

Network Configuration Example

Load Balancing Layer 3 VPN Traffic While
Simultaneously Using IP Header Filtering

Release
13.1



Published: 2013-02-08

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 Load Balancing Layer 3 VPN Traffic While Simultaneously Using IP Header Filtering

Release 13.1

NCE0053

Copyright © 2013, 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
Layer 3 VPN Load Balancing Use Cases	1
Layer 3 VPN Load Balancing Overview	1
Example: Load Balancing Layer 3 VPN Traffic While Simultaneously Using IP Header Filtering	2

Introduction

This document describes how load balancing in a Layer 3 VPN (with internal and external BGP paths) can be configured while simultaneously using IP header filtering.

Layer 3 VPN Load Balancing Use Cases

Load balancing is useful for enhancing network utilization and performance. A load-balanced network provides high availability of critical TCP/IP-based services, such as the Internet and virtual private networking (VPN). Load balancing also ensures detection of device failures and automatic redistribution of traffic to surviving devices in the network.

Critical networks that are required to run at all times need to handle large volumes of client requests with minimal or no delays. Load balancing is essential to support critical applications such as financial transactions, database access, and corporate intranets. In situations where a device failure in a network threatens to disrupt network services, load balancing should be configured.

In a Layer 3 VPN network, a device learns multiple routes to a specific destination through multiple routing protocols and installs the route with the best route preference (also known as the administrative distance value) in its routing table. If multiple routes are received through the same protocol and have the same route preference, the route with the lowest cost (or metric) to the destination is installed in the routing table. If multiple routes are received through a single protocol having the same route preference and cost to a destination, load balancing is required.

The load balancing configured in this example is protocol-independent and allows the forwarding next hops of both the active and alternative routes to be used for load balancing. The type of load balancing configured is known as per-packet load balancing, which ensures equal traffic across all links. Per-packet load balancing avoids overloading of traffic and improves path utilization. To avoid routing loops occurring from traffic exiting the MPLS core re-entering to the core, traffic is filtered by using a VRF label.

Related Documentation

- [Layer 3 VPN Load Balancing Overview on page 1](#)
- [Example: Load Balancing Layer 3 VPN Traffic While Simultaneously Using IP Header Filtering on page 2](#)

Layer 3 VPN Load Balancing Overview

The load balancing feature allows a device to divide incoming and outgoing traffic along multiple paths in order to reduce congestion in the network. Load balancing improves the utilization of various network paths, and provides more effective network bandwidth.

When multiple protocols are in use, the device uses the route preference value (also known as the administrative distance value) to select a route. While using a single routing protocol, the router chooses the path with the lowest cost (or metric) to the destination.

If the device receives and installs multiple paths with the same route preference and same cost to a destination, load balancing must be configured.

In a network with both internal and external BGP paths installed among devices in different autonomous systems, BGP selects only a single best path by default, and does not perform load balancing. A Layer 3 VPN with internal and external BGP paths uses the **multipath** statement for protocol-independent load balancing. When you include the **multipath** statement in a routing instance, protocol-independent load balancing is applied to the default routing table for that routing instance. By using the **vpn-unequal-cost** statement, protocol-independent load balancing is applied to VPN routes. By using the **equal-external-internal** statement, protocol-independent load balancing is applied to both internal and external BGP paths and can be configured in conjunction with IP header filtering (enabled with the **vrf-table-label** statement).

**Related
Documentation**

- [Layer 3 VPN Load Balancing Use Cases on page 1](#)
- [Example: Load Balancing Layer 3 VPN Traffic While Simultaneously Using IP Header Filtering on page 2](#)

Example: Load Balancing Layer 3 VPN Traffic While Simultaneously Using IP Header Filtering

This example shows how to configure load balancing in a Layer 3 VPN (with internal and external BGP paths) while simultaneously using IP header filtering.

- [Requirements on page 2](#)
- [Overview on page 2](#)
- [Configuration on page 5](#)
- [Verification on page 13](#)

Requirements

This example requires the following hardware and software components:

- M Series Multiservice Edge Routers (M120 and M320 only), MX Series 3D Universal Edge Routers, T Series Core Routers, or PTX Series Transport Switches.
- Junos OS Release 12.1 or later

Overview

The following example shows how to configure load balancing while simultaneously using IP header filtering in a Layer 3 VPN.



NOTE: This example demonstrates how load balancing and IP header filtering work together. The testing of IP header filtering is out of the scope of this example.

The Junos OS BGP provides a multipath feature that allows load balancing between peers in the same or different autonomous systems (ASs). This example uses the **equal-external-internal** statement at the **[edit routing-instances instance-name routing-options multipath vpn-unequal-cost]** hierarchy level to perform load balancing. The **vrf-table-label** statement is configured at the **[edit routing-instances instance-name]** hierarchy level to enable IP header filtering.

```
[edit]
routing-instances {
  instance-name {
    vrf-table-label;
    routing-options {
      multipath {
        vpn-unequal-cost {
          equal-external-internal;
        }
      }
    }
  }
}
```



NOTE: These statements are available only in the context of a routing instance.

In this example, Device CE1 is in AS1 and connected to Device PE1. Devices PE1, PE2, PE3, and P are in AS2. Device CE2 is connected to Devices PE2 and PE3 and is in AS3. Device CE3 is connected to Device PE3 and is in AS4. BGP and MPLS are configured through the network. OSPF is the interior gateway protocol (IGP) that is used in this network.

The configuration for Devices PE1, PE2, and PE3 includes the **equal-external-internal** statement at the **[edit routing-instances instance-name routing-options multipath vpn-unequal-cost]** hierarchy level to enable load balancing in the network. IP header filtering is enabled when the **vrf-table-label** statement is configured at the **[edit routing-instances instance-name]** hierarchy level on the PE devices.

[Figure 1 on page 4](#) shows the topology used in this example.

Figure 1: Layer 3 VPN Load Balancing Using IP Header Filtering




Table 1 on page 4 shows the list of IP addresses used in this example for quick reference.

Table 1: Device IP Address Quick Reference

Device	AS	Device ID	Device Interface Units	Device Interface Unit IPs
CE1	1	1.1.1/32	Unit 1	10.1.1/30
PE1	2	1.1.1.2/32	Unit 2	10.1.1.2/30
			Unit 5	10.1.2.5/30
			Unit 9	10.1.3.9/30
PE2	2	1.1.1.3/32	Unit 6	10.1.2.6/30
			Unit 13	10.1.4.13/30
			Unit 21	10.1.6.21/30
PE3	2	1.1.1.4/32	Unit 10	10.1.3.10/30
			Unit 18	10.1.5.18/30
			Unit 25	10.1.7.25/30
			Unit 29	10.1.8.29/30

Table 1: Device IP Address Quick Reference (*continued*)

Device	AS	Device ID	Device Interface Units	Device Interface Unit IPs
P	2	1.1.1.5/32	Unit 14	10.1.4.14/30
			Unit 17	10.1.5.17/30
CE2	3	1.1.1.6/32	Unit 22	10.1.6.22/30
			Unit 26	10.1.7.26/30
CE3	4	1.1.1.7/32	Unit 30	10.1.8.30/30



NOTE: This example was tested using logical systems (logical routers). Therefore all the physical interfaces in the example are the same and the configuration is done on separate logical interfaces. In an non-test network, you will use separate physical routers and separate physical interfaces for the connections to other devices.

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 CE1

```
set interfaces ge-2/1/10 unit 1 family inet address 10.1.1.1/30
set interfaces ge-2/1/10 unit 1 family mpls
set interfaces ge-2/1/10 unit 1 description toPE1
set interfaces lo0 unit 4 family inet address 1.1.1.1/32
set routing-options router-id 1.1.1.1
set routing-options autonomous-system 1
set protocols bgp group toPE1 type external
set protocols bgp group toPE1 export send-direct
set protocols bgp group toPE1 peer-as 2
set protocols bgp group toPE1 neighbor 10.1.1.2
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
```

Device PE1

```
set interfaces ge-2/1/10 unit 2 family inet address 10.1.1.2/30
set interfaces ge-2/1/10 unit 2 family mpls
set interfaces ge-2/1/10 unit 2 description toCE1
set interfaces ge-2/1/10 unit 5 family inet address 10.1.2.5/30
set interfaces ge-2/1/10 unit 5 family mpls
set interfaces ge-2/1/10 unit 5 description toPE2
set interfaces ge-2/1/10 unit 9 family inet address 10.1.3.9/30
```

```
set interfaces ge-2/1/10 unit 9 family mpls
set interfaces ge-2/1/10 unit 9 description toPE3
set interfaces lo0 unit 5 family inet address 1.1.1.2/32
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.5 passive
set protocols ospf area 0.0.0.0 interface ge-2/1/10.5 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/1/10.9 metric 10
set protocols bgp group toInternal type internal
set protocols bgp group toInternal family inet-vpn unicast
set protocols bgp group toInternal local-address 1.1.1.2
set protocols bgp group toInternal neighbor 1.1.1.3
set protocols bgp group toInternal neighbor 1.1.1.4
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 2
set routing-options forwarding-table export lb
set routing-instances purple instance-type vrf
set routing-instances purple interface ge-2/1/10.2
set routing-instances purple route-distinguisher 2:1
set routing-instances purple vrf-target target:2:1
set routing-instances purple vrf-table-label
set routing-instances purple protocols bgp group toCE1 type external
set routing-instances purple protocols bgp group toCE1 peer-as 1
set routing-instances purple protocols bgp group toCE1 neighbor 10.1.1.1
set routing-instances purple routing-options multipath vpn-unequal-cost
    equal-external-internal
set policy-options policy-statement lb then load-balance per-packet
```

Device PE2

```
set interfaces ge-2/1/10 unit 6 family inet address 10.1.2.6/30
set interfaces ge-2/1/10 unit 6 family mpls
set interfaces ge-2/1/10 unit 6 description toPE1
set interfaces ge-2/1/10 unit 13 family inet address 10.1.4.13/30
set interfaces ge-2/1/10 unit 13 family mpls
set interfaces ge-2/1/10 unit 13 description toP
set interfaces ge-2/1/10 unit 21 family inet address 10.1.6.21/30
set interfaces ge-2/1/10 unit 21 family mpls
set interfaces ge-2/1/10 unit 21 description toCE2
set interfaces lo0 unit 6 family inet address 1.1.1.3/32
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.6 passive
set protocols ospf area 0.0.0.0 interface ge-2/1/10.6 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/1/10.13 metric 5
set protocols bgp group toInternal type internal
set protocols bgp group toInternal family inet-vpn unicast
set protocols bgp group toInternal local-address 1.1.1.3
set protocols bgp group toInternal neighbor 1.1.1.2
set protocols bgp group toInternal neighbor 1.1.1.4
set routing-options router-id 1.1.1.3
set routing-options autonomous-system 2
set routing-options forwarding-table export lb
set routing-instances purple instance-type vrf
set routing-instances purple interface ge-2/1/10.21
set routing-instances purple route-distinguisher 2:1
```

```
set routing-instances purple vrf-target target:2:1
set routing-instances purple vrf-table-label
set routing-instances purple protocols bgp group toCE2 type external
set routing-instances purple protocols bgp group toCE2 peer-as 3
set routing-instances purple protocols bgp group toCE2 neighbor 10.1.6.22
set routing-instances purple routing-options multipath vpn-unequal-cost
  equal-external-internal
set policy-options policy-statement lb then load-balance per-packet
```

Device PE3

```
set interfaces ge-2/1/10 unit 10 family inet address 10.1.3.10/30
set interfaces ge-2/1/10 unit 10 family mpls
set interfaces ge-2/1/10 unit 10 description toPE1
set interfaces ge-2/1/10 unit 18 family inet address 10.1.5.18/30
set interfaces ge-2/1/10 unit 18 family mpls
set interfaces ge-2/1/10 unit 18 description toP
set interfaces ge-2/1/10 unit 25 family inet address 10.1.7.25/30
set interfaces ge-2/1/10 unit 25 family mpls
set interfaces ge-2/1/10 unit 25 description toCE2
set interfaces ge-2/1/10 unit 29 family inet address 10.1.8.29/30
set interfaces ge-2/1/10 unit 29 family mpls
set interfaces ge-2/1/10 unit 29 description toCE3
set interfaces lo0 unit 7 family inet address 1.1.1.4/32
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.7 passive
set protocols ospf area 0.0.0.0 interface ge-2/1/10.10 metric 10
set protocols ospf area 0.0.0.0 interface ge-2/1/10.18 metric 5
set protocols bgp group toInternal type internal
set protocols bgp group toInternal local-address 1.1.1.4
set protocols bgp group toInternal family inet-vpn unicast
set protocols bgp group toInternal family route-target
set protocols bgp group toInternal neighbor 1.1.1.2
set protocols bgp group toInternal neighbor 1.1.1.3
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 2
set routing-options forwarding-table export lb
set routing-instances purple instance-type vrf
set routing-instances purple interface ge-2/1/10.25
set routing-instances purple interface ge-2/1/10.29
set routing-instances purple route-distinguisher 2:1
set routing-instances purple vrf-target target:2:1
set routing-instances purple vrf-table-label
set routing-instances purple protocols bgp group toCE2 type external
set routing-instances purple protocols bgp group toCE2 peer-as 3
set routing-instances purple protocols bgp group toCE2 neighbor 10.1.7.26
set routing-instances purple protocols bgp group toCE3 type external
set routing-instances purple protocols bgp group toCE3 peer-as 4
set routing-instances purple protocols bgp group toCE3 neighbor 10.1.8.30
set routing-instances purple routing-options multipath vpn-unequal-cost
  equal-external-internal
set policy-options policy-statement lb then load-balance per-packet
```

Device P

```
set interfaces ge-2/1/10 unit 14 family inet address 10.1.4.14/30
```

```
set interfaces ge-2/1/10 unit 14 family mpls
set interfaces ge-2/1/10 unit 14 description toPE2
set interfaces ge-2/1/10 unit 17 family inet address 10.1.5.17/30
set interfaces ge-2/1/10 unit 17 family mpls
set interfaces ge-2/1/10 unit 17 description toPE3
set interfaces lo0 unit 8 family inet address 1.1.1.5/32
set protocols mpls interface all
set protocols ldp interface all
set protocols ospf area 0.0.0.0 interface lo0.8 passive
set protocols ospf area 0.0.0.0 interface ge-2/1/10.14 metric 5
set protocols ospf area 0.0.0.0 interface ge-2/1/10.17 metric 5
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 2
```

Device CE2

```
set interfaces ge-2/1/10 unit 22 family inet address 10.1.6.22/30
set interfaces ge-2/1/10 unit 22 family mpls
set interfaces ge-2/1/10 unit 22 description toPE2
set interfaces ge-2/1/10 unit 26 family inet address 10.1.7.26/30
set interfaces ge-2/1/10 unit 26 family mpls
set interfaces ge-2/1/10 unit 26 description toPE3
set interfaces lo0 unit 6 family inet address 1.1.1.6/32
set routing-options router-id 1.1.1.6
set routing-options autonomous-system 3
set protocols bgp group toAS2 type internal
set protocols bgp group toAS2 export send-direct
set protocols bgp group toAS2 peer-as 2
set protocols bgp group toAS2 neighbor 10.1.6.21
set protocols bgp group toAS2 neighbor 10.1.7.25
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
```

Device CE3

```
set interfaces ge-2/1/10 unit 30 family inet address 10.1.8.30/30
set interfaces ge-2/1/10 unit 30 family mpls
set interfaces ge-2/1/10 unit 30 description toPE3
set interfaces lo0 unit 7 family inet address 1.1.1.7/32
set routing-options router-id 1.1.1.7
set routing-options autonomous-system 4
set protocols bgp group toPE3 type internal
set protocols bgp group toPE3 export send-direct
set protocols bgp group toPE3 peer-as 2
set protocols bgp group toPE3 neighbor 10.1.8.29
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct then accept
```

Step-by-Step Procedure

The following example requires that you 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 unequal-cost load balancing across the VPN setup:

1. Configure the router ID on Device CE1, and assign the device to its autonomous system.

```
[edit routing-options]
user@CE1# set routing-options router-id 1.1.1.1
user@CE1# set routing-options autonomous-system 1
```

Similarly, configure all other devices.

2. Configure BGP groups for traffic through the entire network.
 - a. Configure the BGP group for traffic to and from the MPLS network (CE devices).

```
[edit protocols bgp group toPE1]
user@CE1# set type external
user@CE1# set peer-as 2
user@CE1# set neighbor 10.1.1.2
```

- b. Configure similar BGP groups (**toAS2** and **toPE3**) on Devices CE2 and CE3 by modifying the **peer-as** and **neighbor** statements accordingly.
 - c. Configure the BGP group for traffic through the MPLS network (PE devices).

```
[edit protocols bgp group toInternal]
user@PE1# set type internal
user@PE1# set family inet-vpn unicast
user@PE1# set local-address 1.1.1.2
user@PE1# set neighbor 1.1.1.3
user@PE1# set neighbor 1.1.1.4
```

- d. Configure the same BGP group (**toInternal**) on Devices PE2 and PE3 by modifying the **local-address** and **neighbor** statements accordingly.
3. Configure a routing policy for exporting routes to and from the MPLS network (**send-direct** policy) and a policy for load balancing traffic network across the MPLS network (**lb** policy).
 - a. Configure a policy (**send-direct**) for exporting routes from the routing table into BGP on Device CE1.

```
[edit policy-options policy-statement send-direct]
user@CE1# set from protocol direct
user@CE1# set then accept

[edit protocols bgp group toPE1]
user@CE1# set export send-direct
```

Similarly, configure the **send-direct** policy on Devices CE2 and CE3.

- b. Configure a policy (**lb**) for exporting routes from the routing table into the forwarding table on Device PE1.

The **lb** policy configures per-packet load balancing, which ensures that all next-hop addresses for a destination are installed in the forwarding table.

```
[edit policy-options policy-statement lb]
user@PE1# set then load-balance per-packet

[edit routing-options]
user@PE1# set forwarding-table export lb
```

Similarly, configure the **lb** policy on Devices PE2, and PE3.

4. Configure the following:

- a. Configure the routing instance on the PE devices for exporting routes through the autonomous systems.
- b. Include the **equal-external-internal** statement at the **[edit routing-instances instance-name routing-options multipath vpn-unequal-cost]** hierarchy level to enable load balancing in the network.
- c. Include the **vrf-table-label** statement at the **[edit routing-instances instance-name]** hierarchy level for filtering traffic prior to exiting the egress device (Device CE3).

Device PE1

```
[edit routing-instances purple]
user@PE1# set instance-type vrf
user@PE1# set interface ge-2/1/10.2
user@PE1# set route-distinguisher 2:1
user@PE1# set vrf-target target:2:1
user@PE1# set vrf-table-label
user@PE1# set protocols bgp group toCE1 type external
user@PE1# set protocols bgp group toCE1 peer-as 1
user@PE1# set protocols bgp group toCE1 neighbor 10.1.1.1
user@PE1# set routing-options multipath vpn-unequal-cost equal-external-internal
```

Device PE2

```
[edit routing-instances purple]
user@PE2# set instance-type vrf
user@PE2# set interface ge-2/1/10.21
user@PE2# set route-distinguisher 2:1
user@PE2# set vrf-target target:2:1
user@PE2# set vrf-table-label
user@PE2# set protocols bgp group toCE2 type external
user@PE2# set protocols bgp group toCE2 peer-as 3
user@PE2# set protocols bgp group toCE2 neighbor 10.1.6.22
user@PE2# set routing-options multipath vpn-unequal-cost equal-external-internal
```

Device PE3

```
[edit routing-instances purple]
user@PE3# set instance-type vrf
user@PE3# set interface ge-2/1/10.25
user@PE3# set interface ge-2/1/10.29
user@PE3# set route-distinguisher 2:1
user@PE3# set vrf-target target:2:1
user@PE3# set vrf-table-label
user@PE3# set protocols bgp group toCE2 type external
user@PE3# set protocols bgp group toCE2 peer-as 3
user@PE3# set protocols bgp group toCE2 neighbor 10.1.7.26
user@PE3# set protocols bgp group toCE3 type external
user@PE3# set protocols bgp group toCE3 peer-as 4
user@PE3# set protocols bgp group toCE3 neighbor 10.1.8.30
user@PE3# set routing-options multipath vpn-unequal-cost equal-external-internal
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-options**, and **show routing-instances**

commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE3# show interfaces
```

```
ge-2/1/10 {  
  unit 10 {  
    description toPE1;  
    family inet {  
      address 10.1.3.10/30;  
    }  
    family mpls  
  }  
  unit 18 {  
    description toP;  
    family inet {  
      address 10.1.5.18/30;  
    }  
    family mpls  
  }  
  unit 25 {  
    description toCE2;  
    family inet {  
      address 10.1.7.25/30;  
    }  
    family mpls  
  }  
  unit 29 {  
    description toCE3;  
    family inet {  
      address 10.1.8.29/30;  
    }  
    family mpls  
  }  
}  
lo0 {  
  unit 7 {  
    family inet {  
      address 1.1.1.4/32;  
    }  
  }  
}
```

```
user@PE3# show protocols
```

```
mpls {  
  interface all;  
}  
bgp {  
  group toInternal {  
    type internal;  
    local-address 1.1.1.4;  
    family inet {  
      unicast;  
    }  
    family inet-vpn {  
      unicast;  
    }  
  }  
}
```

```
        family route-target;
        neighbor 1.1.1.2;
        neighbor 1.1.1.3;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.7 {
            passive;
        }
        interface ge-2/1/10.10 {
            metric 10;
        }
        interface ge-2/1/10.18 {
            metric 5;
        }
    }
}
ldp {
    interface all;
}

user@PE3# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}

user@PE3# show routing-instances
purple {
    instance-type vrf;
    interface ge-2/1/10.25;
    interface ge-2/1/10.29;
    route-distinguisher 2:1;
    vrf-target target:2:1;
    vrf-table-label;
    routing-options {
        multipath {
            vpn-unequal-cost equal-external-internal;
        }
    }
    protocols {
        bgp {
            group toCE2 {
                type external;
                peer-as 3;
                neighbor 10.1.7.26;
            }
            group toCE3 {
                type external;
                peer-as 4;
                neighbor 10.1.8.30;
            }
        }
    }
}
```



```

user@PE3# show routing-options
router-id 1.1.1.4;
autonomous-system 2;
forwarding-table {
    export lb;
}

```

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

Verification

Confirm that the configuration is working properly.

- [Verifying BGP on page 13](#)
- [Verifying Load Balancing on page 14](#)
- [Verifying Load Balancing While Using IP Header Filtering on page 16](#)

Verifying BGP

Purpose Verify that BGP is working.

Action From operational mode, run the **show route protocol bgp** command.

```

user@PE3> show route protocol bgp

inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

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

purple.inet.0: 9 destinations, 14 routes (9 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32      *[BGP/170] 04:47:14, localpref 100, from 1.1.1.2
                AS path: 1 I
                > to 10.1.3.9 via ge-2/1/10.10, Push 16
1.1.1.6/32      @[BGP/170] 00:13:28, localpref 100
                AS path: 3 I
                > to 10.1.7.26 via ge-2/1/10.25
                [BGP/170] 00:10:36, localpref 100, from 1.1.1.3
                AS path: 3 I
                > to 10.1.5.17 via ge-2/1/10.18, Push 16, Push 299776(top)
1.1.1.7/32      *[BGP/170] 00:10:56, localpref 100
                AS path: 4 I
                > to 10.1.8.30 via ge-2/1/10.29
10.1.1.0/30     *[BGP/170] 04:47:14, localpref 100, from 1.1.1.2
                AS path: I
                > to 10.1.3.9 via ge-2/1/10.10, Push 16
10.1.6.20/30    *[BGP/170] 04:47:03, localpref 100, from 1.1.1.3
                AS path: I
                > to 10.1.5.17 via ge-2/1/10.18, Push 16, Push 299776(top)
                [BGP/170] 00:13:28, localpref 100
                AS path: 3 I
                > to 10.1.7.26 via ge-2/1/10.25
10.1.7.24/30    [BGP/170] 00:13:28, localpref 100
                AS path: 3 I
                > to 10.1.7.26 via ge-2/1/10.25
10.1.8.28/30    [BGP/170] 00:10:56, localpref 100

```

```

AS path: 4 I
> to 10.1.8.30 via ge-2/1/10.29

mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:1:1.1.1.1/32
*[BGP/170] 04:47:14, localpref 100, from 1.1.1.2
AS path: 1 I
> to 10.1.3.9 via ge-2/1/10.10, Push 16

2:1:1.1.1.6/32
*[BGP/170] 00:10:36, localpref 100, from 1.1.1.3
AS path: 3 I
> to 10.1.5.17 via ge-2/1/10.18, Push 16, Push 299776(top)

2:1:10.1.1.0/30
*[BGP/170] 04:47:14, localpref 100, from 1.1.1.2
AS path: I
> to 10.1.3.9 via ge-2/1/10.10, Push 16

2:1:10.1.6.20/30
*[BGP/170] 04:47:03, localpref 100, from 1.1.1.3

```

The output lists the BGP routes installed into the routing table. The lines of output that start with **1.1.1.1/32**, **10.1.1.0/30**, and **2:1:1.1.1/32** show the BGP routes to Device CE1, which is in AS1. The lines of output that start with **1.1.1.6/32**, **2:1:1.1.6/32**, and **2:1:10.1.6.20/30** show the BGP routes to Device CE2, which is in AS3. The line of output that starts with **1.1.1.7/32** shows the BGP route to Device CE3, which is in AS4.

Meaning BGP is functional in the network.

Verifying Load Balancing

Purpose Verify that forwarding is taking place in both directions by checking:

- If both next hops are installed in the forwarding table for a route.
- If external BGP routes are installed in the forwarding table for a route.

Action From operational mode, run the **show route forwarding-table** and **show route forwarding-table destination <destination IP>** commands.

```
user@PE3> show route forwarding-table
```

```

Router: PE3
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
1.1.1.2/32       user  1 10.1.3.9      ucst  999   8 ge-2/1/10.10
1.1.1.3/32       user  1 10.1.5.17     ucst  1243  12 ge-2/1/10.18
1.1.1.4/32       intf  0 1.1.1.4       locl  895   1
1.1.1.5/32       user  1 10.1.5.17     ucst  1243  12 ge-2/1/10.18
10.1.2.4/30      user  0
                  10.1.3.9      ucst  999   8 ge-2/1/10.10
                  10.1.5.17     ucst  1243  12 ge-2/1/10.18

```

```

10.1.3.8/30      intf    0                rslv   899    1 ge-2/1/10.10
10.1.3.8/32      dest    0 10.1.3.8       recv   897    1 ge-2/1/10.10
10.1.3.9/32      dest    0 0.6.80.3.0.21.59.d.c5.d9.0.21.59.d.c5.da.8.0
                                         ucst   999    8 ge-2/1/10.10
10.1.3.10/32     intf    0 10.1.3.10       locl   898    2
10.1.3.10/32     dest    0 10.1.3.10       locl   898    2
10.1.3.11/32     dest    0 10.1.3.11       bcst   896    1 ge-2/1/10.10
10.1.4.12/30     user    0 10.1.5.17       ucst  1243   12 ge-2/1/10.18
10.1.5.16/30     intf    0                rslv   903    1 ge-2/1/10.18
10.1.5.16/32     dest    0 10.1.5.16       recv   901    1 ge-2/1/10.18
10.1.5.17/32     dest    0 0.e.80.3.0.21.59.d.c5.d9.0.21.59.d.c5.da.8.0
                                         ucst  1243   12 ge-2/1/10.18
10.1.5.18/32     intf    0 10.1.5.18       locl   902    2
10.1.5.18/32     dest    0 10.1.5.18       locl   902    2
10.1.5.19/32     dest    0 10.1.5.19       bcst   900    1 ge-2/1/10.18
224.0.0.0/4      perm    2                mdsc   592    1
224.0.0.1/32     perm    0 224.0.0.1       mcst   576    3
224.0.0.5/32     user    1 224.0.0.5       mcst   576    3
255.255.255.255/32 perm    0                bcst   577    1

```

Router: PE3

Routing table: __master.anon__.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	909	1	
0.0.0.0/32	perm	0		dscd	907	1	
224.0.0.0/4	perm	0		mdsc	908	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	904	1	
255.255.255.255/32	perm	0		bcst	905	1	

Router: PE3

Routing table: purple.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	918	1	
0.0.0.0/32	perm	0		dscd	916	1	
1.1.1.1/32	user	0		indr	1048576	3	
			10.1.3.9	Push	16	1187	2 ge-2/1/10.10
1.1.1.6/32	user	0		ulst	1048587	2	
			10.1.7.26	ucst	1239	4	ge-2/1/10.25
			10.1.5.17	indr	1048579	3	
				Push	16, Push	299776(top)	1306
2 ge-2/1/10.18							
1.1.1.7/32	user	0	10.1.8.30	ucst	1299	4	ge-2/1/10.29
299808(S=0)	user	0	10.1.5.17	Pop	1304	2	ge-2/1/10.18
...							

In the **default.inet** routing table, which is the forwarding table, the line of output that starts with **10.1.2.4/30** shows that for a route to Device PE2 in the same AS, two next hops are installed in the table: **10.1.3.9** and **10.1.5.17**.

In the **purple.inet** routing table, which is the external routing table, the line of output that starts with **1.1.1.6/32** shows that for a route to Device CE2 in AS3, an internal next hop of **10.1.5.17** and an external next hop of **10.1.7.26** are installed in the table. This indicates that both internal and external BGP routes are operational in the network.

user@PE3> show route forwarding-table destination 10.1.2.6

Router: PE3

Routing table: default.inet

```

Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
10.1.2.4/30      user  0          10.1.3.9          ucst  999    8 ge-2/1/10.10
                  10.1.5.17        ucst 1243   12 ge-2/1/10.18

```

```

Router: PE3
Routing table: __master.anon__.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0          rjct  909    1

```

```

Router: PE3
Routing table: purple.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0          rjct  918    1

```

The line of output that starts with **10.1.2.4/30** shows that for a route from Device PE3 to Device PE2 in the same AS, two next hops are installed in the table: **10.1.3.9** through the **ge-2/1/10.10** interface, and **10.1.5.17** through the **ge-2/1/10.18** interface.

Meaning Multiple next hops for a route, including external BGP routes, are installed in the forwarding tables.

Verifying Load Balancing While Using IP Header Filtering

Purpose Verify that filtered traffic reaches the egress CE devices after load balancing has been configured on the PE devices.

Action Configure a firewall filter on Device PE3 on the interface connecting to Device CE2.

```

[edit firewall family inet filter filterPE3 term a]
user@PE3# set from protocol tcp
user@PE3# set from source-port-except bgp
user@PE3# set from destination-port-except bgp
user@PE3# set then count filterPE3
user@PE3# set then accept

```

```

[edit firewall family inet filter filterPE3 term b]
user@PE3# set then accept

```

```

[edit interfaces ge-2/1/10 unit 25]
user@PE3# set family inet filter output filterPE3

```

Similarly, configure a firewall filter on Device PE3 on the interface facing Device CE3, and another on Device PE2 on the interface facing Device CE2.

Count the packets exiting the egress interfaces on Devices PE2 and PE3 by using the **show firewall filter <filter name> counter <counter name>** operational mode command. The output confirms if load balancing takes place with IP header filtering configured (enabled by the **vrf-table-label** statement). If all transmitted packets have been load-balanced between the paths PE3->CE2, PE3->CE3, and PE2->CE2, then it means that the IP header filtering feature works in a load-balanced Layer 3 network.

You can clear the counter by using the **clear firewall filter <filter name> counter <counter name>** operational mode command.

Meaning Load balancing takes place with IP header filtering configured.

Related Documentation

- Configuring Protocol-Independent Load Balancing in Layer 3 VPNs
- Example: Load Balancing BGP Traffic
- Load Balancing and IP Header Filtering for Layer 3 VPNs

