5;
  }
```

```
}
ospf {
   traffic-engineering;
   area 0.0.0.0 {
      interface ge-1/2/0.1;
      interface ge-1/2/1.5;
      interface lo0.2 {
         passive;
      }
   }
}
ldp {
   interface ge-1/2/0.1;
   interface ge-1/2/1.5;
   interface lo0.2;
}

user@PE1# show routing-instances
green {
   instance-type vpls;
   interface ge-0/3/1.0;
   interface ge-0/3/3.0;
   route-distinguisher 1.1.1.2:100;
   l2vpn-id l2vpn-id:100:100;
   vrf-target target:100:100;
   protocols {
      vpls {
         no-tunnel-services;
         oam {
            ping-interval 600;
            bfd-liveness-detection {
               minimum-interval 200;
            }
         }
         multi-homing {
            site test {
               identifier 1;
               interface ge-0/3/1.0;
            }
         }
      }
   }
}

user@PE1# show routing-options
router-id 1.1.1.2;
autonomous-system 100;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying That Multihoming Is Operational on page 13
- Checking the Multihoming Routes on page 14

- Checking the BFD Sessions on page 15
- Pinging the Remote PE Router in the VPLS Domain on page 15

Verifying That Multihoming Is Operational

Purpose Verify that multihoming is operational.

Action From operational mode, enter the **show vpls connections extensive** command.

```
user@PE1> show vpls connections extensive
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not present
CM -- control-word mismatch      -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range               Up -- operational
OL -- no outgoing label          Dn -- down
LD -- local site signaled down   CF -- call admission control failure
RD -- remote site signaled down  SC -- local and remote site ID collision
LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch               MI -- Mesh-Group ID not available
BK -- Backup connection          ST -- Standby connection
PF -- Profile parse failure      PB -- Profile busy
RS -- remote site standby        SN -- Static Neighbor
LB -- Local site not best-site   RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: green
  L2vpn-id: 100:100
  Local-id: 1.1.1.2
    Number of local interfaces: 2
    Number of local interfaces up: 2
    ge-0/3/1.0
    ge-0/3/3.0
    lsi.101711873                  Intf - vpls green local-id 1.1.1.2 remote-id
1.1.1.4 neighbor 1.1.1.4
    Remote-id                 Type  St     Time last up          # Up trans
    1.1.1.4                   rmt   Up     Jan 31 13:49:52 2012           1
      Remote PE: 1.1.1.4, Negotiated control-word: No
      Incoming label: 262146, Outgoing label: 262146
      Local interface: lsi.101711873, Status: Up, Encapsulation: ETHERNET
       Description: Intf - vpls green local-id 1.1.1.2 remote-id 1.1.1.4 neighbor
 1.1.1.4
    Connection History:
        Jan 31 13:49:52 2012  status update timer
        Jan 31 13:49:52 2012  PE route changed
        Jan 31 13:49:52 2012  Out lbl Update                    262146
        Jan 31 13:49:52 2012  In lbl Update                     262146
        Jan 31 13:49:52 2012  loc intf up               lsi.101711873
  Multi-home:
    Local-site                Id     Pref   State
```

```
test                      1      100    Up
  Number of interfaces: 1
  Number of interfaces up: 1
    ge-0/3/1.0
  Received multi-homing advertisements:
    Remote-PE         Pref   flag   Description
    1.1.1.4           100    0x0
```

Meaning The output shows the status of multihoming for routing instance green.

Checking the Multihoming Routes

Purpose Verify that the expected routes are identified as multihoming.

Action From operational mode, enter the **show route table bgp.l2vpn.0** and **show route table green.l2vpn.0** commands.

```
user@PE1> show route table bgp.l2vpn.0
bgp.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.4:100:1.1.1.4/96 AD
                   *[BGP/170] 1d 03:10:45, localpref 100, from 1.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
1.1.1.4:100:1:0/96 MH
                   *[BGP/170] 1d 03:10:45, localpref 100, from 1.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
```

```
user@PE1> show route table green.l2vpn.0
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
1.1.1.4:100:1.1.1.4/96 AD
                   *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
1.1.1.4:100:1:0/96 MH
                   *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
                   *[VPLS/7] 1d 03:11:02, metric2 1
                    >    via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
                   *[LDP/9] 1d 03:11:02
                         Discard
```

Meaning **MH** in the output indicates a multihoming route. **AD** indicates autodiscovery.

Checking the BFD Sessions

Purpose Verify that the BFD session status is operational.

Action From operational mode, enter the **show bfd session** command.

```
user@PE1> show bfd session
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
127.0.0.1                Up        ge-1/2/1.5     0.600     0.200        3
127.0.0.1                Up        ge-1/2/0.1     0.600     0.200        3

2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```

Meaning **Up** in the **State** field indicates that BFD is working.

Pinging the Remote PE Router in the VPLS Domain

Purpose Check the operability of the MPLS Layer 2 virtual private network (VPN) connection.

Action From operational mode, enter the **ping mpls l2vpn** command with the **fec129** option.

```
user@PE1> ping mpls l2vpn fec129 instance green remote-id 1.1.1.5 remote-pe-address 1.1.1.5
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Meaning The output shows that the ping operation is successful, meaning that the LSP for a FEC 129 Layer 2 VPN connection is reachable.

Related Documentation

- VPLS Multihoming Overview on page 2
- Advantages of Using Autodiscovery for VPLS Multihoming on page 1