JNCIE-SP Certification Self-Study Bundle

Revision 1.b

Student Workbook and Lab Guide

JUNIPER NETWORKS | Education Services

1133 Innovation Way
Sunnyvale, CA 94089 USA
408-745-2000
www.juniper.net

Course Number: EDU-JNCIE-SP-BOOK

This document is produced by Juniper Networks, Inc.

This document or any part thereof may not be reproduced or transmitted in any form under penalty of law, without the prior written permission of Juniper Networks Education Services.

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.

JNCIE-SP Certification Self-Study Bundle, Revision 1.b

Copyright © 2020 Juniper Networks, Inc. All rights reserved.

Printed in USA.

Revision History:

Revision 1.a—February 2019
Revision 1.b—March 2020

The information in this document is current as of the date listed above, has been carefully verified, and is believed to be accurate as of this publication date. Juniper Networks assumes no responsibilities for any inaccuracies that may appear in this document. In no event will Juniper Networks be liable for direct, indirect, special, exemplary, incidental, or consequential damages resulting from any defect or omission in this document, even if advised of the possibility of such damages.

Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

YEAR 2000 NOTICE

Juniper Networks hardware and software products do not suffer from Year 2000 problems and hence are Year 2000 compliant. The Junos operating system has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

SOFTWARE LICENSE

The terms and conditions for using Juniper Networks software are described in the software license provided with the software, or to the extent applicable, in an agreement executed between you and Juniper Networks, or Juniper Networks agent. By using Juniper Networks software, you indicate that you understand and agree to be bound by its license terms and conditions. Generally speaking, the software license restricts the manner in which you are permitted to use the Juniper Networks software, may contain prohibitions against certain uses, and may state conditions under which the license is automatically terminated. You should consult the software license for further details.

Contents

Contents 3
General Information 9
About This Publication 9
Education Services Offerings 9
Technical Publications 9
Juniper Networks Support 9
Target Audience 10
Exam Strategy 10
Lab Preparation 10
Day Before the Exam 11
Exam Day 11
After the Exam 12
Workbook and Configuration Files 12
Chapter 1: General System Features 14
Task 1: Initial System Settings 17
Task 2: SNMP Configuration 19
Task 3: Firewall Filters 20
Task 4: Interface Configuration 21
Task 5: Scripting 24
Chapter 2: IGP Configuration and Troubleshooting 26

© 2020 Juniper Networks, Inc. All rights reserved.

Task 1: OSPF Troubleshooting 26
Task 2: IS-IS Troubleshooting 28
Task 3: IGP Rollout 32
Chapter 3: BGP and Routing Policy 36
Task 1: IBGP and Confederation 36
Task 2: EBGP Configuration 36
Task 3: Routing Policies 38
Task 4: IBGP and Route Reflection 40
Chapter 4: MPLS Configuration 42
Task 1: LDP Configuration 43
Task 2: RSVP Configuration 43
Task 3: RSVP Protection 48
Task 4: IPv6 Tunneling with 6PE 48
Chapter 5: Layer 3 VPN Configuration 49
Task 1: Layer 3 VPN Configuration 49
Task 2: Multicast in Layer 3 VPNs 53
Task 3: IPv6 Tunneling with 6VPE 54
Chapter 6: Layer 2 VPN Services 55
Task 1: Layer 2 VPN Configuration 55
Task 2: VPLS Configuration 56
Task 3: EVPN Configuration 59
Chapter 7: Inter-Provider VPN Configuration 61

© 2020 Juniper Networks, Inc. All rights reserved.

Task 1: Inter-Provider VPN Option B 61
Task 2: Inter-Provider VPN Option C 61
Chapter 8: Class of Service 63
Task 1: Forwarding Classes, Queues, and Schedulers 63
Task 2: Classification, Policing, and Marking 64
Super Lab 1 66
Task 1: Initial System Configuration 67
Task 2: Building the Network 69
Task 3: IGP Configuration 71
Task 4: BGP Configuration 74
Task 5: MPLS Configuration 75
Task 6: VPN Configuration 78
Task 7: Class of Service Configuration 81
Super Lab 2 84
Task 1: System Features 85
Task 2: Troubleshooting and Configuring IGP 87
Task 3: Troubleshooting and Configuring BGP 89
Task 4: MPLS Configuration 90
Task 5: IPv6 Configuration 91
Task 6: MPLS VPNs Configuration 92
Appendices 94
Appendix - Additional Theory 95

© 2020 Juniper Networks, Inc. All rights reserved.

OSPF Adjacency Troubleshooting 95
BGP Adjacency Troubleshooting 99
BGP IPV6 NLRI over IPV4 Peering 103
Troubleshooting: Multicast Traffic Engineering Using RIB-groups 111
Advanced Firewall Filtering 114
Appendix - Topology Diagrams 117
Appendix - Chapter 1: General System Features 134
Task 1: Initial System Configuration 136
Task 2: SNMP Configuration 151
Task 3: Firewall filters 156
Task 4: Interface Configuration 167
Task 5: Scripting 179
Appendix - Chapter 2: IGP Configuration and Troubleshooting 185
Task 1: OSPF Troubleshooting 185
Task 2: IS-IS Troubleshooting 252
Task 3: IGP Rollout 307
Appendix - Chapter 3: BGP and Routing policy 374
Task 1: Creating Clusters – Initial Setup 374
Task 2: EBGP Configuration 390
Task 3: Routing Policies 436
Task 4: IBGP and Route Reflection 486
Appendix - Chapter 4: MPLS Configuration 530

© 2020 Juniper Networks, Inc. All rights reserved.

MPLS Overview 530
LDP Overview 531
Task 1: LDP Configuration 532
Task 2: RSVP Configuration 553
Task 3: RSVP Protection 631
Task 4: IPv6 Tunneling with 6PE 661
Appendix - Chapter 5: L3VPN Configuration 671
Task 1: L3VPN Configuration 671
Task 2: Multicast in L3VPNs 729
Task 3: IPv6 Tunneling with 6VPE 755
Appendix - Chapter 6: Layer 2 VPN Services 764
Task 1: L2VPN Configuration 764
Task 2: VPLS Configuration 785
Task 3. EVPN Configuration 815
Appendix - Chapter 7: Inter-Provider VPN Configuration 823
Task 1: Inter-Provider VPN Option B 823
Task 2: Inter-Provider VPN Option C 834
Appendix - Chapter 8: Class of Service 848
Task 1: Forwarding Classes, Queues and Schedulers 848
Task 2: Classification, Policing and Marking 858
Appendix – Super Lab 1 869
Task 1: Initial System Configuration 869

© 2020 Juniper Networks, Inc. All rights reserved.

Task 2: Building the Network 876
Task 3: IGP Configuration 884
Task 4: BGP Configuration 934
Task 5: MPLS Configuration 987
Task 6: VPN Configuration 1071
Task 7: Class of Service Configuration 1128
Appendix – Super Lab 2 1146
Task 1: System Features 1148
Task 2: Troubleshooting and Configuring IGP 1168
Task 3: Troubleshooting and Configuring BGP 1224
Task 4: MPLS Configuration 1288
Task 5: IPv6 Configuration 1327
Task 6: MPLS VPNs Configuration 1336

© 2020 Juniper Networks, Inc. All rights reserved.

General Information

About This Publication

The JNCIE-SP lab exam workbook is designed for candidates to practice technologies covered in the JNCIE-SP lab exam. It also helps candidates practice skills relating to task prioritization, dependencies, and correlations.

This workbook is split into two parts. The first part of the guide outlines the high level tasks for each chapter and practice labs (referred to as Super Labs). The second part is the appendix, which contains the detailed explanations and configuration steps to complete each task.

The JNCIE-SP lab exam preparation workbook consists of 11 virtual MX-series devices running Junos version 18.3, and a Linux based student-desktop device.

Education Services Offerings

You can obtain information on the latest Education Services offerings, course dates, and class locations from the World Wide Web by pointing your Web browser to: http://www.juniper.net/training/education/.

Technical Publications

You can print technical manuals and release notes directly from the Internet in a variety of formats:

- Go to http://www.juniper.net/techpubs/.
- Locate the specific software or hardware release and title you need and choose the format in which you want to view or print the document.

Documentation sets and CDs are available through your local Juniper Networks sales office or account representative.

Juniper Networks Support

For technical support, contact Juniper Networks at http://www.juniper.net/customers/support/, or at 1-888-314-JTAC (within the United States) or 408-745-2121 (outside the United States).

© 2020 Juniper Networks, Inc. All rights reserved.

Target Audience

This workbook is developed for experienced network engineers who are preparing for the JNCIE-SP lab exam. Although not required, it is highly recommended that you have passed the JNCIS-SP and JNCIP-SP written exams before you start using this workbook. The *JNCIE-SP Certification Self-Study Bundle* is a hands-on guide to validate the skills needed to pass the official JNCIE-SP lab exam. This guide is based on official JNCIE-SP exam objectives and is targeted at JNCIP-SP certified engineers who are studying for the expert-level certification and need extra help preparing for the exam. Each chapter covers several technologies with expert-level configuration tasks and detailed answers. In this workbook you will find several technology introductions. However, do not expect a full explanation about OSPF, BGP, MPLS, etc. because there are other resources available to learn the prerequisite knowledge. This guide also contains two practice exams to give you a similar experience to the real JNCIE-SP exam.

Exam Strategy

Lab Preparation

To earn the JNCIE-SP credential you must first pass the following Juniper Networks written exams: JNCIA-Junos, JNCIS-SP, and JNCIP-SP. Finally, you must also pass the very challenging 8 hour JNCIE-SP lab exam. You must have extensive hands-on experience with service provider routing and switching on vMX platforms. When preparing for the JNCIE-SP lab exam ensure you fully understand the exam topics and objectives. The latest update for the JNCIE-SP lab exam topics and objectives can be found on the Juniper Networks certification website.

Use this workbook to verify if you are ready for the JNCIE-SP lab exam as it covers very challenging tasks specifically written for JNCIE-SP lab preparation. The workbook tasks are written at the same difficulty level as the JNCIE-SP lab exam and follows the same structure.

We recommend completing all the workbook's chapter and practice exam tasks at least twice. The first time, work through all the tasks and use the solutions as often as you like. The second time try to solve all the chapter and practice lab tasks without looking at the solutions.

While using this workbook, ensure that you take study notes and record the time spent on each task. Time is of the essence for the JNCIE-SP lab exam.

© 2020 Juniper Networks, Inc. All rights reserved.

Day Before the Exam

We advise you not to do extensive lab preparation anymore on this day as you should be ready by now, perhaps some review of your study notes, but that's it! On the day before your exam your goal is to stay relaxed and rest. When travelling multiple hours to the JNCIE lab testing location, please ensure you book a hotel in advance and arrive on the day before of your exam. If possible, visit the JNCIE lab testing center and familiarize yourself with the location. Go to bed early as tomorrow it will be a long day.

Exam Day

When you wake up, ensure you dress comfortable and have a good breakfast. This will be a very long day! Ensure you arrive 30 minutes before the exam start time. Ensure you do not forget to take a valid ID with you, exam registration confirmation and if you prefer your own keyboard and mouse. If you forget your ID, the JNCIE lab exam proctor will not allow you to take the exam. When it's time to enter the exam room, the JNCIE lab exam proctor will direct you to your workstation. Using the SecureCRT software you can access the JNCIE console server, which provides console access to the exam devices. There are Juniper Networks Service Provider related PDF files installed on your computer, which you can use to lookup certain commands/ features. You do not have internet access. When the proctor tells you that the exam begins, you are allowed then to open the exam documents. Before immediately starting with the first task, we would like to advise you to read the entire exam and study the topology first to fully understand all requirements and how tasks might correlate to each other.

Be familiar with the Windows notepad application as it is a real time saver. For example, if you apply BGP configuration to multiple devices, with just some small changes then notepad is your friend.

Timing is important! Be quick but pay attention to detail. A simple typo in a L3VPN configuration can cause you to fail a complete task, even if the L3VPN is working. Ensure that you meet **all** requirements as mentioned in the exam documents. You will not be given partial points for tasks.

Be sure there is plenty of time at the end of the day to verify every task for the second time. This is important because you might have broken parts of the network during the exam. In addition to this, you might have missed something in the question when you read it the first time. And finally, after going through the entire exam, you might have a different (or better) understanding of what it is that you are supposed to do.

If you are struggling with a task for more than 10-15 minutes, check the task and determine whether it will have impact on the rest of your exam. If it does not impact any

© 2020 Juniper Networks, Inc. All rights reserved.

other task on the exam, move to the next task and revisit the question again later during the exam.

Do not assume any other requirement then then ones mentioned in the tasks. The JNCIE exam is about fulfilling the tasks. It is not about displaying your understanding of the best-practices, scalable network design or anything other than fulfilling the task. So, it might be that your route-reflector is not redundant, the cluster-id is completely random, the active hold-time is 90 seconds and the sessions are not authenticated. If the tasks do not require any of these things, no one cares! Sometimes there are multiple ways to complete a task. In that case you can choose your preferred option. If you are unsure about a task requirement then ask the proctor. Do not be shy. While the JNCIE lab exam proctor cannot assist you in solving the tasks, he or she can certainly help you in clarifying the task requirements.

Notify the JNCIE lab exam proctor immediately if you suspect a hardware issue. If it is, the proctor will fix the issue and might extend your time. If not, the proctor will tell you that the issue is caused by your own configuration and that no additional time will be given.

Make sure that you have at least ***one*** task correct in ***each*** section of the exam, if not you will fail the entire exam even if your total score is above the passing requirement!

Be sure to understand that a configuration mistake might result in point loss for multiple sections of the exam. As an example: if you apply a misconfigured firewall filter it might bring your BGP, IGP neighbors, and perhaps also other protocols down, which results in point loss for multiple sections.

After the Exam

That's it! There is nothing more you can do for now. It's time to wait and hope you passed the JNCIE lab exam. It usually takes 5 business days before you receive your exam results. If you pass, great! If not, we recommend you retake the exam again as soon as possible as all topics are still fresh in your memory! Some candidates wait a couple of months between attempts. This usually results in losing your configuration speed and it takes more time and energy to relearn/memorize all the topics again.

Workbook and Configuration Files

Together with this workbook you can access the configuration files on each device that you will need while working through the chapters of the course. The workbook configurations files consist of initial, part, and final configurations for each chapter. These configuration files are stored on the respective device under the /var/home/lab directory.

© 2020 Juniper Networks, Inc. All rights reserved.

Note: This workbook has been built to be serial in nature meaning that each chapter builds on previous chapter configurations. You can load the start configuration on all devices in the topology to ensure the environment is correct for each chapter and part. For example, if you would like to start with Chapter 4, you must load the ch4-start.config configuration file. In addition to start configurations, end configurations have been provided for each chapter and can be used to verify that all tasks are completed and working properly.

Important: To access the devices in this environment please use the jncie/jncie123 user credentials.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 1: General System Features

***Tip**: Throughout the workbook before you begin a chapter, we recommend that you read the entire chapter before starting with the first task.*

This chapter will focus on initial system configuration and general system features. You will configure various features, such as host names, management network access, management user authentication and authorization, NTP, SNMP, Syslog, RE protection firewall filters, network interfaces, and VRRP. You will be operating 8 devices, R1 through R8 referred to as your routers in this workbook.

© 2020 Juniper Networks, Inc. All rights reserved.

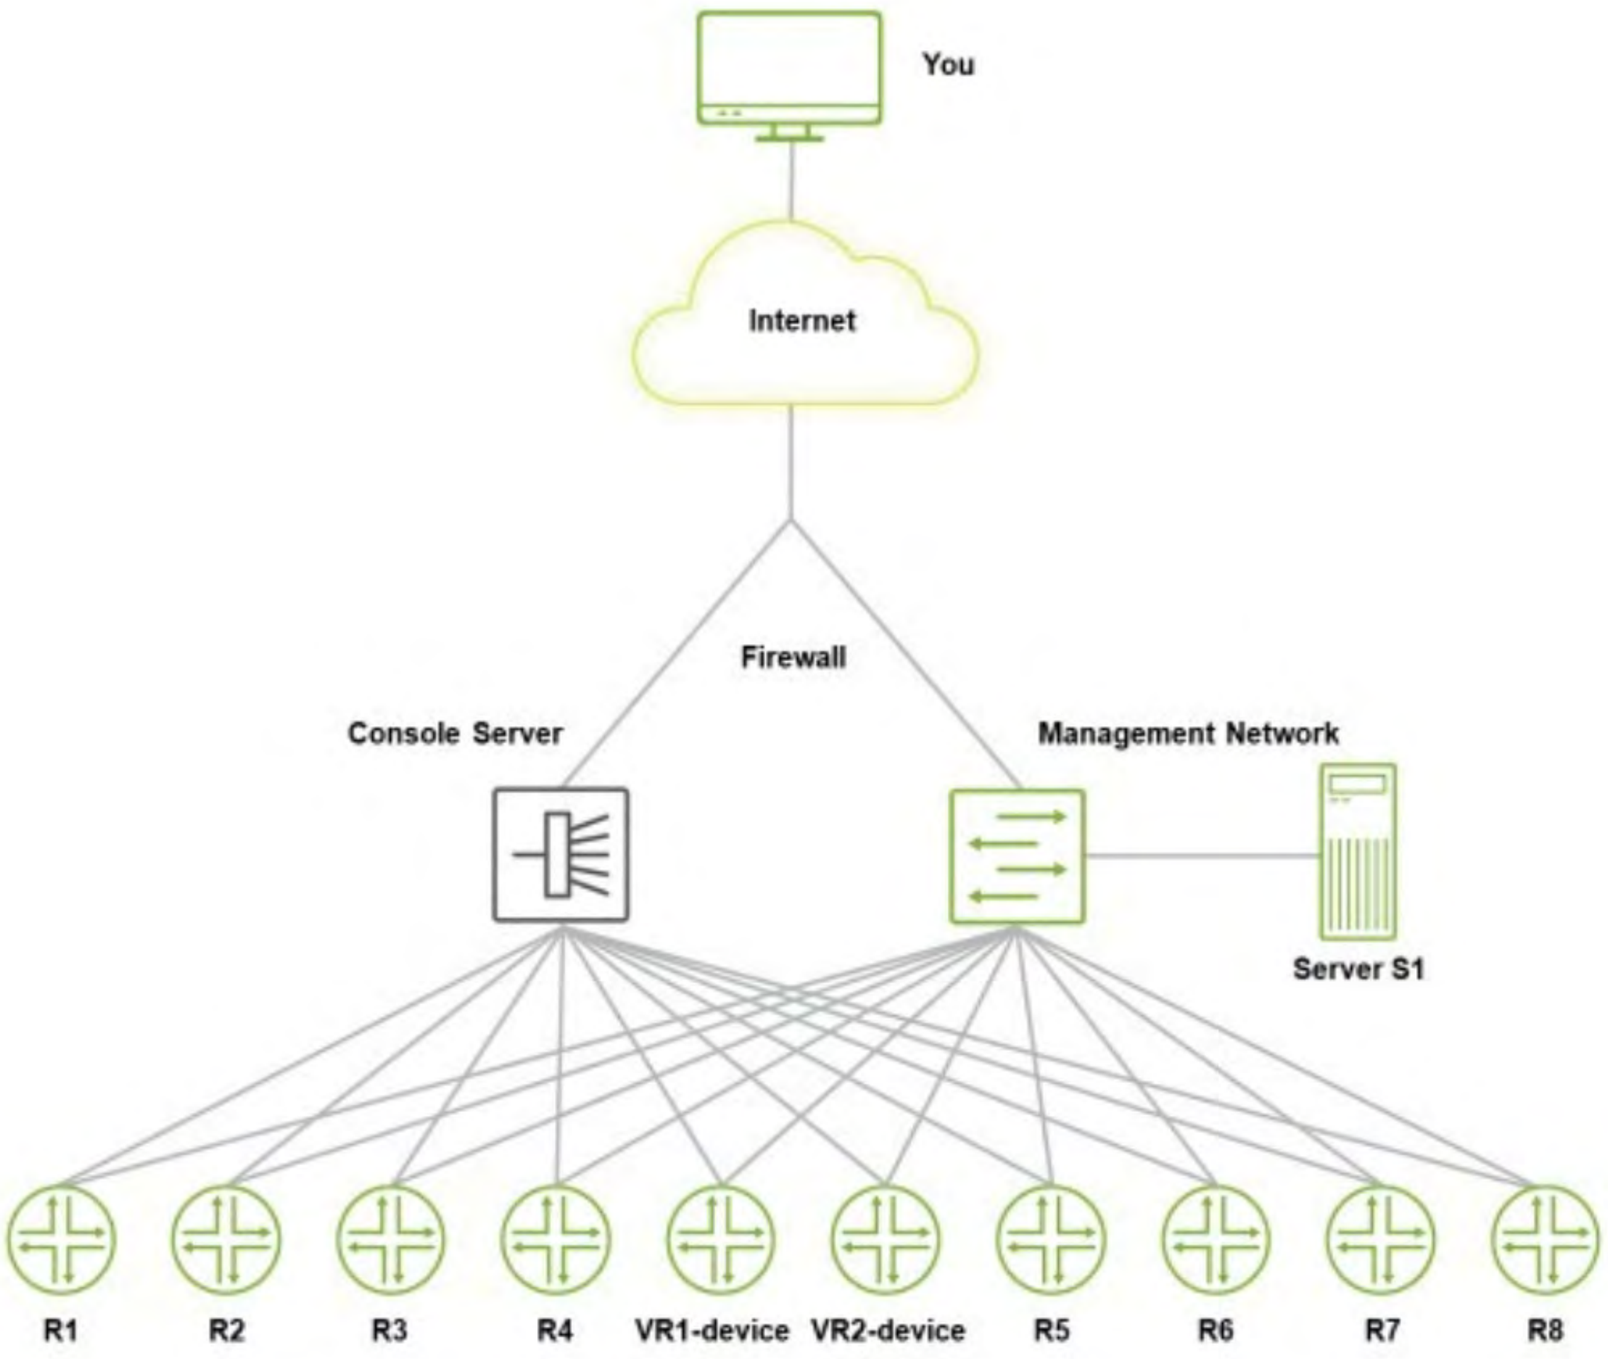
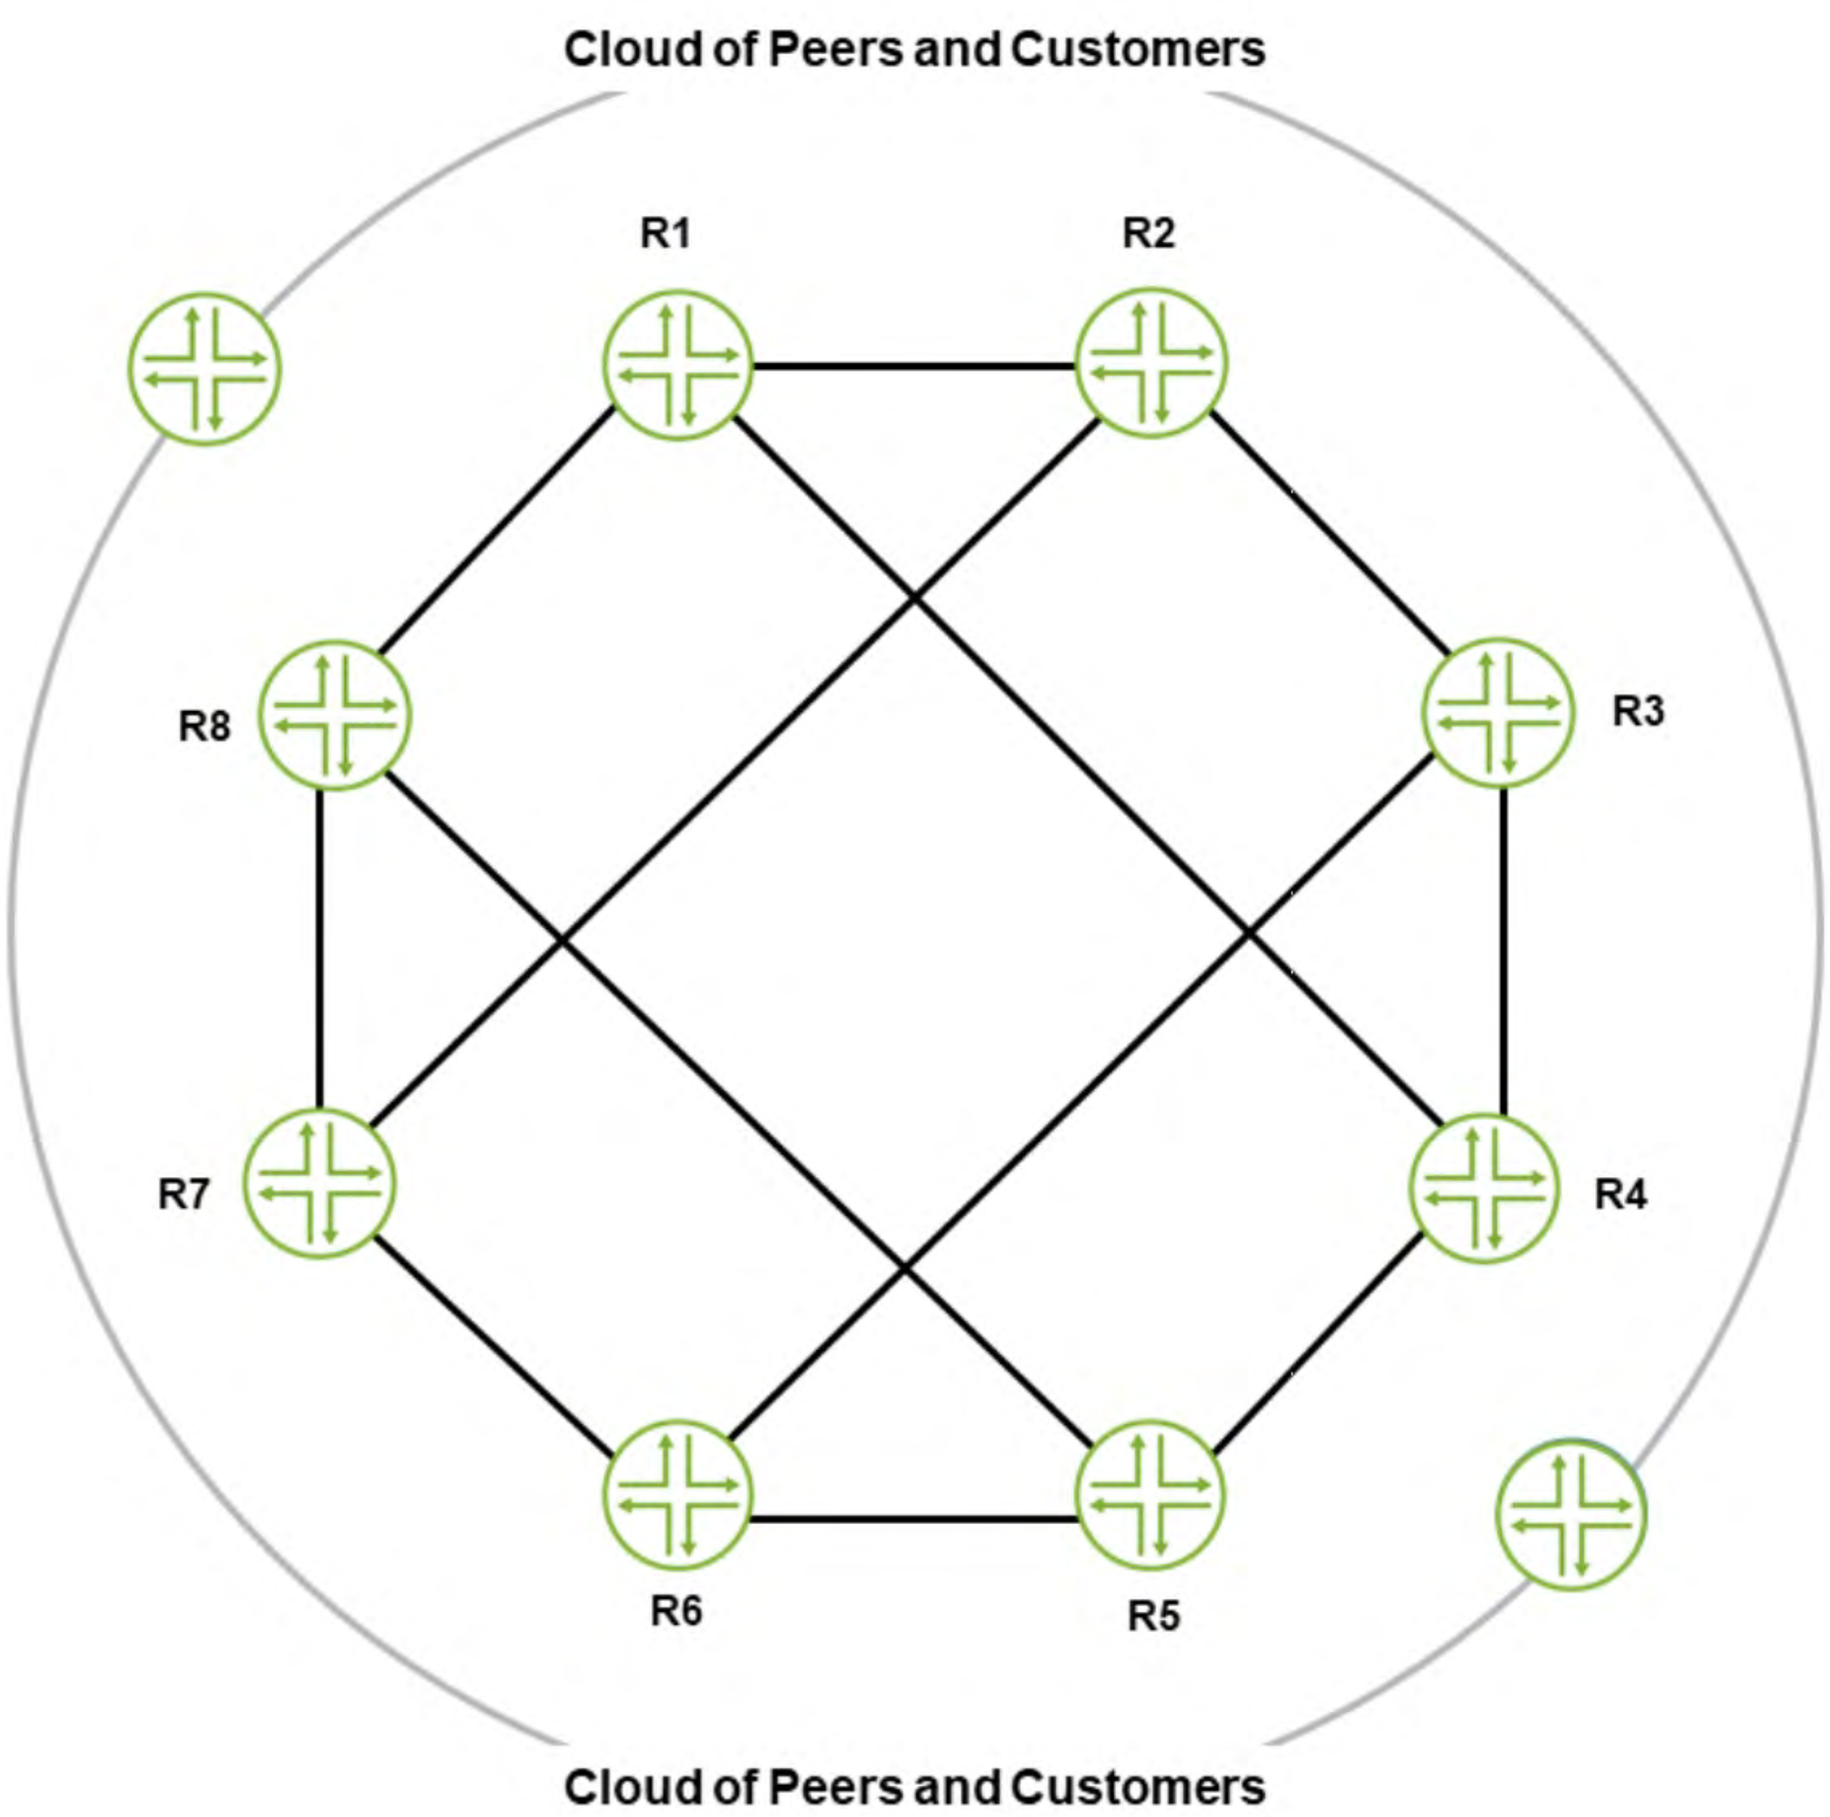


Figure 1

© 2020 Juniper Networks, Inc. All rights reserved.



Generic Network Layout

Figure 2

© 2020 Juniper Networks, Inc. All rights reserved.

Task 1: Initial System Settings

In this task, you will configure your devices' host names, root passwords, the OoB management interfaces, management services, static routing and DNS.

Note: The lab uses two dedicated VR-devices to emulate external systems interacting with your domain. The devices are reachable at 10.10.1.9 and .10 IP addresses using username "lab" and password "lab123".

Note: The student desktop provides all the NTP/FTP/SNMP/Syslog/RADIUS/DNS proxy server related functions for the network. The server is reachable at a 10.10.1.100 alias IP address.

Load the chapter 1 start configurations on the devices and use lab password lab123 on every router. Please do not change the root password on our devices to prevent unnecessary password recovery. Many of the tasks in this section will already be part of the start configuration. You will need to configure these items again to ensure that there are no hidden or intentional mistakes.

1) Configure the host names according to Table 1.

Router	Router Type	Host Name
vmx1	vMX	R1
vmx2	vMX	R2
vmx3	vMX	R3
vmx4	vMX	R4
vmx5	vMX	R5
vmx6	vMX	R6
vmx7	vMX	R7
vmx8	vMX	R8

Table 1

© 2020 Juniper Networks, Inc. All rights reserved.

2) Configure the OoB management interface for each router with the appropriate IP addresses. The routers and their respective IP addresses are listed in Table 2. Set the interface descriptions to your preference.

Host Name	OoB Interface Name	OoB Interface IP Address
R1	fxp0	10.10.1.1/24
R2	fxp0	10.10.1.2/24
R3	fxp0	10.10.1.3/24
R4	fxp0	10.10.1.4/24
R5	fxp0	10.10.1.5/24
R6	fxp0	10.10.1.6/24
R7	fxp0	10.10.1.7/24
R8	fxp0	10.10.1.8/24

Table 2

3) Enable each router to accept management connections for the SSH, Telnet and FTP protocols.
4) Configure a static route for the remote management network 10.10.10/24 with the next-hop 10.10.1.254. Make sure the network is never redistributed into any dynamic routing protocol. Ensure the router is reachable while RPD is not running.
5) Configure the routers to use server S1 as the DNS server.
6) Set the time zone to Europe/Amsterdam on all your devices.
7) Ensure that all routers synchronize their time with the NTP server S1. Configure the devices to synchronize time with the S1 at boot time. Ensure that all the NTP exchanges are authenticated using MD5 with the password workbook.
8) Configure all your devices to transfer their configuration to the FTP server S1 each time the configuration is committed. Use username lab and password lab123 for the FTP server access.
9) Configure the authentication method in such a way that the router first tries to authenticate users on the RADIUS server and then, if not successful, with local password. Use S1 as the RADIUS server. Configure the RADIUS server with retry attempts 1 and a timeout of 2 seconds. Use workbook as the RADIUS shared secret.
10) Create on every router a new user lab, with the password lab123, that will have super user privileges.

© 2020 Juniper Networks, Inc. All rights reserved.

***Tip**: From this point on we recommend you operate routers using the user lab account.*

11) Configure additional users on all the devices as defined in Table 3. Note that word "any" in the Table 3 is used literally, i.e. a user can have any username.

Username	Password	Privileges
any	-	Permissions "view" and "view-configuration". Authenticated only by the RADIUS
ops	ops123	Permissions "clear", "network", "reset", "trace" and "view"
noc	noc123	Permissions "all". Additionally, cannot execute any of the "clear", "configure", "edit" or "start shell" commands

Table 3

12) Configure the Syslog settings on all your devices as indicated in Table 4.

Receiver	Message Type
File "jncie-sp-messages"	All info level messages
Syslog server S1	Interactive commands
	Configuration changes
	All notice level messages
File "user-commands"	All user interactive commands
User "ops"	All warning level messages

Table 4

13) Set the Syslog archive size to 3 files with 100 Kb each.

Task 2: SNMP Configuration

In this task, you will configure SNMPv3 for secure NMS interaction.

1) Configure SNMPv3 view parameters according to Table 5. Make sure that SNMPv3 provides read only access.

© 2020 Juniper Networks, Inc. All rights reserved.

Parameter	Value
USM username	lab
USM user authentication	SHA
USM user authentication password	workbook
USM user encryption	3DES
USM user encryption password	workbook
VACM security model	usm
VACM user	lab
VACM security level	privacy
VACM read view OID	.1

Table 5

2) Configure SNMPv3 notification parameters according to Table 6.

Parameter	Value
Target address	lab
Target processing model	SHA
Target security model	workbook
Target security level	3DES
Target security name	workbook
Notification OID filter	usm
Notification type	lab

Table 6

Task 3: Firewall Filters

In this task, you will configure a firewall filter for Routing Engine (RE) protection.

1) Configure an IPv4 firewall filter allowing protocol messages from AH, BFD, VRRP, RIP, OSPF, RSVP, LDP, PIM, IGMP, MSDP protocols.
2) Configure the firewall filter so that BGP messages are accepted only from configured BGP neighbors. Make sure that a configured BGP neighbor is automatically allowed in the firewall filter.
3) Configure the firewall filter to accept NTP, RADIUS, DNS, SNMP, SSH, Telnet, FTP protocols only from the 10.10.1/24 management network.

© 2020 Juniper Networks, Inc. All rights reserved.

4) Configure the firewall filter to accept ICMP and traceroute messages. Ensure that the flow of the messages is limited to 100kbps with a burst size of 25K. The excess traffic must be dropped.
5) Configure the firewall filter to discard any other traffic, increment a named drop counter and send a log message.
6) Apply the firewall filter such as to ensure that it is used for RE protection.

Task 4: Interface Configuration

In this task, you are configuring the network interfaces, aggregated Ethernet interfaces and VRRP.

1) Build the network as shown in Figure 3. The interface parameters can be found in Table 7. Configure interfaces on R1 and R2, and R5 and R6 to form an aggregated Ethernet bundle. Enable LACP continuity checking on the AE interface. Configure the logical interface descriptions.

© 2020 Juniper Networks, Inc. All rights reserved.

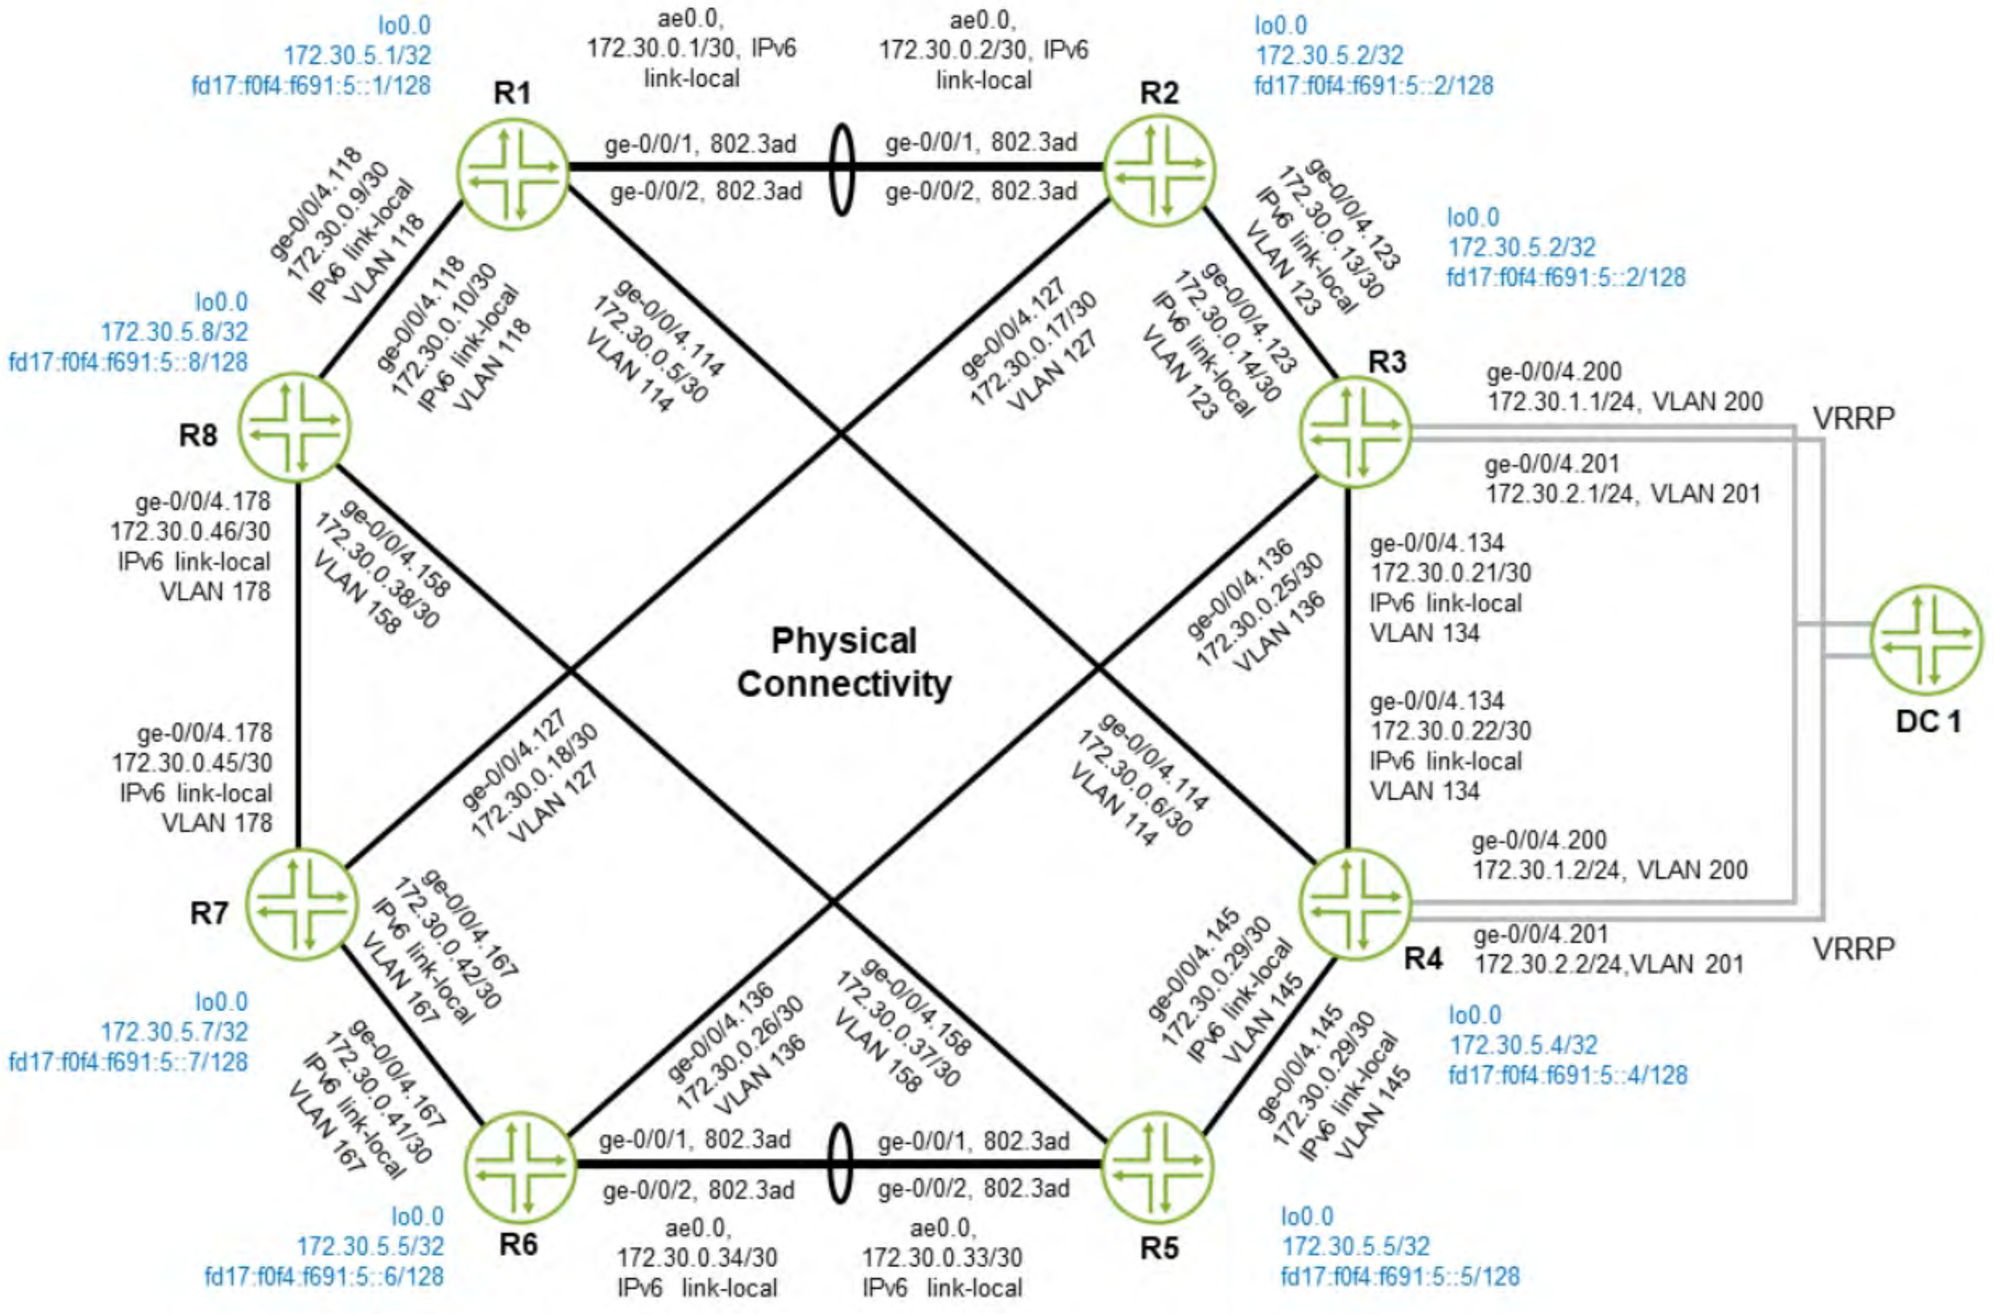


Figure 3

Note: The interface unit numbers match the VLAN tags.

© 2020 Juniper Networks, Inc. All rights reserved.

Router	Interface Name	IP Address	IPv6 Address
R1	ge-0/0/1	802.3ad	
	ge-0/0/4.114	172.30.0.5/30	
	ge-0/0/4.118	172.30.0.9/30	link-local
	ge-0/0/2	802.3ad	
	ae0.0	172.30.0.1/30	link-local
	lo0.0	172.30.5.1/32	fd17:f0f4:f691:5::1/128
R2	ge-0/0/1	802.3ad	
	ge-0/0/4.127	172.30.0.17/30	
	ge-0/0/4.123	172.30.0.13/30	link-local
	ge-0/0/2	802.3ad	
	ae0.0	172.30.0.2/30	link-local
	lo0.0	172.30.5.2/32	fd17:f0f4:f691:5::2/128
R3	ge-0/0/4.134	172.30.0.21/30	link-local
	ge-0/0/4.136	172.30.0.25/30	
	ge-0/0/4.123	172.30.0.14/30	link-local
	ge-0/0/4.200	172.30.1.1/24	
	ge-0/0/4.201	172.30.2.1/24	
	lo0.0	172.30.5.3/32	fd17:f0f4:f691:5::3/128
R4	ge-0/0/4.134	172.30.0.22/30	link-local
	ge-0/0/4.114	172.30.0.6/30	
	ge-0/0/4.145	172.30.0.29/30	link-local
	ge-0/0/4.200	172.30.1.2/24	
	ge-0/0/4.201	172.30.2.2/24	
	lo0.0	172.30.5.4/32	fd17:f0f4:f691:5::4/128
R5	ge-0/0/1	802.3ad	
	ge-0/0/4.158	172.30.0.37/30	
	ge-0/0/4.145	172.30.0.30/30	link-local
	ge-0/0/2	802.3ad	
	ae0.0	172.30.0.33/30	link-local
	lo0.0	172.30.5.5/32	fd17:f0f4:f691:5::5/128
R6	ge-0/0/1	802.3ad	
	ge-0/0/4.136	172.30.0.26/30	

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/4.167	172.30.0.41/30	link-local
	ge-0/0/2	802.3ad	
	ae0.0	172.30.0.34/30	link-local
	lo0.0	172.30.5.6/32	fd17:f0f4:f691:5::6/128
R7	ge-0/0/4.178	172.30.0.45/30	link-local
	ge-0/0/4.127	172.30.0.18/30	
	ge-0/0/4.167	172.30.0.42/30	link-local
	lo0.0	172.30.5.7/32	fd17:f0f4:f691:5::7/128
R8	ge-0/0/4.178	172.30.0.46/30	link-local
	ge-0/0/4.158	172.30.0.38/30	
	ge-0/0/4.118	172.30.0.10/30	link-local
	lo0.0	172.30.5.8/32	fd17:f0f4:f691:5::8/128

Table 7

2) On R3 and R4 configure VRRP such as R3 is the VRRP master on ge-0/0/4.200 interface and R4 is the VRRP master on ge-0/0/4.201 interface. Use .254 Virtual Router IP address on the subnets.
3) Make sure that R3 and R4 track their uplink interfaces (R3 = ge-0/0/4.136, ge-0/0/4.123) (R4 = ge-0/0/4.114, ge-0/0/4.145) so that if both the interfaces go down the device resigns from its VRRP mastership.
4) Make sure that VRRP messages are authenticated with MD5. Use workbook as the authentication key.

Task 5: Scripting

In this task, you will download and apply operational, event and commit scripts.

Note: These are example scripts written by Juniper Networks and available in the public domain. Writing your own scripts is beyond the scope of this workbook.

1) Download the op script called and “show-interfaces.slax” from the FTP server S1 to all your routers.

***Tip**: This op script adds descriptions and protocol filtering to the normal "show interfaces terse" command. Two arguments (interface and protocol) provide additional filtering.*

2) Download the commit script called “interface-mask-check.slax” from the FTP server S1 to all your routers.

© 2020 Juniper Networks, Inc. All rights reserved.

***Tip**: This commit script verifies that the ipv4 address on each interface has a network mask of 24 or greater. If the mask is less than /24 then a warning is issued.*

3) Download the event script called and "syslog-int-desc-on-link-change.slax" from the FTP server S1 to all your routers.

***Tip**: This event script generates a new syslog message based on the triggering syslog message of SNMP_TRAP_LINK_DOWN or SNMP_TRAP_LINK_UP. It collects the related interface information from the syslog message and grabs the interface description to form a new syslog message.*

4) Enable the scripts.
5) Verify that the scripts are operational.
6) **Save your configuration on all your devices in a named file F1 in order to use it as the baseline configuration for subsequent labs.**

Note: You can call the file anything. F1 is used here as a reference name.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 2: IGP Configuration and Troubleshooting

This chapter contains three independent tasks: OSPF troubleshooting, IS-IS troubleshooting and the new IGP rollout. NOTE: You need the final configurations you have saved in the previous chapter or you can also use the ch1-end.config configuration file.

Task 1: OSPF Troubleshooting

In this task, you load a broken OSPF configuration, troubleshoot it and fix the errors. The network diagram is shown in Figure 4. Table 8 shows interface to area designation.

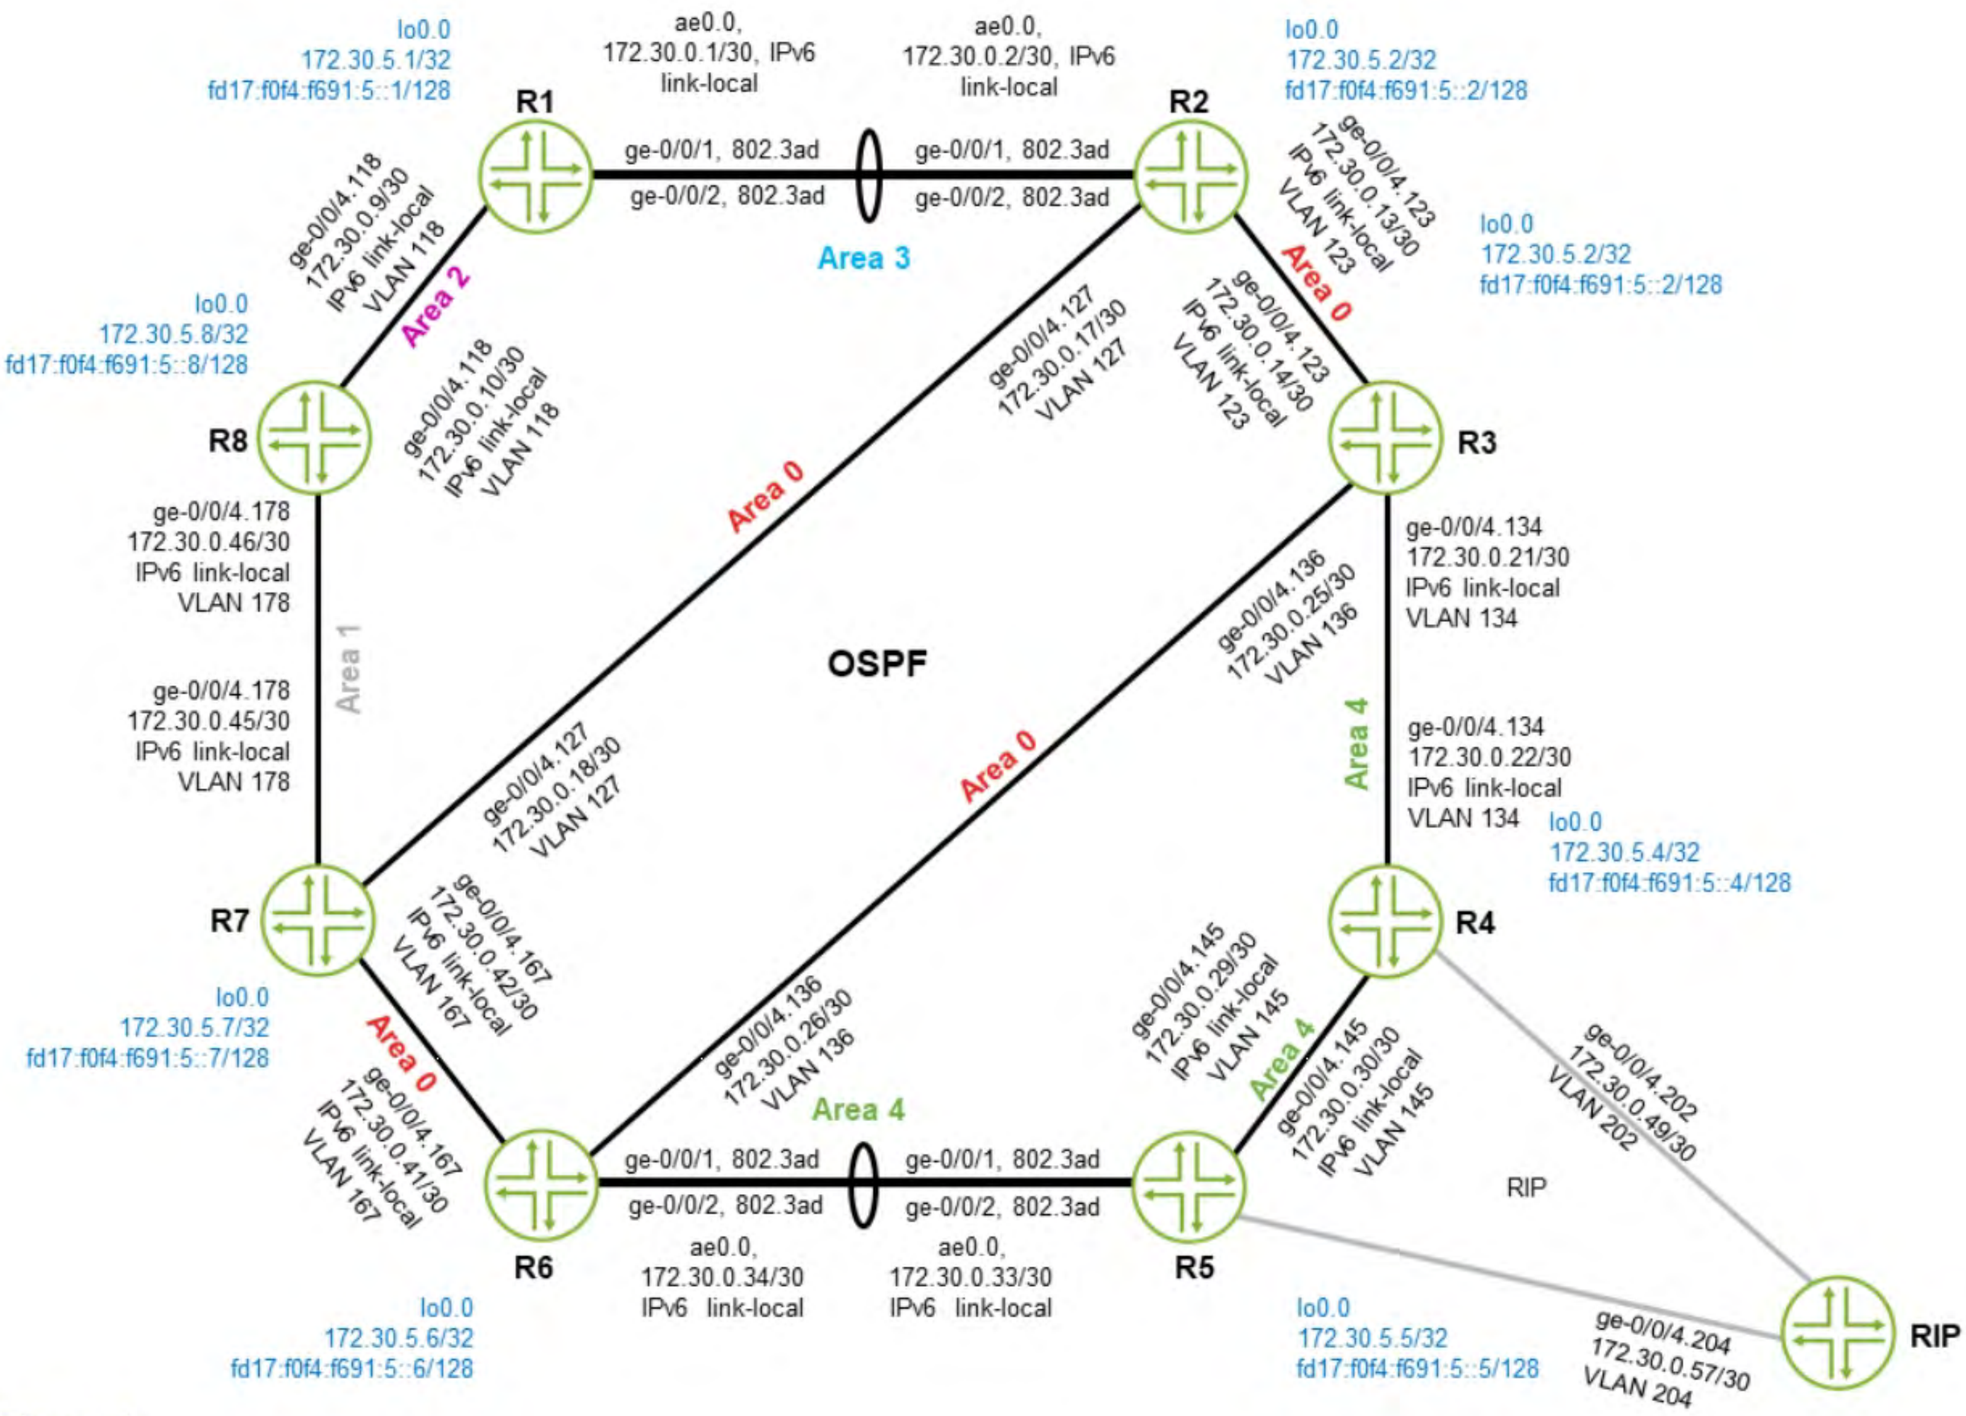


Figure 1

© 2020 Juniper Networks, Inc. All rights reserved.

Router	Interface	Area
R1	ae0.0	0.0.0.3
	ge-0/0/4.118	0.0.0.2
	lo0.0	0.0.0.2
R2	ae0.0	0.0.0.3
	ge-0/0/4.127	0.0.0.0
	ge-0/0/4.123	0.0.0.0
	lo0.0	0.0.0.0
R3	ge-0/0/4.134	0.0.0.4
	ge-0/0/4.136	0.0.0.0
	ge-0/0/4.123	0.0.0.0
	lo0.0	0.0.0.0
R4	ge-0/0/4.134	0.0.0.4
	ge-0/0/4.145	0.0.0.4
	lo0.0	0.0.0.4
R5	ae0.0	0.0.0.4
	ge-0/0/4.145	0.0.0.4
	lo0.0	0.0.0.4
R6	ae0.0	0.0.0.4
	ge-0/0/4.136	0.0.0.0
	ge-0/0/4.167	0.0.0.0
	lo0.0	0.0.0.0
R7	ge-0/0/4.178	0.0.0.1
	ge-0/0/4.127	0.0.0.0
	ge-0/0/4.167	0.0.0.0
	lo0.0	0.0.0.0
R8	ge-0/0/4.178	0.0.0.1
	ge-0/0/4.118	0.0.0.2
	lo0.0	0.0.0.2

Table 1

The OSPF network must meet the following criteria:

- All OSPF adjacencies are full.
- All your routers can reach all other routers loopbacks.
- No routing loops are allowed anywhere.

© 2020 Juniper Networks, Inc. All rights reserved.

- All routers must use MD5 authentication on all OSPF interfaces.
- All RIP routes must be seen in area 4.
- The backbone area must have a single summarized route to RIP destinations.
- The default route must be advertised to the RIP router. The RIP router must prefer R4 updates.
- No Type 2, 3, 4 and 5 LSA's are allowed in area 4.
- Any ABR failure must not have any area isolated.
- Any ASBR failure must not result in RIP routes disappearing from the OSPF domain or the default route disappearing from the RIP domain.
- No static routing is allowed.

1) Load and override your routers' configuration with the task reset configuration.
2) Using operational and configuration mode commands, troubleshoot the OSPF network, and fix the errors.
3) Write a summary report on all the issues found.

Task 2: IS-IS Troubleshooting

In this task, you load a broken IS-IS configuration, troubleshoot it and fix the errors. The network diagram is shown in Figure 5. Table 9 shows interface to level designation. Table 10 shows router to area designation.

© 2020 Juniper Networks, Inc. All rights reserved.

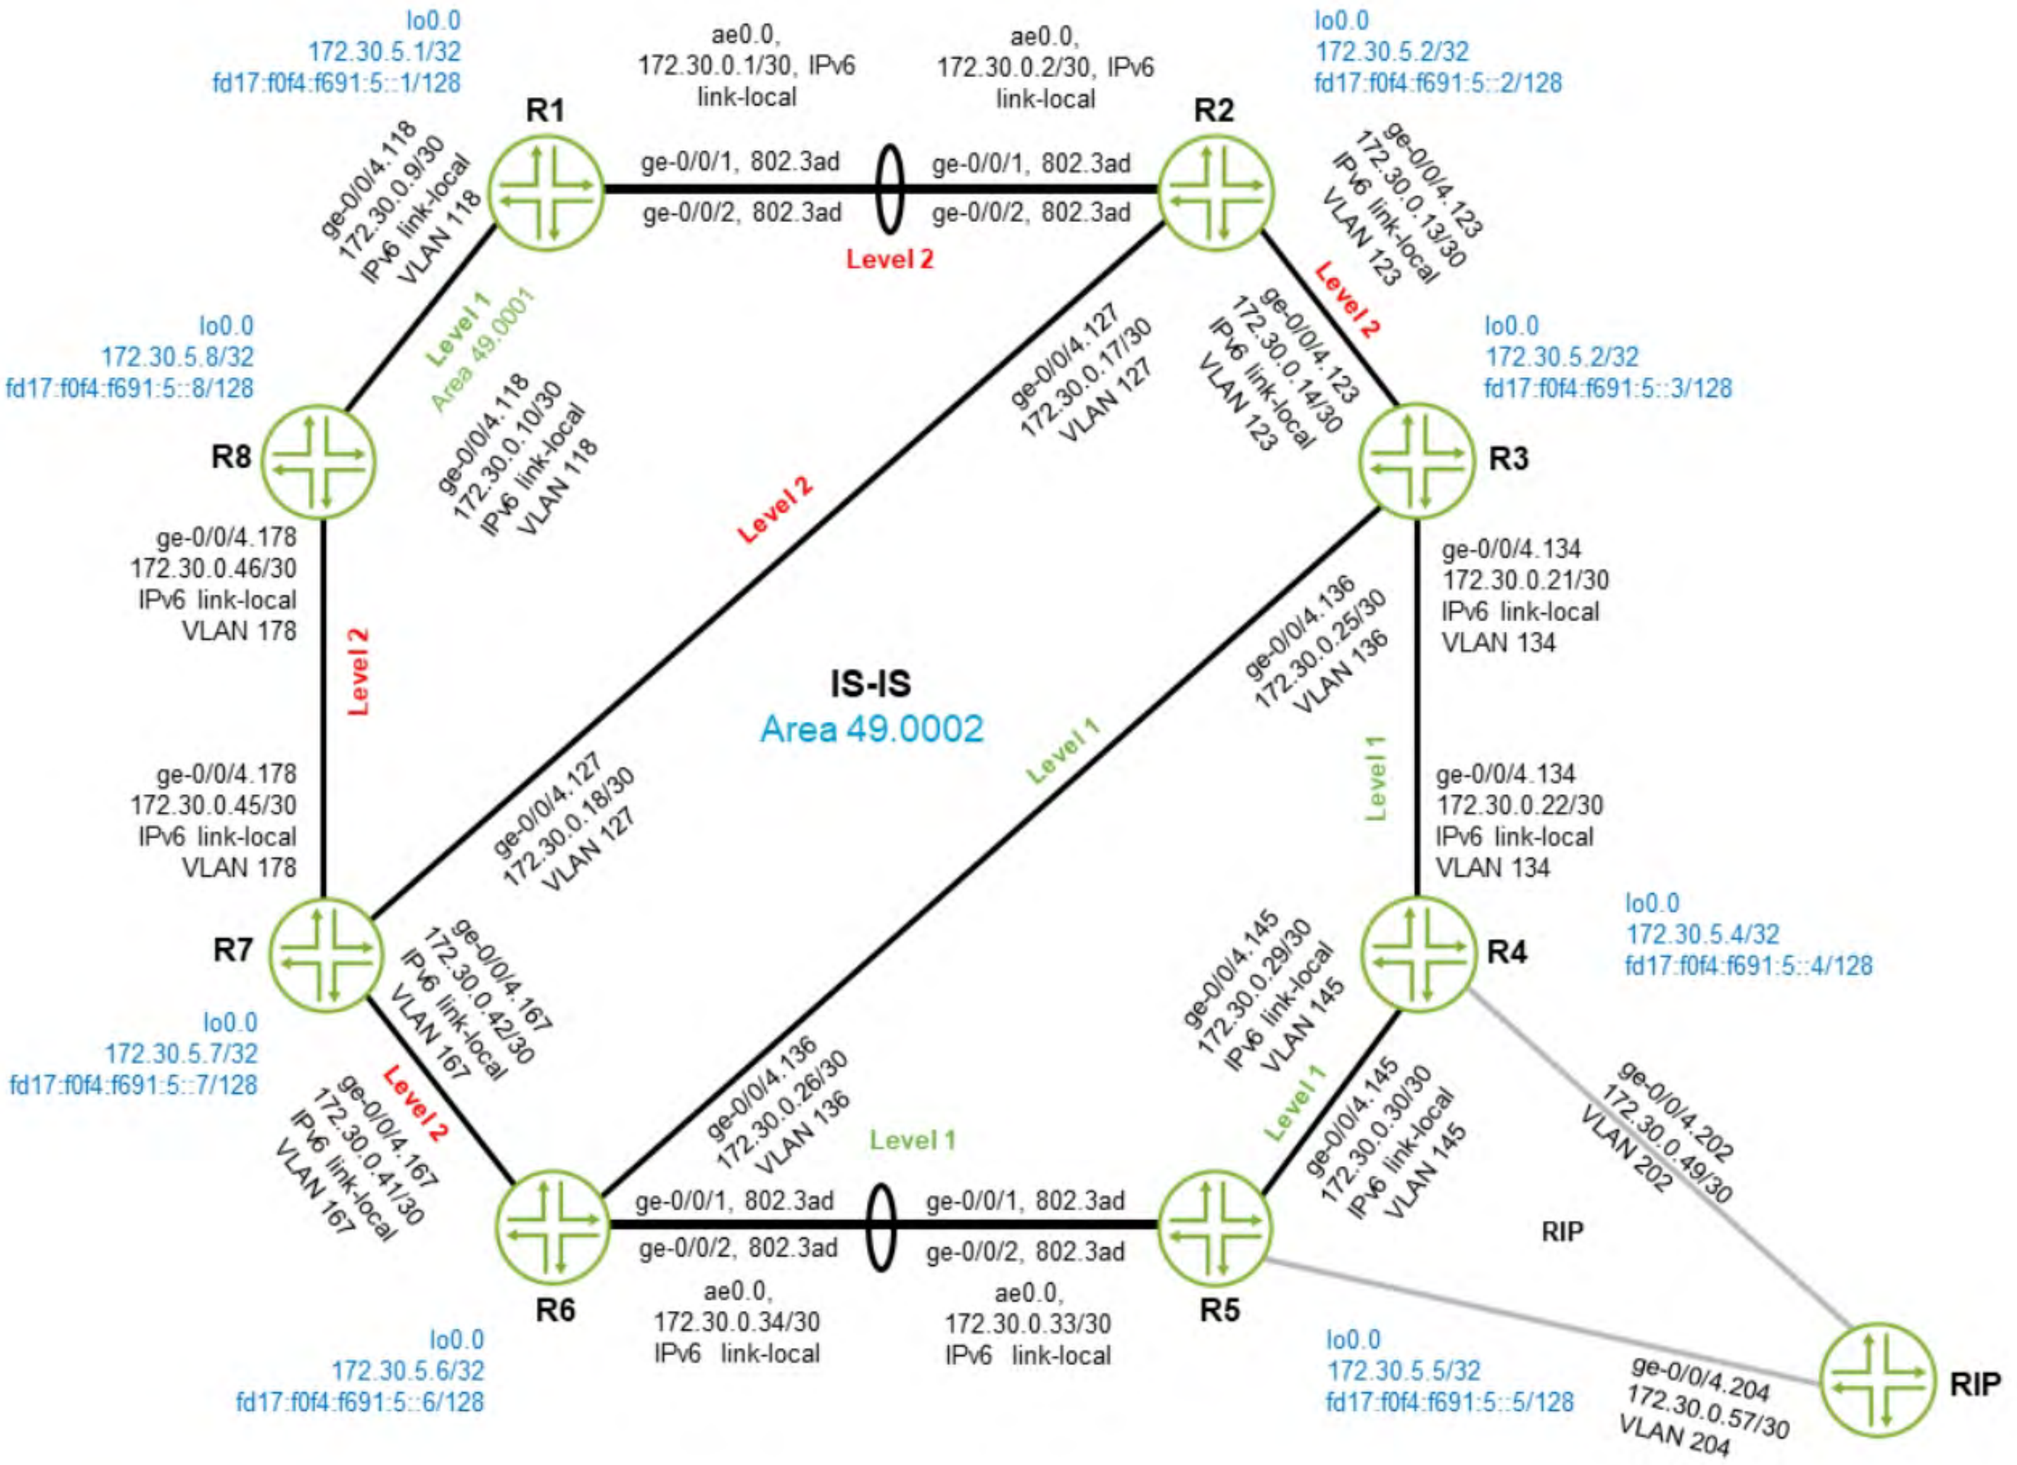


Figure 2

© 2020 Juniper Networks, Inc. All rights reserved.

Router	Interface	Level
R1	ae0.0	2
	ge-0/0/4.118	1
	lo0.0	1
R2	ae0.0	2
	ge-0/0/4.127	2
	ge-0/0/4.123	2
	lo0.0	2
R3	ge-0/0/4.134	1
	ge-0/0/4.136	1
	ge-0/0/4.123	2
	lo0.0	1
R4	ge-0/0/4.134	1
	ge-0/0/4.145	1
	lo0.0	1
R5	ae0.0	1
	ge-0/0/4.145	1
	lo0.0	1
R6	ae0.0	1
	ge-0/0/4.136	1
	ge-0/0/4.167	2
	lo0.0	1
R7	ge-0/0/4.178	2
	ge-0/0/4.127	2
	ge-0/0/4.167	2
	lo0.0	2
R8	ge-0/0/4.178	2
	ge-0/0/4.118	1
	lo0.0	1

Table 9

Router	Area
R1	49.0001
R2	49.0002
R3	49.0002

© 2020 Juniper Networks, Inc. All rights reserved.

R4	49.0002
R5	49.0002
R6	49.0002
R7	49.0002
R8	49.0001

Table 10

The IS-IS network must meet the following criteria:

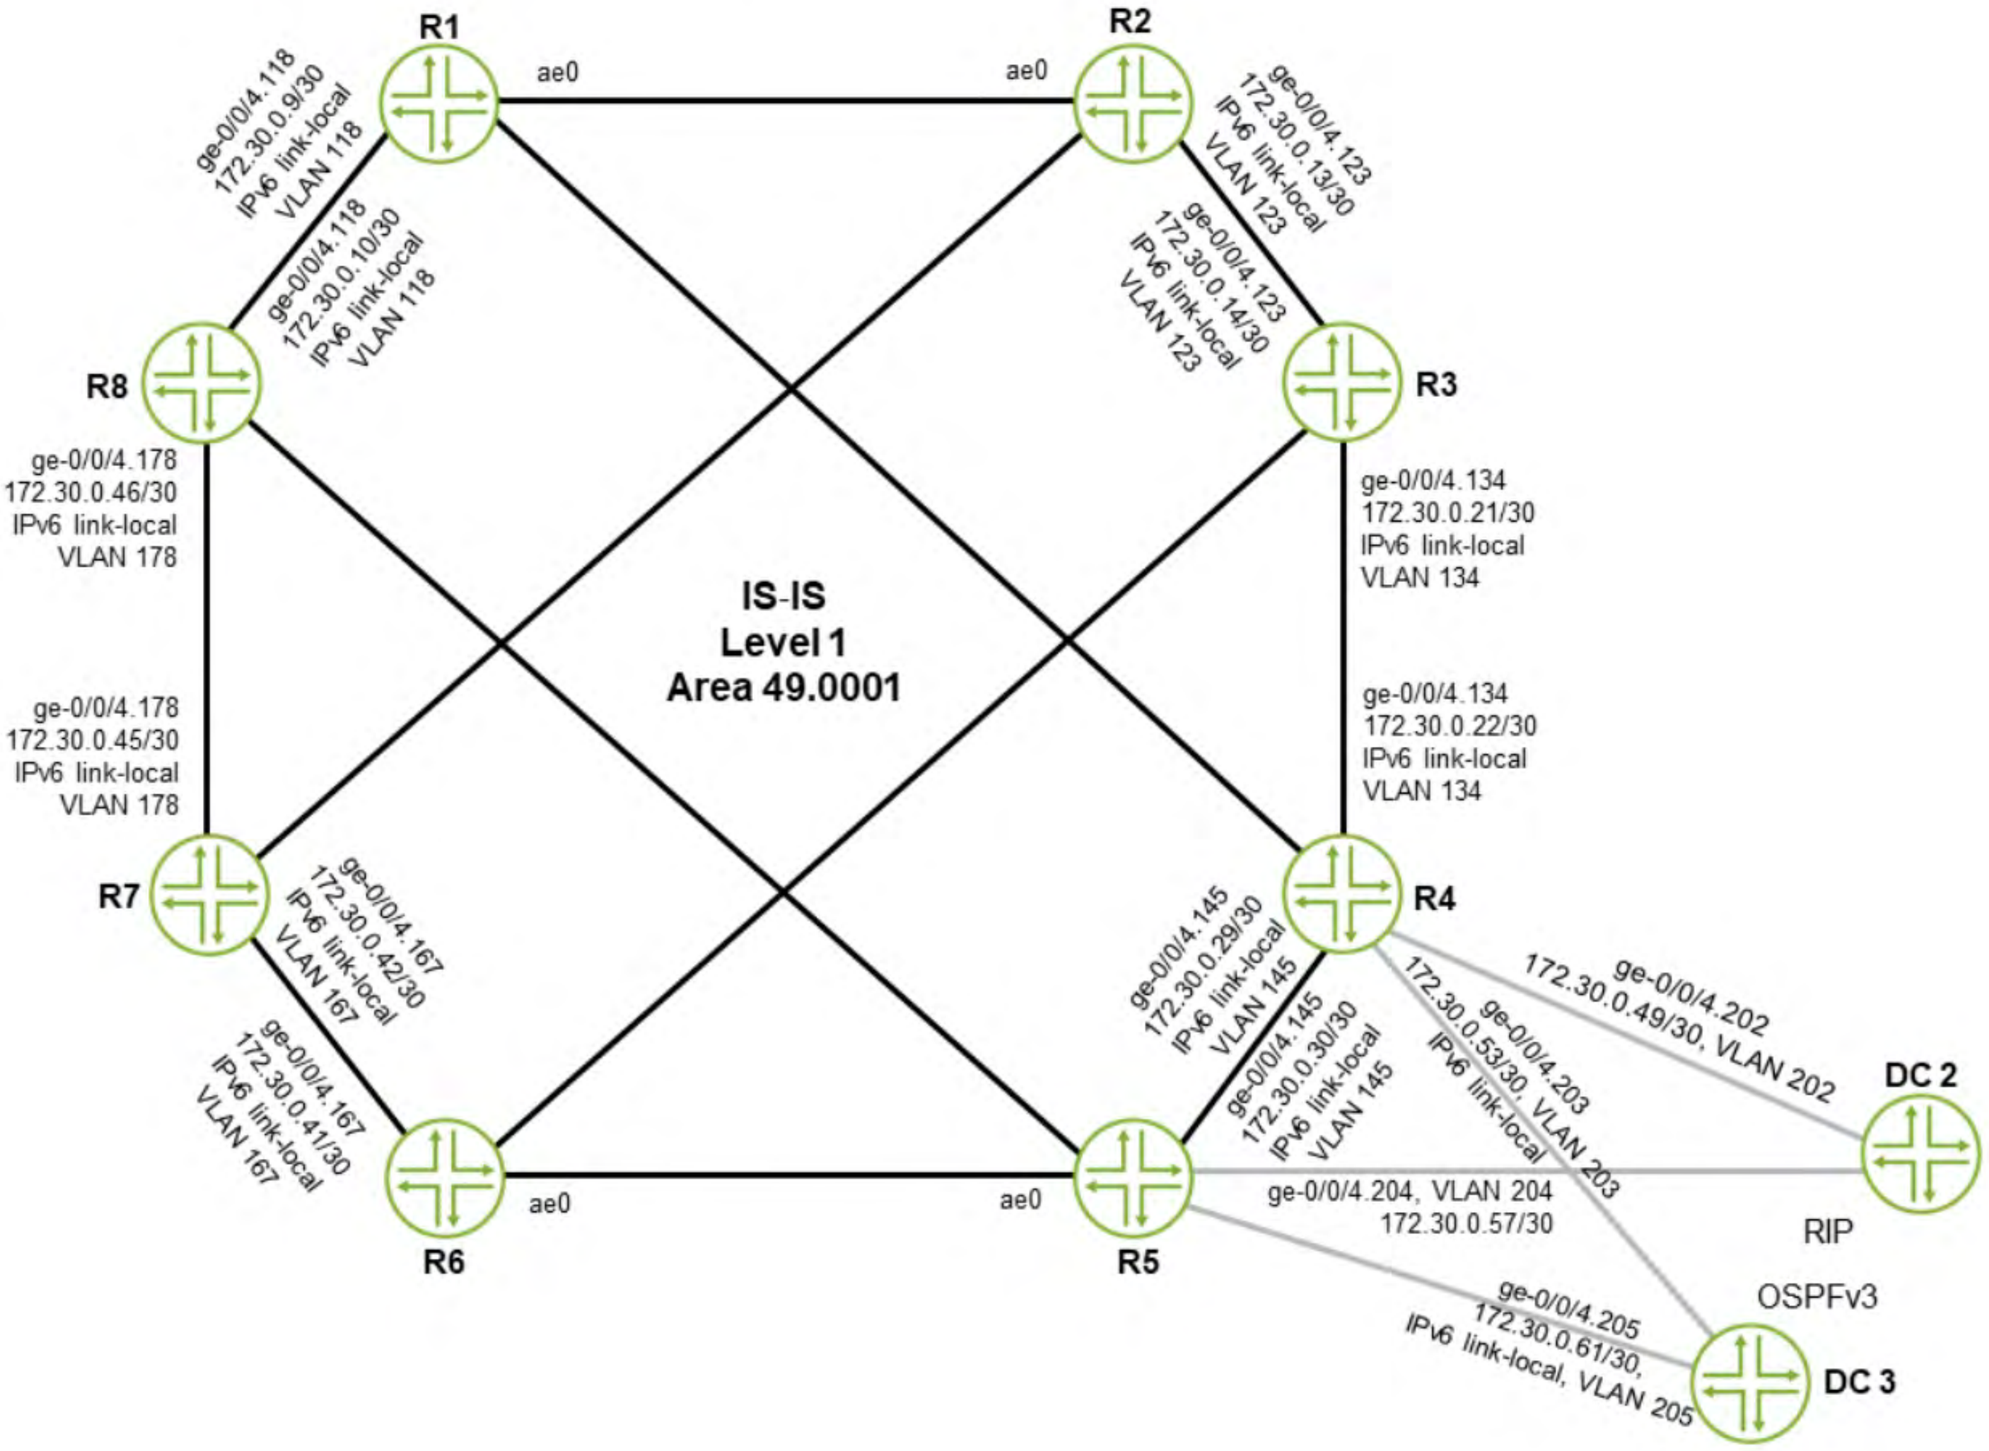
- All IS-IS adjacencies are up.
- All your routers can reach all other routers loopbacks.
- No routing loops are allowed anywhere.
- Each IS-IS interface must have no more than one adjacency.
- All routers must use MD5 authentication for Hello IS-IS PDU only on all IS-IS interfaces.
- L2 interfaces must not elect DIS.
- All RIP routes must be seen in all L1 routers database in area 49.0002.
- The level 2 must have a single summarized route to RIP destinations.
- All IS-IS routes must be advertised to the RIP router. The RIP router must prefer R4 updates.
- Any L1/L2 router failure must not have any L1 area isolated.
- Any ASBR failure must not result in RIP routes disappearing from the IS-IS domain or the default route disappearing from the RIP domain.
- No static routing is allowed.
- Load and override your routers' configuration with the task reset configuration.

1) Using operational and configuration mode commands troubleshoot the IS-IS network and fix the errors.
2) Write a summary report on all the issues found.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 3: IGP Rollout

In this task, you will configure a flat single area single level IS-IS network that will be used as a foundation for the subsequent tasks.



© 2020 Juniper Networks, Inc. All rights reserved.

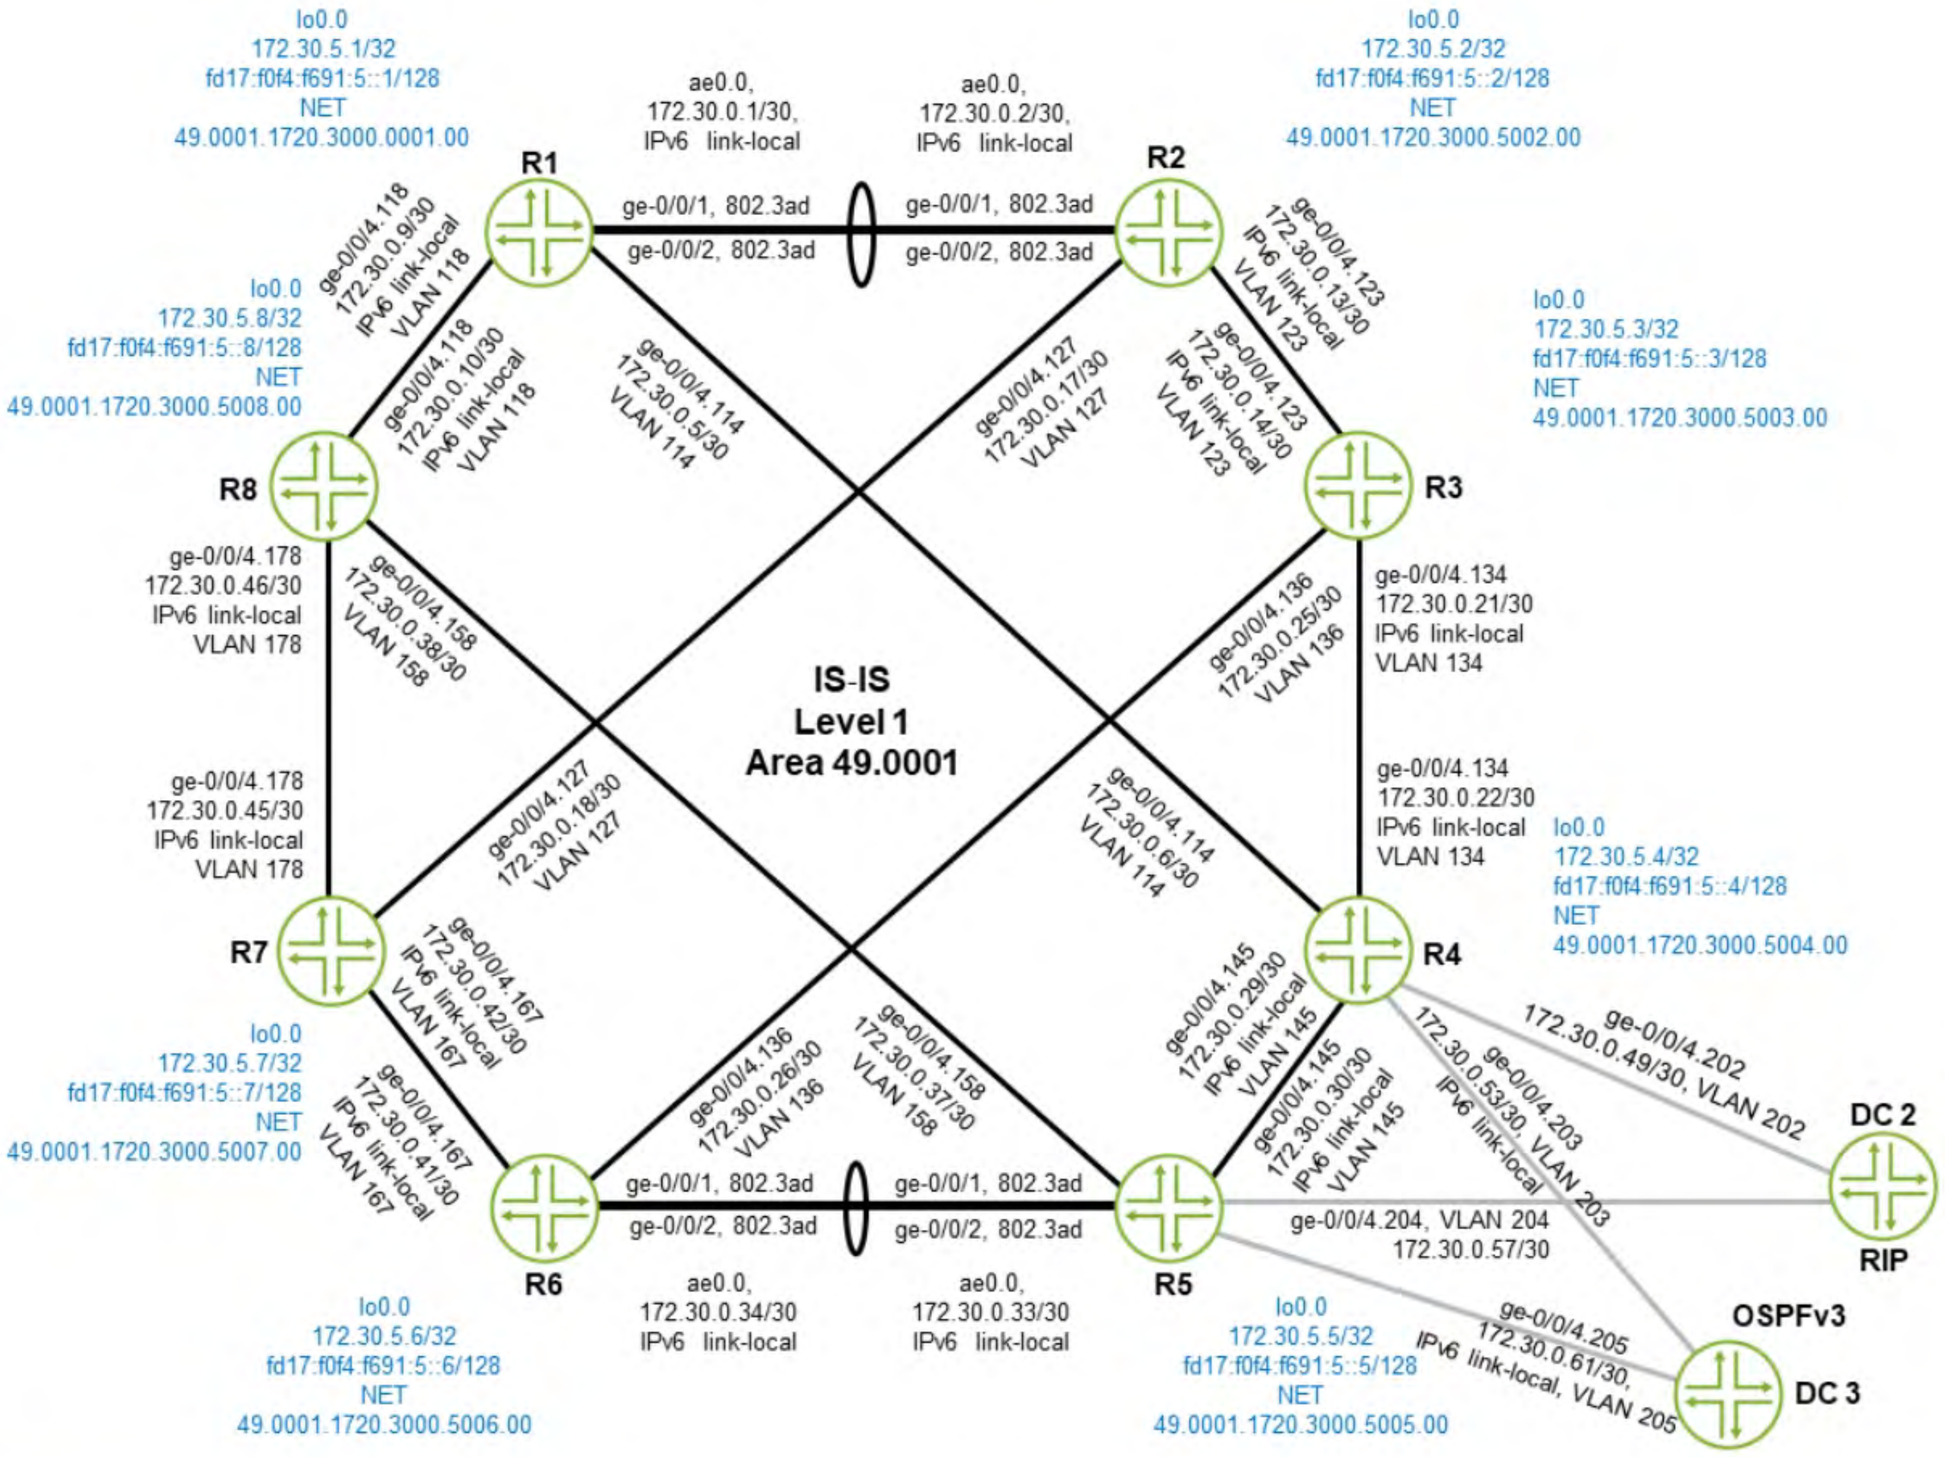


Figure 3

Note: You are not allowed to use static routes in this and all subsequent chapter tasks unless indicated explicitly.

1) Load and override your routers' configuration with that of saved in the file(s) F1.
2) Configure additional interfaces on your routers as indicated in Table 11. Set the interfaces description.

Router	Interface Name	IP Address	IPv6 Address
R4	ge-0/0/4.202	172.30.0.49/30	
	ge-0/0/4.203	172.30.0.53/30	link-local
R5	ge-0/0/4.204	172.30.0.57/30	
	ge-0/0/4.205	172.30.0.61/30	link-local

Table 2

© 2020 Juniper Networks, Inc. All rights reserved.

3) Use figure 6 and Table 12 to configure the listed NET addresses on all routers.

Router	NET
R1	49.0001.1720.3000.5001.00
R2	49.0001.1720.3000.5002.00
R3	49.0001.1720.3000.5003.00
R4	49.0001.1720.3000.5004.00
R5	49.0001.1720.3000.5005.00
R6	49.0001.1720.3000.5006.00
R7	49.0001.1720.3000.5007.00
R8	49.0001.1720.3000.5008.00

Table 12

4) Make sure that Router IDs are set explicitly on all your routers equal to the loopback IP address.
5) Make sure that both the VRRP subnets appear in the IS-IS domain but the IS-IS adjacencies are not formed on them. Make sure that any of the R3 or R4 failure will not result in the VRRP subnets disappearing from the IS-IS domain.
6) Make sure that no pseudo nodes enter into the IS-IS database.
7) Configure MD5 authentication on all IS-IS enabled interfaces for all IS-IS PDUs.
8) Configure all routers to automatically calculate metrics based on interface bandwidth. Make sure that narrow metrics are not used.
9) Make sure that IS-IS neighbors can detect the adjacency loss within 6 seconds. You are not allowed to alter the IS-IS timers.
10) Make sure that all adjacencies are up, and all routers can reach all other routers' IPv4 loopback addresses.
11) Configure RIP on R4 (ge-0/0/4.202) and R5 (ge-0/0/4.204) interfaces respectively.
12) Advertise only the default route to the RIP router. Make sure that any of the R4 or R5 failure will not result in the default route disappearing from the RIP domain.
13) Advertise the received RIP routes to IS-IS. Make sure that any of the R4 or R5 failure will not result in the RIP routes disappearing from the IS-IS domain.
14) Make sure that the default route received from RIP is not installed into the routing table.
15) Make sure that all your routers can reach all other routers' IPv6 loopback addresses.

© 2020 Juniper Networks, Inc. All rights reserved.

16) Configure OSPFv3 area 0 on R4 (ge-0/0/4.203) and R5 (ge-0/0/4.205) interfaces respectively. Make sure that OSPFv3 supports both IPv4 and IPv6 routing.
17) Advertise IPv4 and IPv6 IS-IS routes to OSPFv3. Advertise IPv4 and IPv6 OSPFv3 routes to IS-IS. Make sure that any of the R4 or R5 failure will not disrupt the routing between the IS-IS and OSPFv3 domains.
18) Advertise RIP routes to OSPFv3. Advertise IPv4 OSPFv3 routes to RIP. Make sure that any of the R4 or R5 failure will not disrupt the routing between the OSPFv3 and RIP domains.
19) No routing loops or suboptimal routing are allowed anywhere.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 3: BGP and Routing Policy

In this chapter, you will create the BGP network including IBGP with Route Reflection and Confederation, and multiple EBGP sessions with peers and customers emulating a typical ISP setup. You will also configure multiple routing policies to achieve high accuracy control over BGP routing exchange and path selection.

Task 1: IBGP and Confederation

In this task, you build an IBGP confederation network.

1) Configure a confederation network. Make sure that no router has more than 2 IBGP neighbors. An arbitrary number of CBGP sessions are allowed.
2) Make sure that the IBGP sessions use the loopback interface for peering.
3) Make sure that any of the routers failure will not result in any of the Sub-AS isolated.
4) Configure MD5 authentication for all IBGP and CBGP sessions.
5) Ensure that all the IBGP and CBGP session state change is logged to syslog.

Task 2: EBGP Configuration

In this task, you configure IPv4 and IPv6 EBGP peering.

© 2020 Juniper Networks, Inc. All rights reserved.

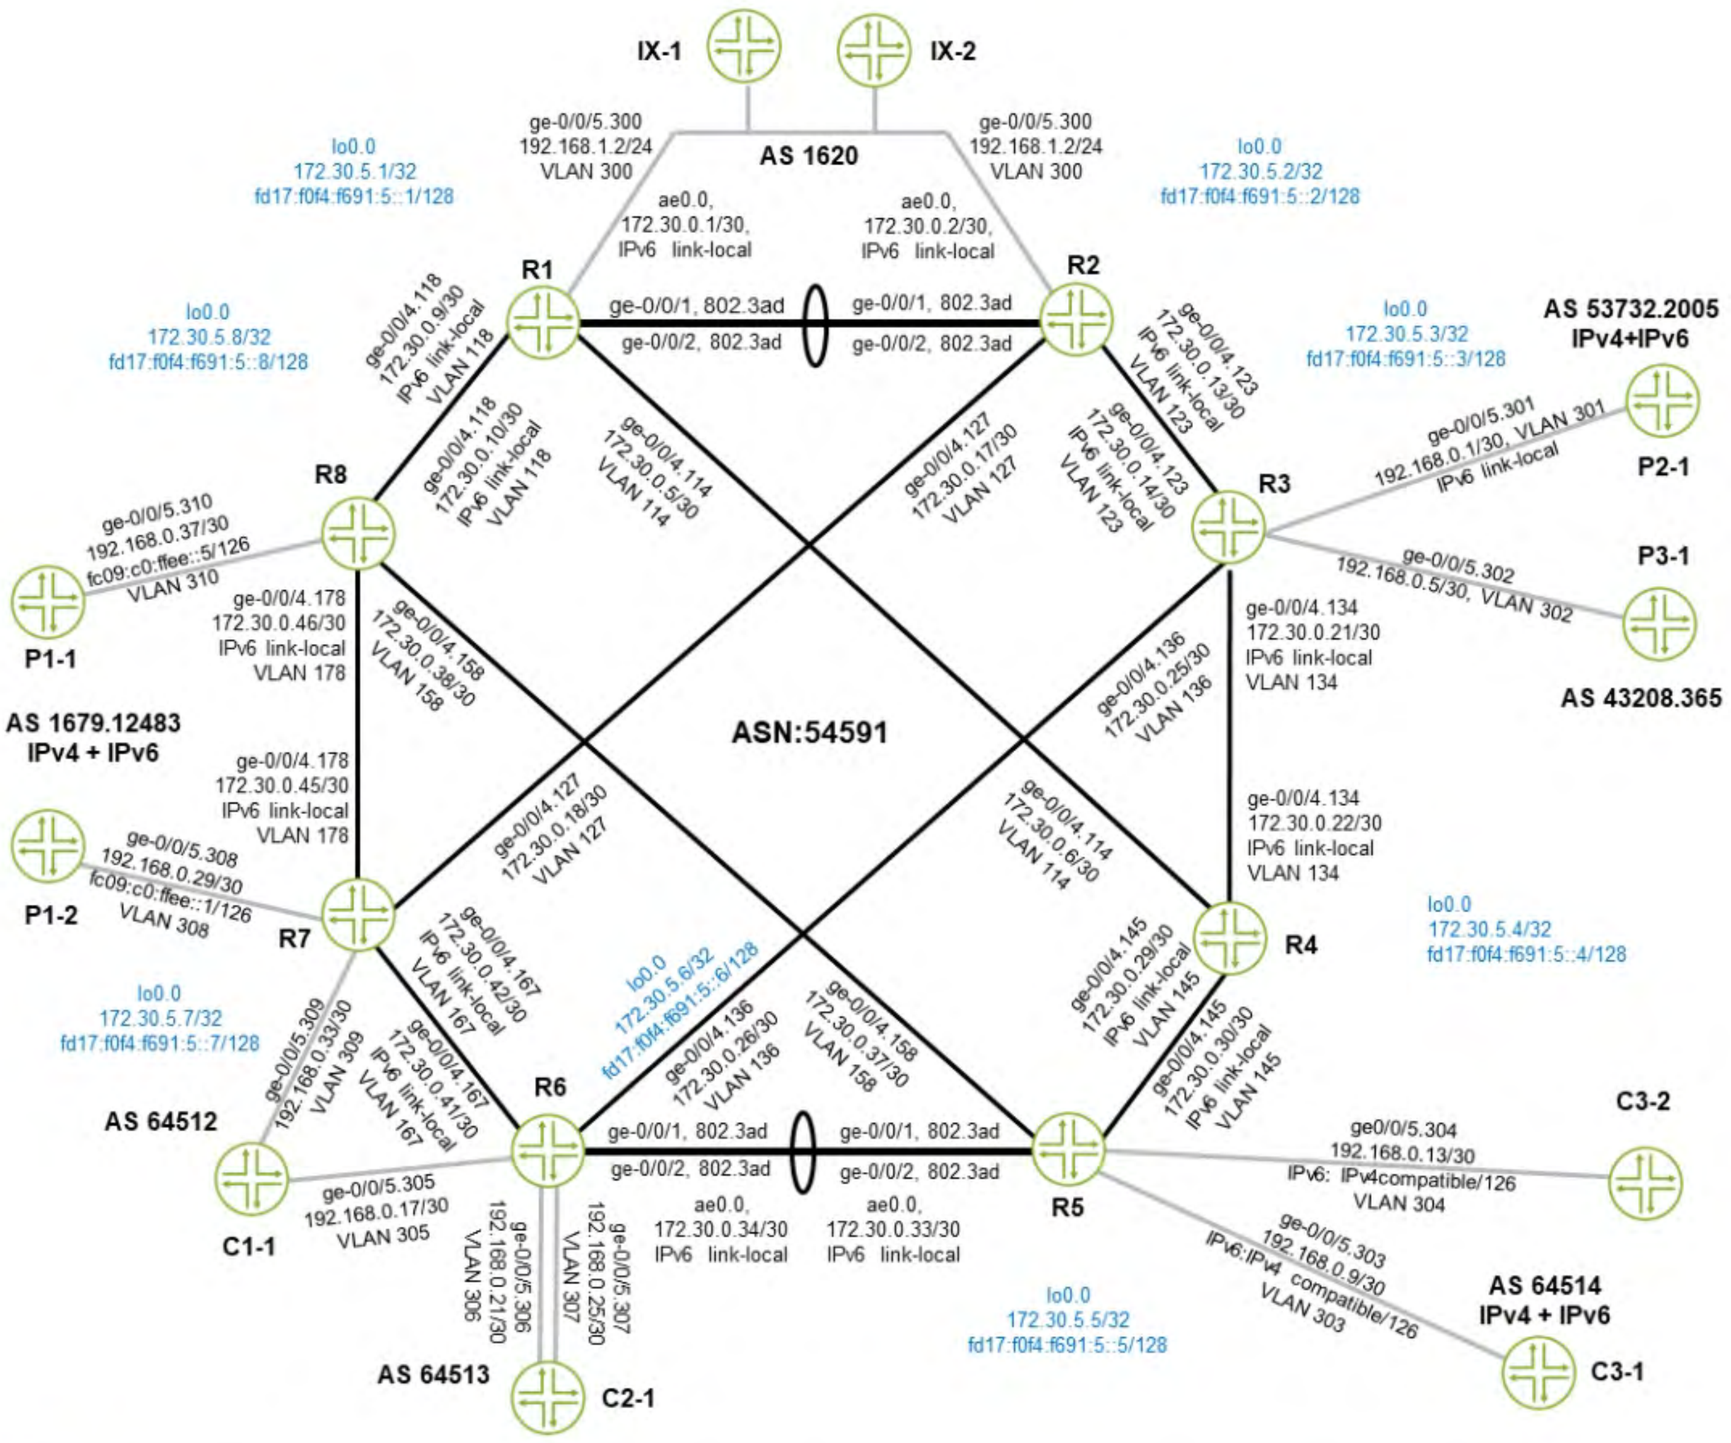


Figure 4

1) Configure the additional interfaces on your routers as indicated in Table 13. Configure the interface description.

Router	Interface Name	IP Address	IPv6 Address
R1	ge-0/0/5.300	192.168.1.1/24	
R2	ge-0/0/5.300	192.168.1.2/24	
R3	ge-0/0/5.301	192.168.0.1/30	link-local
	ge-0/0/5.302	192.168.0.5/30	
R5	ge-0/0/5.303	192.168.0.9/30	IPv4 compatible/126
	ge-0/0/5.304	192.168.0.13/30	IPv4 compatible/126
R6	ge-0/0/5.305	192.168.0.17/30	
	ge-0/0/5.306	192.168.0.21/30	
	ge-0/0/5.307	192.168.0.25/30	
R7	ge-0/0/5.308	192.168.0.29/30	fc09:c0:ffee::1/126

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/5.309	192.168.0.33/30	
R8	ge-0/0/5.310	192.168.0.37/30	fc09:c0:ffee::5/126

Table 13

2) Configure IPv4 EBGP sessions as shown in Figure 7.
3) Ensure that all the EBGP session state changes are logged to syslog.
4) Make sure that both R1 and R2 peer with both IX-1 and IX-2 routers. The IX-1 peering address is 192.168.1.3 and IX-2 is 192.168.1.4.
5) Use loopback interface peering for R6 to C2-1 session. Make sure that a single interface failure of the R6's ge-0/0/5.306 or ge-0/0/5.307 interfaces will not break the EBGP session down. Use RIP protocol to get the C2-1 loopback address.
6) Configure R5 to load balance over the two EBGP sessions to C3-1 and C3-2.
7) Make sure that no more than 20 prefixes are accepted from C1-1. If this limit is exceeded the session should be torn down and remain down for 3 minutes.
8) Configure native IPv6 EBGP peering with the P1 and P2 peers. Use link-local address for the session at R3. Find out the P2-1 IPv6 link-local address by using router monitoring tools.
9) Configure the IPv4 EBGP sessions to C3 to support IPv6 routing.
10) All routes received from customers C1 and C3 should be damped in case of flapping. Modify three damping parameters to make C1 damping more aggressive.
11) Make sure that all IPv4 routes received by all ASBRs over EBGP present in all other routers' routing tables.
12) Make sure that R1 and R2 do not use policy to resolve the BGP Next Hop problem.

Task 3: Routing Policies

In this task, you configure BGP routing policies to get precise handling of IPv4 routing exchanges across your AS.

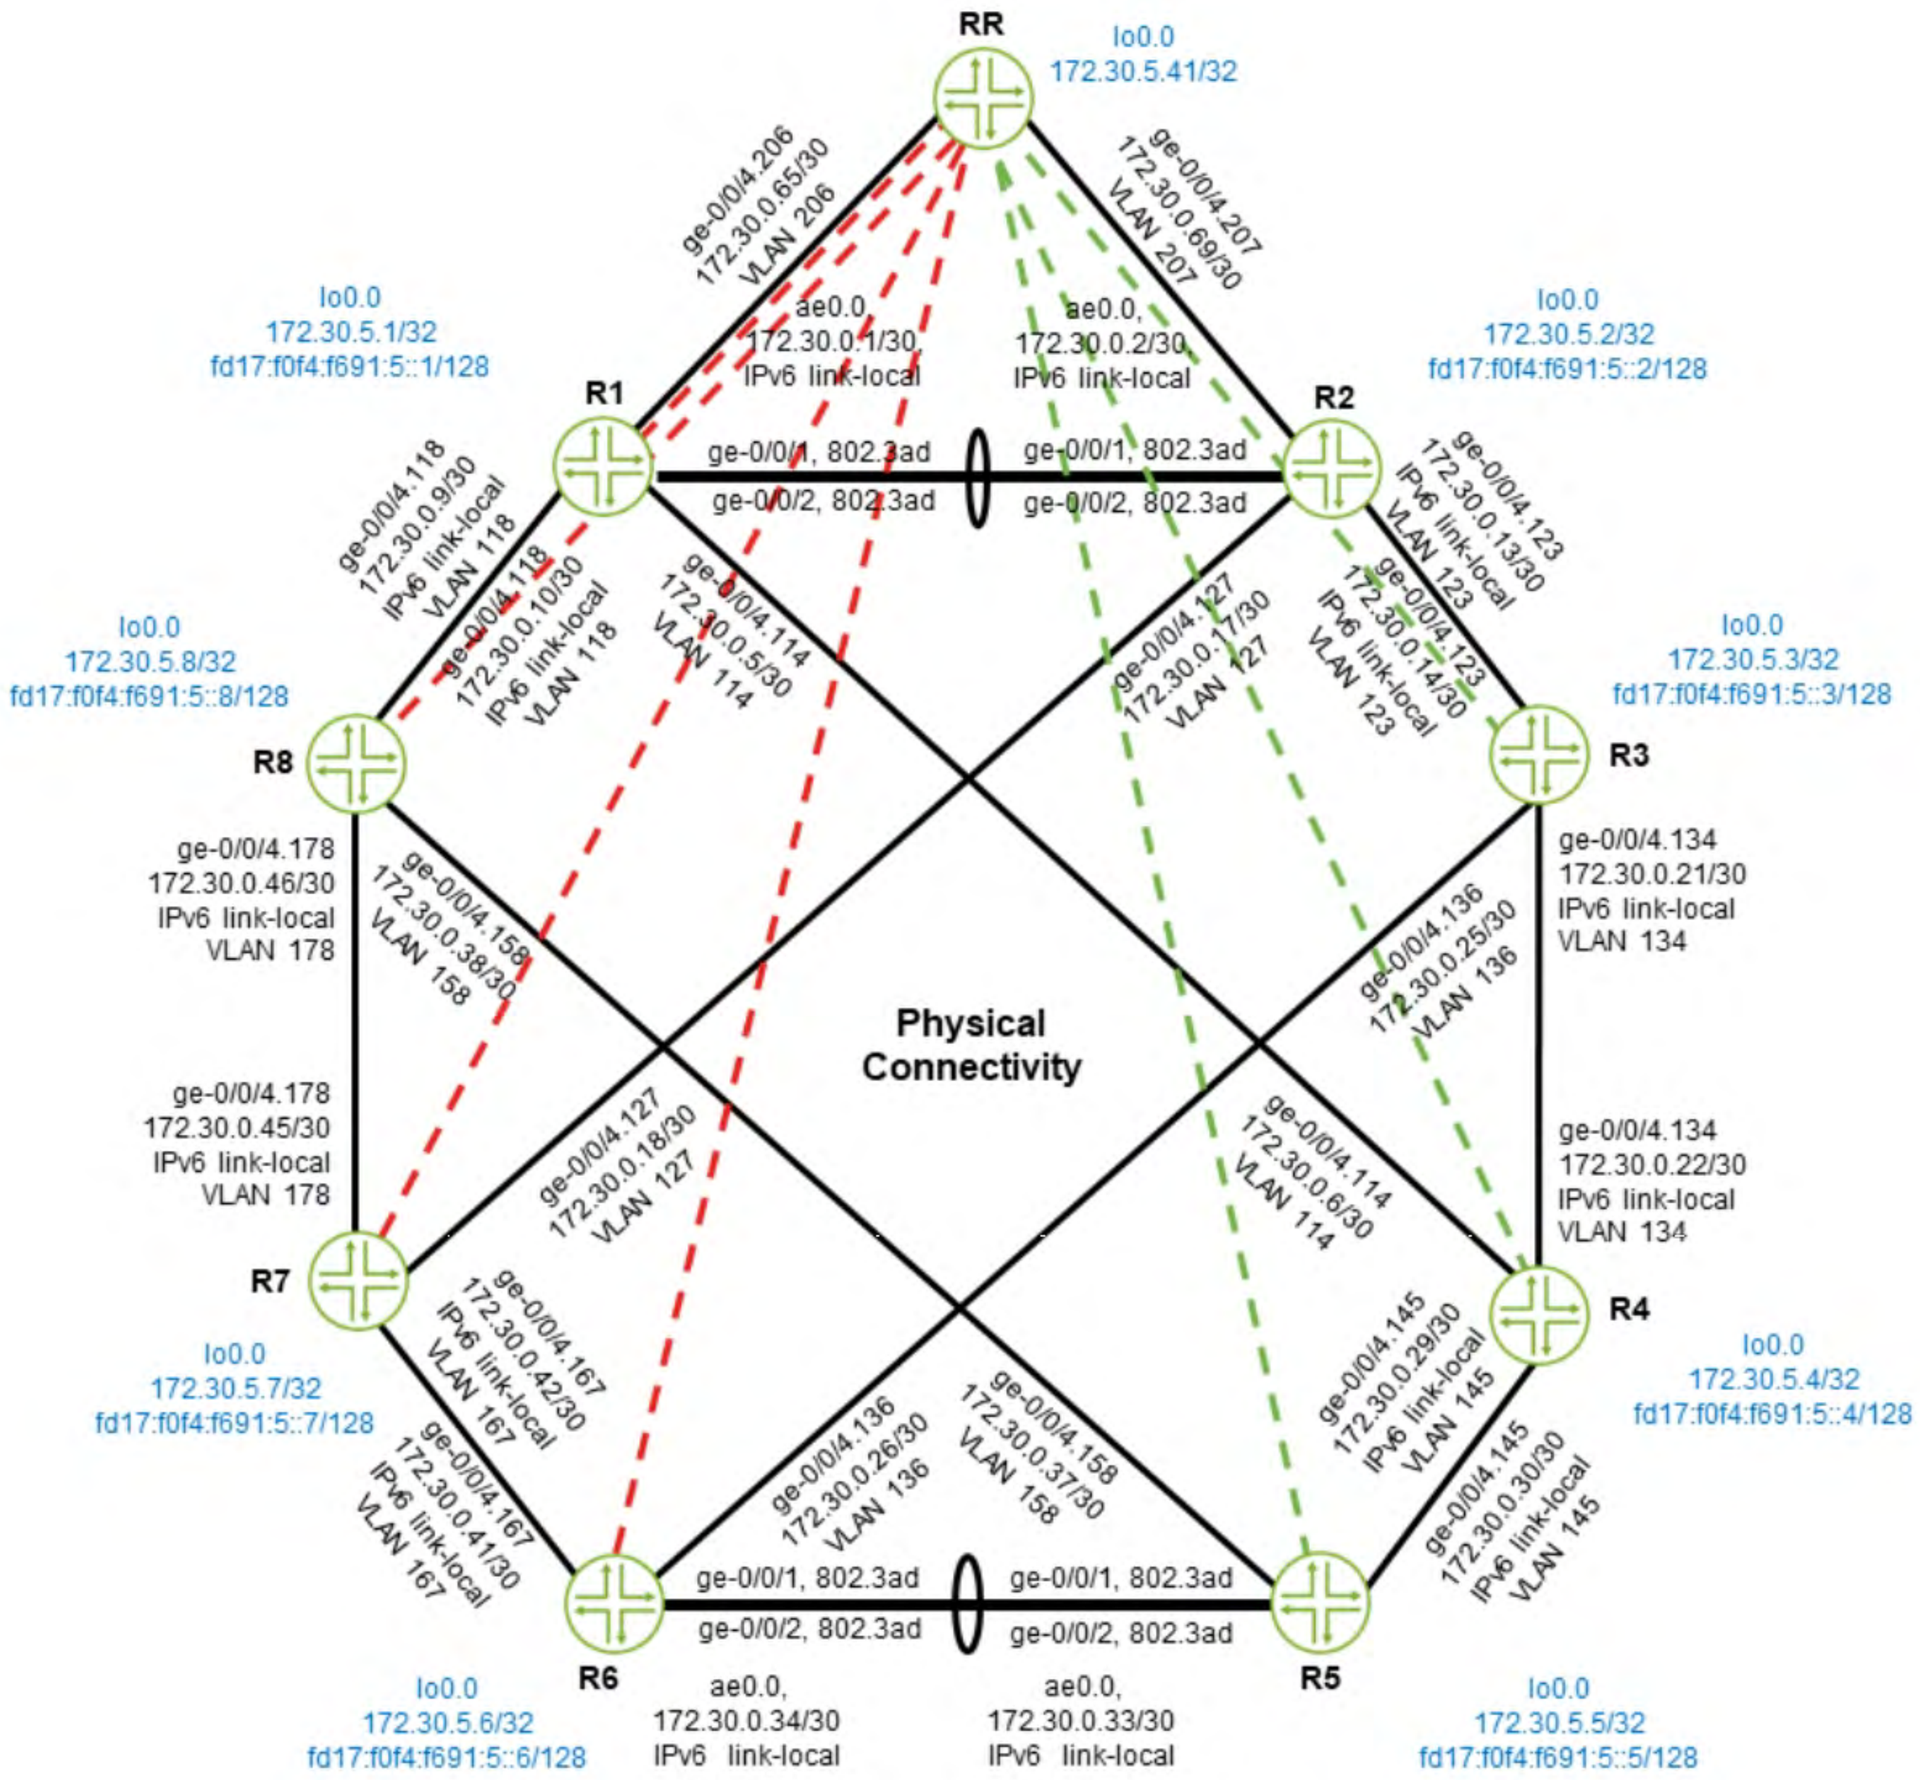
1) Make sure that the customer C1, C2 and C3 IPv4 routes are advertised to all EBGP peers.
2) Make sure that routes received from IX-1 or IX-2 are not advertised to P1 AS and vice versa.
3) Do not accept any IPv4 prefixes that are not originated in P1 AS from the P1 neighbors.

© 2020 Juniper Networks, Inc. All rights reserved.

4) Make sure that routes received from IX routers are less preferred than the same routes learned from either of P1, P2 or P3 peers.
5) Advertise only the default route to customer C2.
6) If a route is learned directly from a customer (C1, C2 or C3), it should be preferred to the same route learned from any other peer, however if a customer advertises a route with a community of "<Customer AS>:90" the route should be less preferred.
7) Do not accept IPv4 routes that have a mask shorter than /8 or longer than /24 from anywhere. You may accept routes with mask /32 originated in AS 43208.365.
8) Do not accept the 0.0.0.0 route with any mask length from any of the peers or customers.
9) Make sure that you use standard communities to identify IPv4 routes received from any of your neighboring AS's.
10) Advertise a single summary IPv4 route that aggregates your AS local routes including the RIP and OSPF routes to all your EBGP peers except C2.
11) Advertise parts of your AS summary route to P1 neighbors such as to achieve equal per-prefix load balancing for the traffic entering your AS from the P1 AS. When advertising these parts make sure that P1 does not re-advertise them outside of its AS using a well-known community.
12) Make sure that R8 is the preferred exit point for P1 destinations.
13) Make sure that R6 is preferred for both inbound and outbound traffic for the C1 customer.
14) Make sure that IX peers prefer routes advertised by R1 router.
15) Make sure that if a customer advertises an IPv4 route with a community of "<Customer AS>:666" the traffic to that destination is black holed.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 4: IBGP and Route Reflection



In this task, you will redesign your IBGP network to use route reflection instead of confederation. There is an extra virtual router referred to as RR that will act as route reflector “on a stick” in your network.

Note: The route reflector is configured on a stand-alone router. You can reach the router at its OoB management port at 10.10.1.19 address. Feel free to modify the RR settings as needed.

© 2020 Juniper Networks, Inc. All rights reserved.

Note: Assume the route reflector does not support 4-byte AS numbers.

1) Remove all IBGP settings.
2) Configure the RR facing interfaces at R1 and R2 as indicated in Table 14. Set the interfaces description.

Router	Interface Name	IP Address
R1	ge-0/0/4.206	172.30.0.65/30
R2	ge-0/0/4.207	172.30.0.69/30

Table 3

3) Configure IBGP route reflection. There must be two clusters and any client may be a member of one cluster only.
4) Clients can only have IBGP sessions with the route reflector.
5) Make sure that IBGP sessions use loopback interface peering. The RR loopback address is 172.30.5.41.
6) Make sure that the route reflection does not result in suboptimal routing.
7) Configure MD5 authentication for all the IBGP sessions.
8) Enable BFD neighbor continuity checking for all the IBGP sessions. Use the default BFD timer values.
9) Ensure that all the IBGP session state changes are logged to syslog.
10) No unresolved IPv4 routes are allowed anywhere.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 4: MPLS Configuration

In this chapter, you will create core MPLS network. The chapter tasks include configuration of LDP-signaled LSPs, RSVP-signaled LSPs, traffic engineering, traffic protection and optimization, and LDP tunneling.

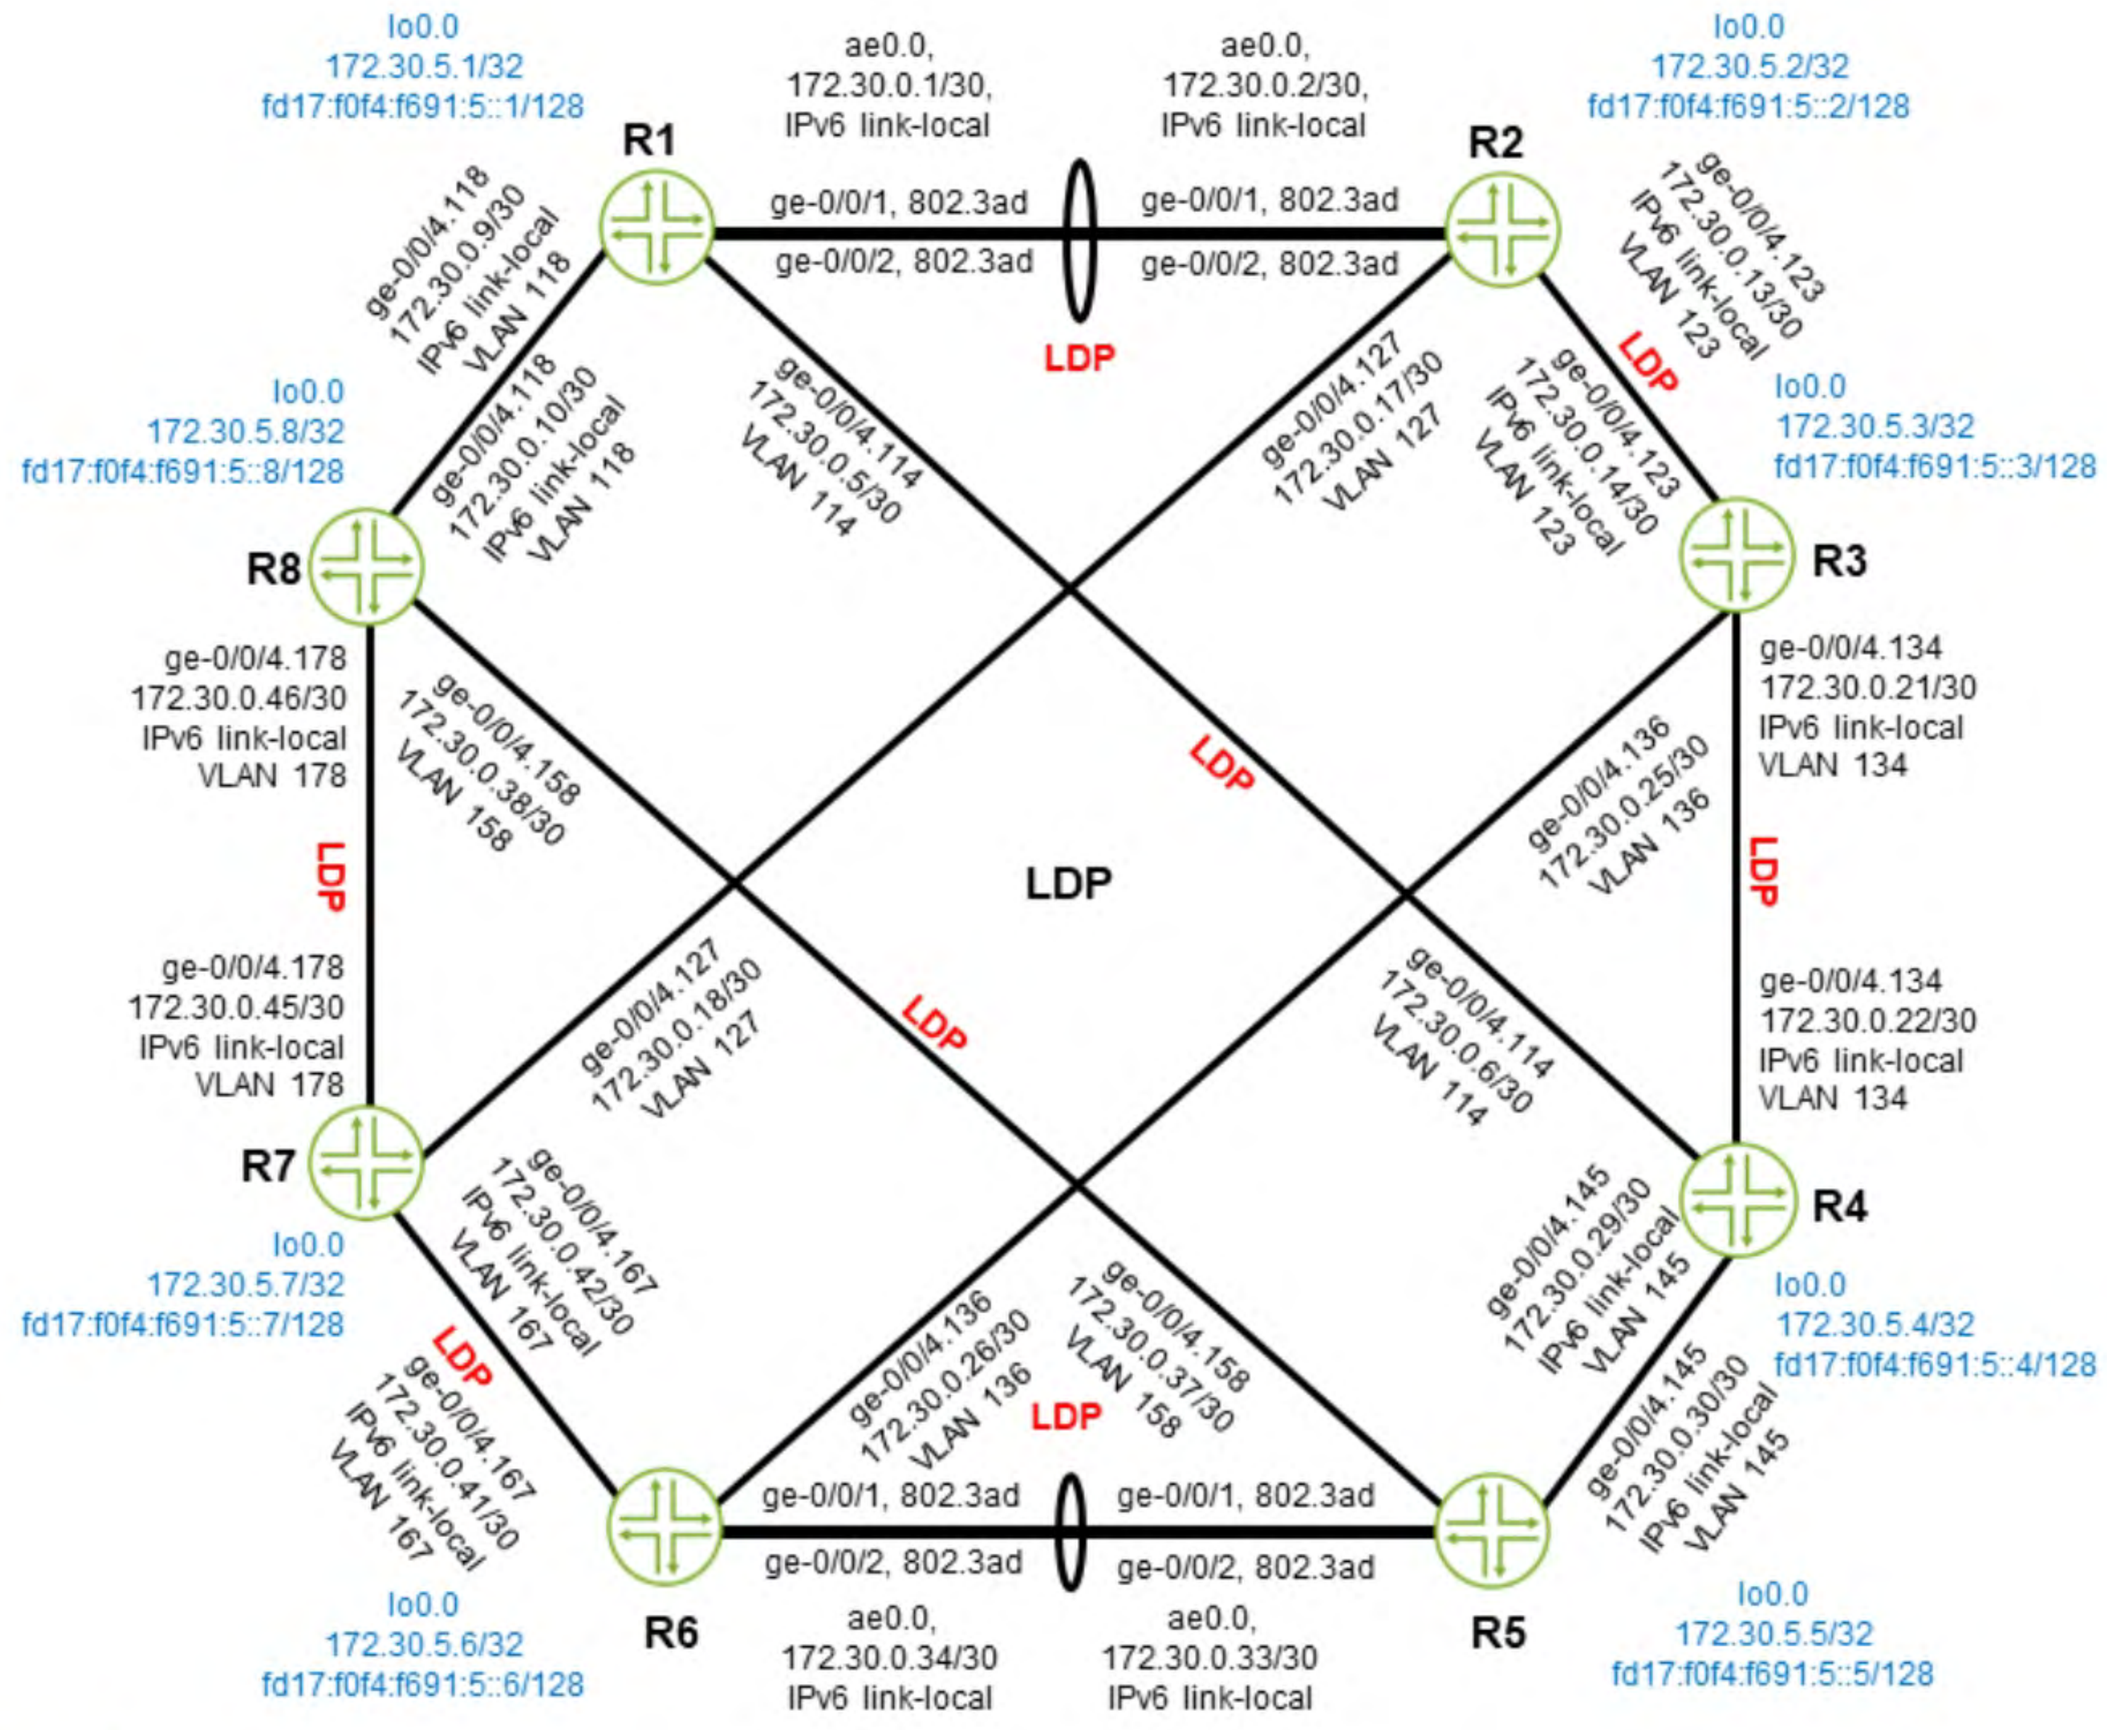


Figure 5

© 2020 Juniper Networks, Inc. All rights reserved.

Task 1: LDP Configuration

In this task, you configure LDP-signaled MPLS LSPs.

1) Configure LDP as shown in Figure 8.
2) Configure MD5 authentication for all LDP sessions.
3) Configure IS-IS to track the LDP operational status on all LDP-enabled interfaces.
4) Configure R1 and R2 to inject the IX facing subnet into LDP. Make sure that each FEC advertised by R1 or R2 is reachable by a separate LSP.
5) Make sure that LDP LSPs show the same metric as the IGP paths they follow.
6) Make sure that LDP labels are popped by the egress routers.

Note: You will join the LDP islands with LDP tunneling in the RSVP configuration tasks.

Task 2: RSVP Configuration

In this task, you configure RSVP-signaled MPLS LSPs, implement RSVP traffic engineering, configure RSVP optimization, LDP tunneling, and LSP load balancing.

1) Enable RSVP on all routers' core facing interfaces.
2) Configure all RSVP-enabled interfaces but the ae0 Ethernet bundles to allow 333 Mbps of bandwidth reservation.
3) Configure link administrative groups as shown in Table 15.

© 2020 Juniper Networks, Inc. All rights reserved.

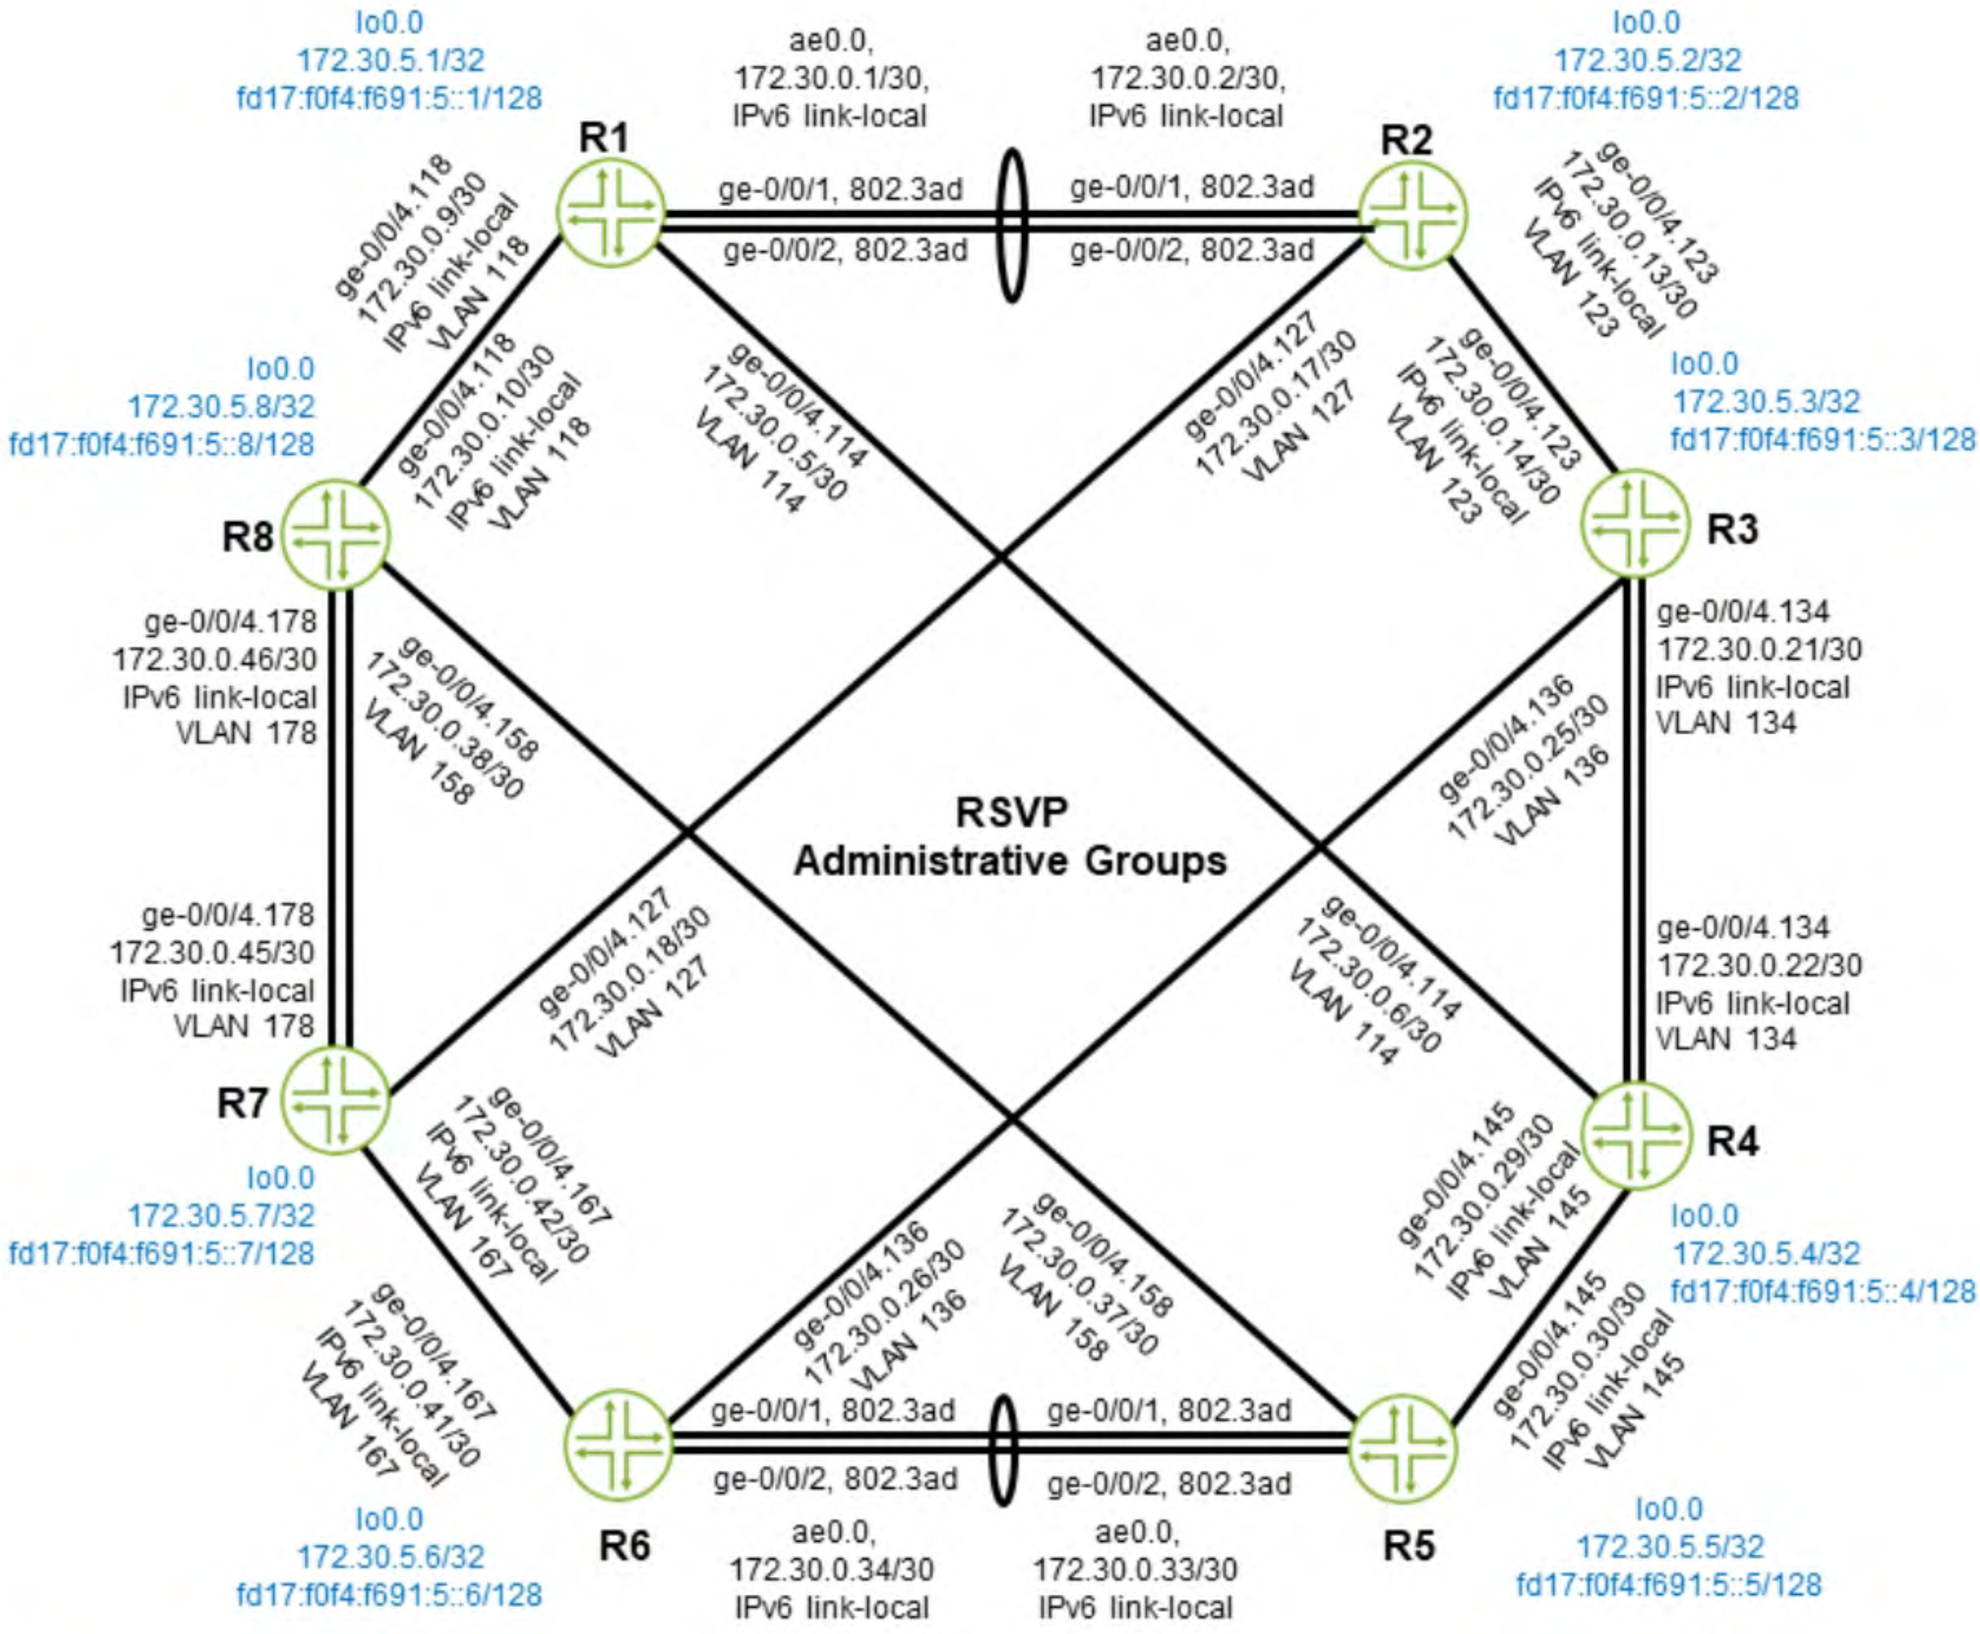


Figure 6

Router	Interface	Admin Group
R1	ge-0/0/4.114	green
	ge-0/0/4.118	red
	ae0.0	green, red
R2	ge-0/0/4.127	green
	ge-0/0/4.123	red
	ae0.0	green, red
R3	ge-0/0/4.134	green, red
	ge-0/0/4.136	green

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/4.123	red
R4	ge-0/0/4.134	green, red
	ge-0/0/4.114	green
	ge-0/0/4.145	red
R5	ge-0/0/4.158	green
	ge-0/0/4.145	red
	ae0.0	green, red
R6	ge-0/0/4.136	green
	ge-0/0/4.167	red
	ae0.0	green, red
R7	ge-0/0/4.178	green, red
	ge-0/0/4.127	green
	ge-0/0/4.167	red
R8	ge-0/0/4.178	green, red
	ge-0/0/4.158	green
	ge-0/0/4.118	red

Table 4

4) Configure RSVP-signaled LSPs as shown in Table 16.

© 2020 Juniper Networks, Inc. All rights reserved.

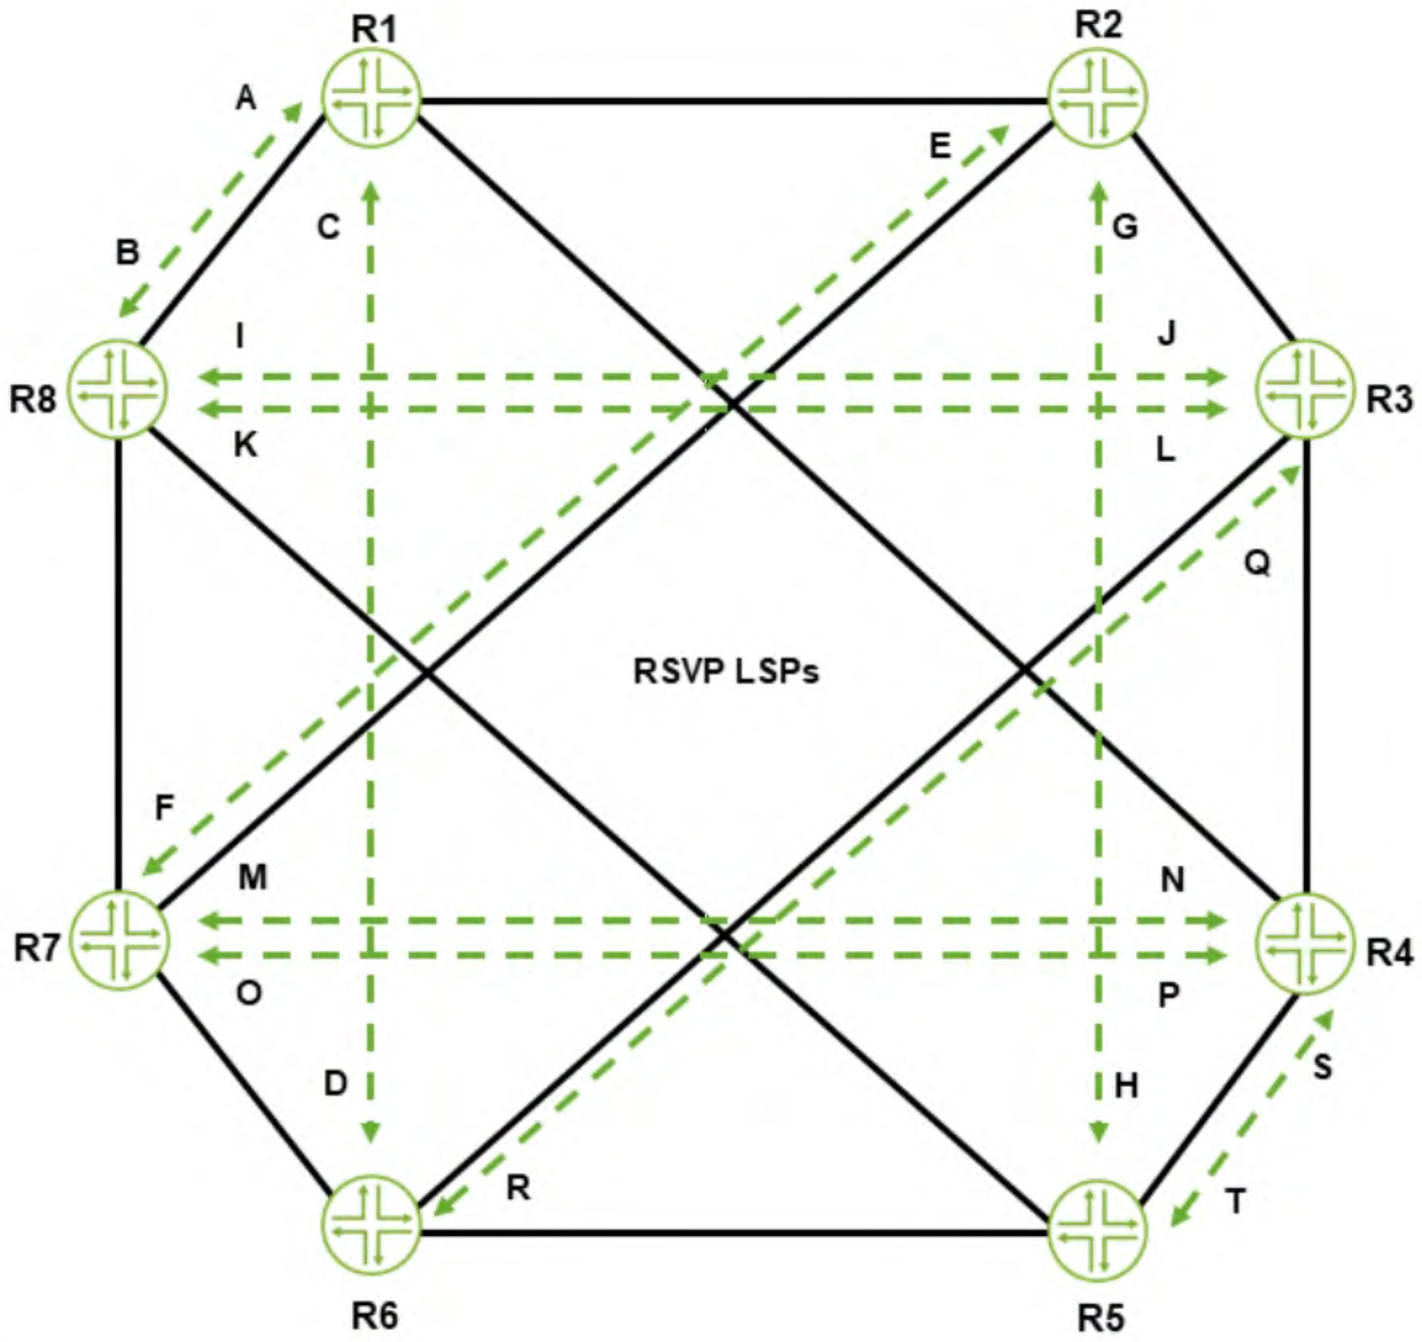


Figure 7

Ingress	Egress	LSP ID
R1	R8	A
R1	R6	C
R2	R7	E
R2	R5	G
R3	R8	J
R3	R8	L
R3	R6	Q
R4	R7	N
R4	R7	P
R4	R5	S
R5	R2	H
R5	R4	T
R6	R1	D

© 2020 Juniper Networks, Inc. All rights reserved.

R6	R3	R
R7	R2	F
R7	R4	M
R7	R4	O
R8	R1	B
R8	R3	I
R8	R3	K

Table 5

Note: The LSP IDs are used here as reference names only.

5) Configure MD5 authentication for all RSVP sessions.
6) Enable BFD continuity checking for all the RSVP sessions.
7) Make sure that LSPs E, F, Q and R use only links belonging to “red” administrative group.
8) Make sure that LSPs A, B, S and T use only links belonging to “green” administrative group.
9) Configure LSPs I and K, and LSPs J and L so that they use two distinct physical paths to the egress node. The paths should take 3 hops each. You may not use administrative groups in this step.
10) Configure LSPs M and O, and LSPs N and P so that they use two distinct physical paths to the egress node. LSPs M and O should use only “green” links and LSPs N and P should use only “red” links.
11) Configure all LSPs except A, B, S, T to reserve 60Mbps of bandwidth.
12) Configure LSPs A, B, S, T to automatically re-signal the LSP once in 48 hours based on the average bandwidth usage. Make sure that the LSPs can use not less than 30Mbps and not more than 120Mbps.
13) Configure LSPs A, B, E, F, I, J, Q, R, S, T to ensure that they have higher priority for bandwidth reservation than the remaining LSPs.
14) Make sure that if LSPs K, L, O, P must be preempted, the ingress router will attempt to re-signal the LSP before tearing it down.
15) Configure automatic optimization for the LSPs I, J, K, L, M, N, O, P. Set the optimize timer to 8 hours. Make sure that the ingress routers attempt to re-signal the LSP before tearing it down.
16) Make sure that R5 and R6 prefer RSVP LSPs as the next hops for IPv4 BGP routes advertised by IX peers.
17) Configure LDP tunnels between R3 and R8, and R4 and R7. Make sure that any router in your AS has an LDP-signaled LSP to any other router.

© 2020 Juniper Networks, Inc. All rights reserved.

18) Make sure that IPv4 traffic at R8 from P1 to P2 uses LSP I and traffic from P1 to P3 uses LSP K.
19) Configure per flow load balancing over LSPs N and P. Vice versa configure per flow load balancing over LSPs M and O.
20) Make sure that MPLS paths in your network are hidden from external traceroute utilities.

Task 3: RSVP Protection

In this task, you implement different LSP protection mechanisms.

1) Configure a backup protection path for all RSVP-signaled LSPs but K, L, O, P.
2) Make sure that for the LSPs C, D, G, H the protection path is established in advance, before the primary path fails.
3) Configure all the protection paths to inherit the bandwidth settings from the primary ones. Make sure that for LSPs C, D, G, H the bandwidth is shared between the primary and protection paths.
4) Configure LSPs E, F, Q and R to not revert to the primary path if a switchover to the protection path occurred.
5) Configure LSPs C, D, G, H to use fast reroute protection mechanism. Make sure that the detour LSPs do not inherit either bandwidth or administrative group settings from the main LSP. The detour LSPs must transit not more than 5 hops.
6) Configure LSPs A, B, E, F, Q, R, S, T to use link protection mechanism.
7) Configure LSPs I, J, M, N to use link and node protection mechanism.

Task 4: IPv6 Tunneling with 6PE

This task focuses on 6PE implementation.

1) Enable IPv6 over MPLS tunneling in your network using 6PE technique. You may not use native IPv6 forwarding anywhere within your AS for transit packets.
2) You may not have any MPLS LSPs on the route reflector. A static route is allowed on the RR if needed.
3) Make sure that end-to-end IPv6 communication is provided among C3, P1 and P2 over your MPLS network.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 5: Layer 3 VPN Configuration

In this chapter, you implement L3VPNs. The tasks include L3VPN configuration with customers running either OSPF or BGP, dual-homed customer sites, customer Internet access, multicasting in VPNs and IPv6 tunneling with 6VPE.

Task 1: Layer 3 VPN Configuration

In this task, you deploy L3VPN for with customers running either OSPF or BGP.

1) Configure additional interfaces on your routers as indicated in Table 17. Set the interfaces description.

Router	Interface Name	IP Address	IPv6 Address
R1	ge-0/0/5.311	192.168.0.41/30	
	ge-0/0/5.312	192.168.0.45/30	
	ge-0/0/5.313	192.168.0.49/30	
	lo0.1	172.30.5.9/32	
	lo0.2	172.30.5.10/32	
R2	ge-0/0/5.314	192.168.0.53/30	
	ge-0/0/5.315	192.168.0.57/30	
	ge-0/0/5.316	192.168.0.61/30	
	lo0.1	172.30.5.13/32	
	lo0.2	172.30.5.14/32	
R3	ge-0/0/5.317		fc09:c0:ffee::9/126
	ge-0/0/5.318	192.168.0.69/30	
	lo0.1	172.30.5.17/32	
	lo0.2	172.30.5.18/32	fd17:f0f4:f691:5::12/128
R4	ge-0/0/5.319	192.168.0.73/30	
	ge-0/0/5.320	192.168.0.77/30	
	lo0.1	172.30.5.21/32	
	lo0.2	172.30.5.22/32	
R5	ge-0/0/5.321	192.168.0.81/30	
	lo0.1	172.30.5.25/32	
R6	ge-0/0/5.322	192.168.0.85/30	
	lo0.1	172.30.5.29/32	
R7	ge-0/0/5.323	192.168.0.89/30	

© 2020 Juniper Networks, Inc. All rights reserved.

R8	lo0.1	172.30.5.33/32	
	ge-0/0/5.324	192.168.0.93/30	
	ge-0/0/5.325		fc09:c0:ffee::d/126
	lo0.1	172.30.5.37/32	
	lo0.2	172.30.5.38/32	fd17:f0f4:f691:5::26/128

Table 6

© 2020 Juniper Networks, Inc. All rights reserved.

Configure L3VPNs as shown in Figure 11. Table 18 specifies the L3VPN details.

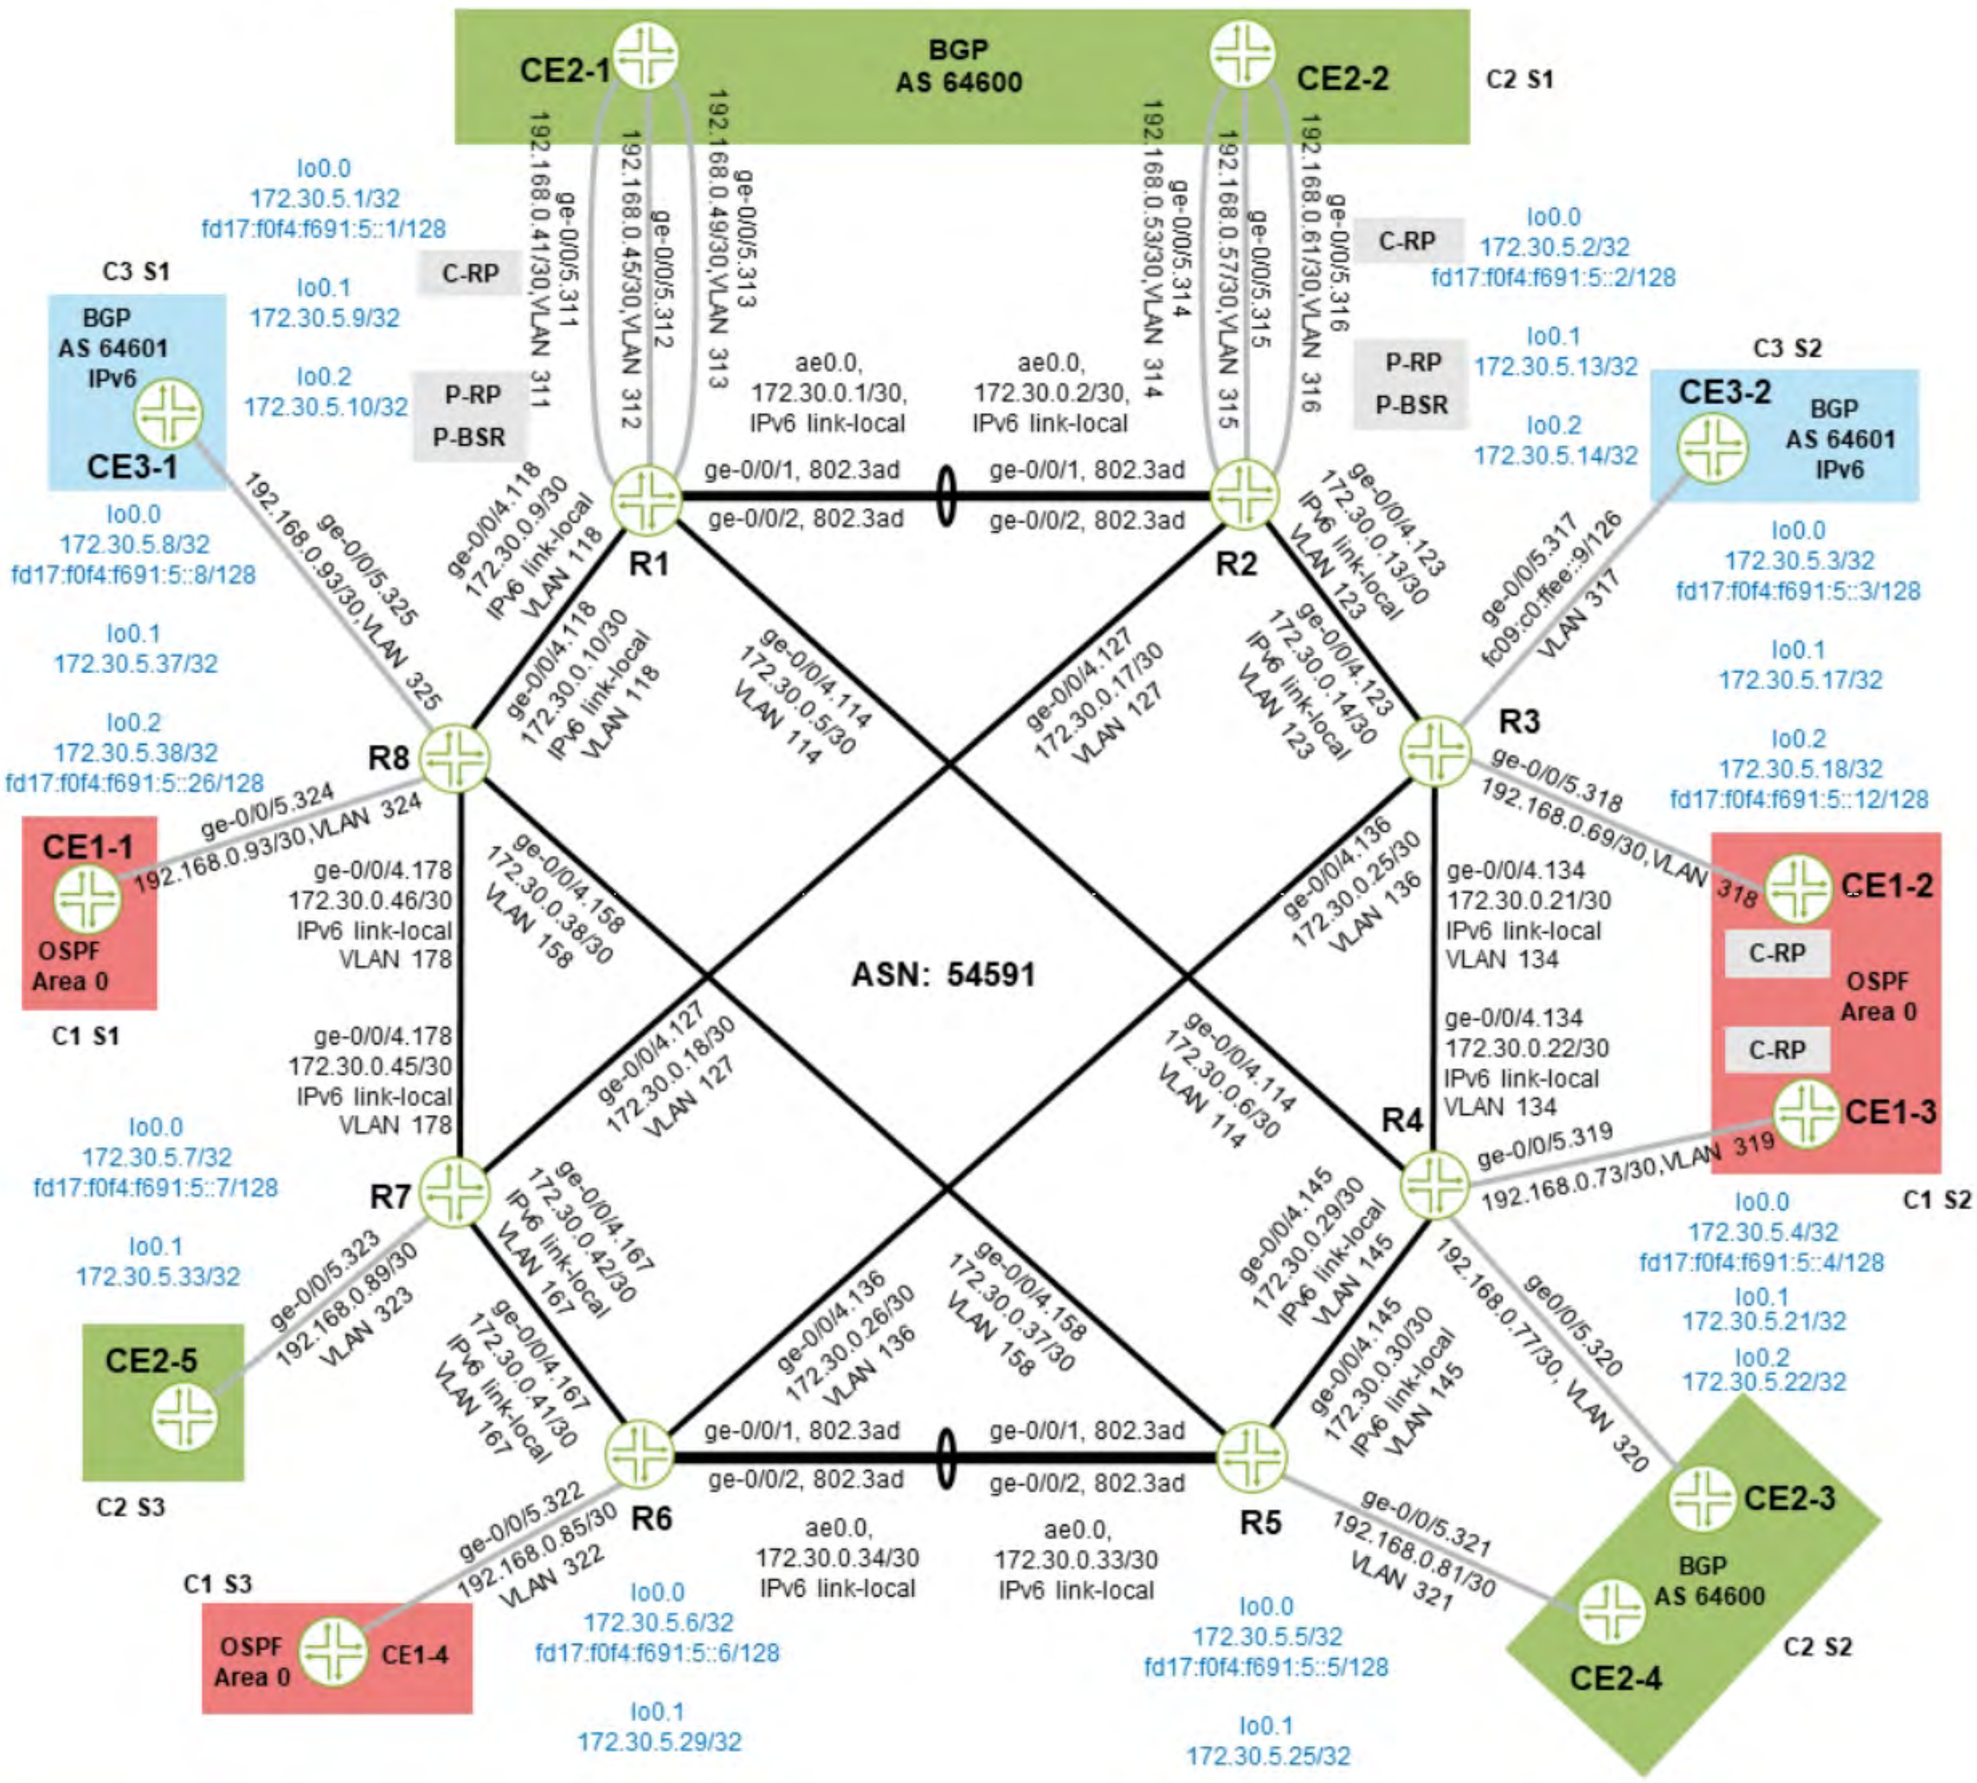


Figure 8

© 2020 Juniper Networks, Inc. All rights reserved.

Customer	Site	Router	PE-CE Protocol	Protocol details
C1	S1	CE1-1	OSPF	Area 0.0.0.0
	S2	CE1-2	OSPF	Area 0.0.0.0
		CE1-3	OSPF	Area 0.0.0.0
	S3	CE1-4	OSPF	Area 0.0.0.0
C2	S1	CE2-1	BGP	AS 64600
		CE2-2	BGP	AS 64600
	S2	CE2-3	BGP	AS 64600
		CE2-4	BGP	AS 64600
	S3	CE2-5	BGP	AS 64600

Table 7

2) You might not have any MPLS LSPs on route reflector. A static route is allowed on the RR if needed.
3) Make sure that the customer C1 OSPF area 0 appears as a contiguous area without ABRs.
4) Customer C1 has some backdoor OSPF connections but prefers that your MPLS network would be used for traffic forwarding between the customer sites.
5) Make sure that your MPLS network can be used as a backup path between CE1-2 and CE1-3.
6) Make sure that once customer C1 disables its backdoor connections any of the R3 or R4 PE failure will not result in any of the customer sites become isolated.
7) Customer C2 requires that the customer site S1 is used as a central transit site for all traffic exchanges among all the customer sites in a hub-and-spoke fashion.
8) Make sure that if a route is originated in customer C2 site S1 or S2, it is never advertised back to the same site.
9) Make sure that PE-CE link subnets in customer C2 VPN are advertised to the customer remote VPN sites.
10) Make sure that all PE routers receive only the routes with those targets that they specifically request for.
11) Allow local communication between customer C1 site S2 and customer C2 site S2 at R4. Make sure that the routes exchanged between the local VRFs are not advertised to any of the remote PE routers.
12) Customer C1 must be provided with Internet access at the customer site S2 using single customer-facing interface. Make sure that any of the R3 or R4 failure will not have customer C1 site S2 isolated from the Internet.

© 2020 Juniper Networks, Inc. All rights reserved.

Note: The customer IP ranges are assumed to be globally routable or NATed outside of your network.

13) Customer C2 must be provided with Internet access at the customer site S1, using a dedicated interface at both R1 (ge-0/0/5.313) and R2 (ge-0/0/5.316) routers. All other customer sites should be able to reach the Internet via the site S1.

Task 2: Multicast in Layer 3 VPNs

In this task, you implement multicast and Next Generation multicast in the L3VPNs.

Note: Customer C2 uses 239.0.0.0/24 multicast range.

1) Enable PIM sparse mode ASM in your AS. Make sure that R1 and R2 act as anycast RPs. You may not use MSDP in your network.
2) Use bootstrap RP mapping in your network. Make sure that R1 is the active BSR and R2 will take over the BSR role if R1 fails.
3) Configure your network to use inet.2 table for multicast RPF.
4) Configure NG MVPN in customer C2 VPN. The customer site S1 acts as a sender site only and sites S2 and S3 as receiver sites. Make sure that P2MP RSVP-signaled LSP is used as the PMSI.
5) Customer C2 outsources its RP to your network. Make sure that your routers R1 and R2 act as the customer anycast RPs.
6) Enable selective PMSI's in customer C2 site S1 for multicast groups 239.0.0.1 and 239.0.0.2 from any source in range 172.31.64.0/21. Make sure that the site uses inclusive PMSI for the remaining multicast groups in the customer range.
7) Make sure that customer C2 site S1 inclusive PMSI establishes automatically using parameters defined in Table 19 and selective PMSI's establish automatically using parameters defined in Table 20. Set the selective PMSI's threshold to 100Mb. No more than 5 selective PMSI's may be signaled.
8) Make sure that the customer C2 receiver sites join only source-based multicast distribution trees.

© 2020 Juniper Networks, Inc. All rights reserved.

Parameter	Value
Bandwidth	30Mbps
Priority	better than the higher priority LSPs configured so far
Protection	link protection
Hop limit	5

Table 8

Parameter	Value
Bandwidth	60Mbps
Priority	same as for the inclusive PMSI
Protection	link protection
Hop limit	5

Table 9

Task 3: IPv6 Tunneling with 6VPE

In this task, you implement IPv6 tunneling with 6VPE.

1) Establish native IPv6 EBGP sessions with customer C3 CE routers at R3 and R8.
2) Provide customer C3 with traffic forwarding between the customer sites. You may not use native IPv6 IBGP peering in your network.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 6: Layer 2 VPN Services

In this chapter, you implement L2VPN, VPLS, and EVPN services in your network.

Task 1: Layer 2 VPN Configuration

In this task, you configure LDP- and BGP-signaled L2VPN services.

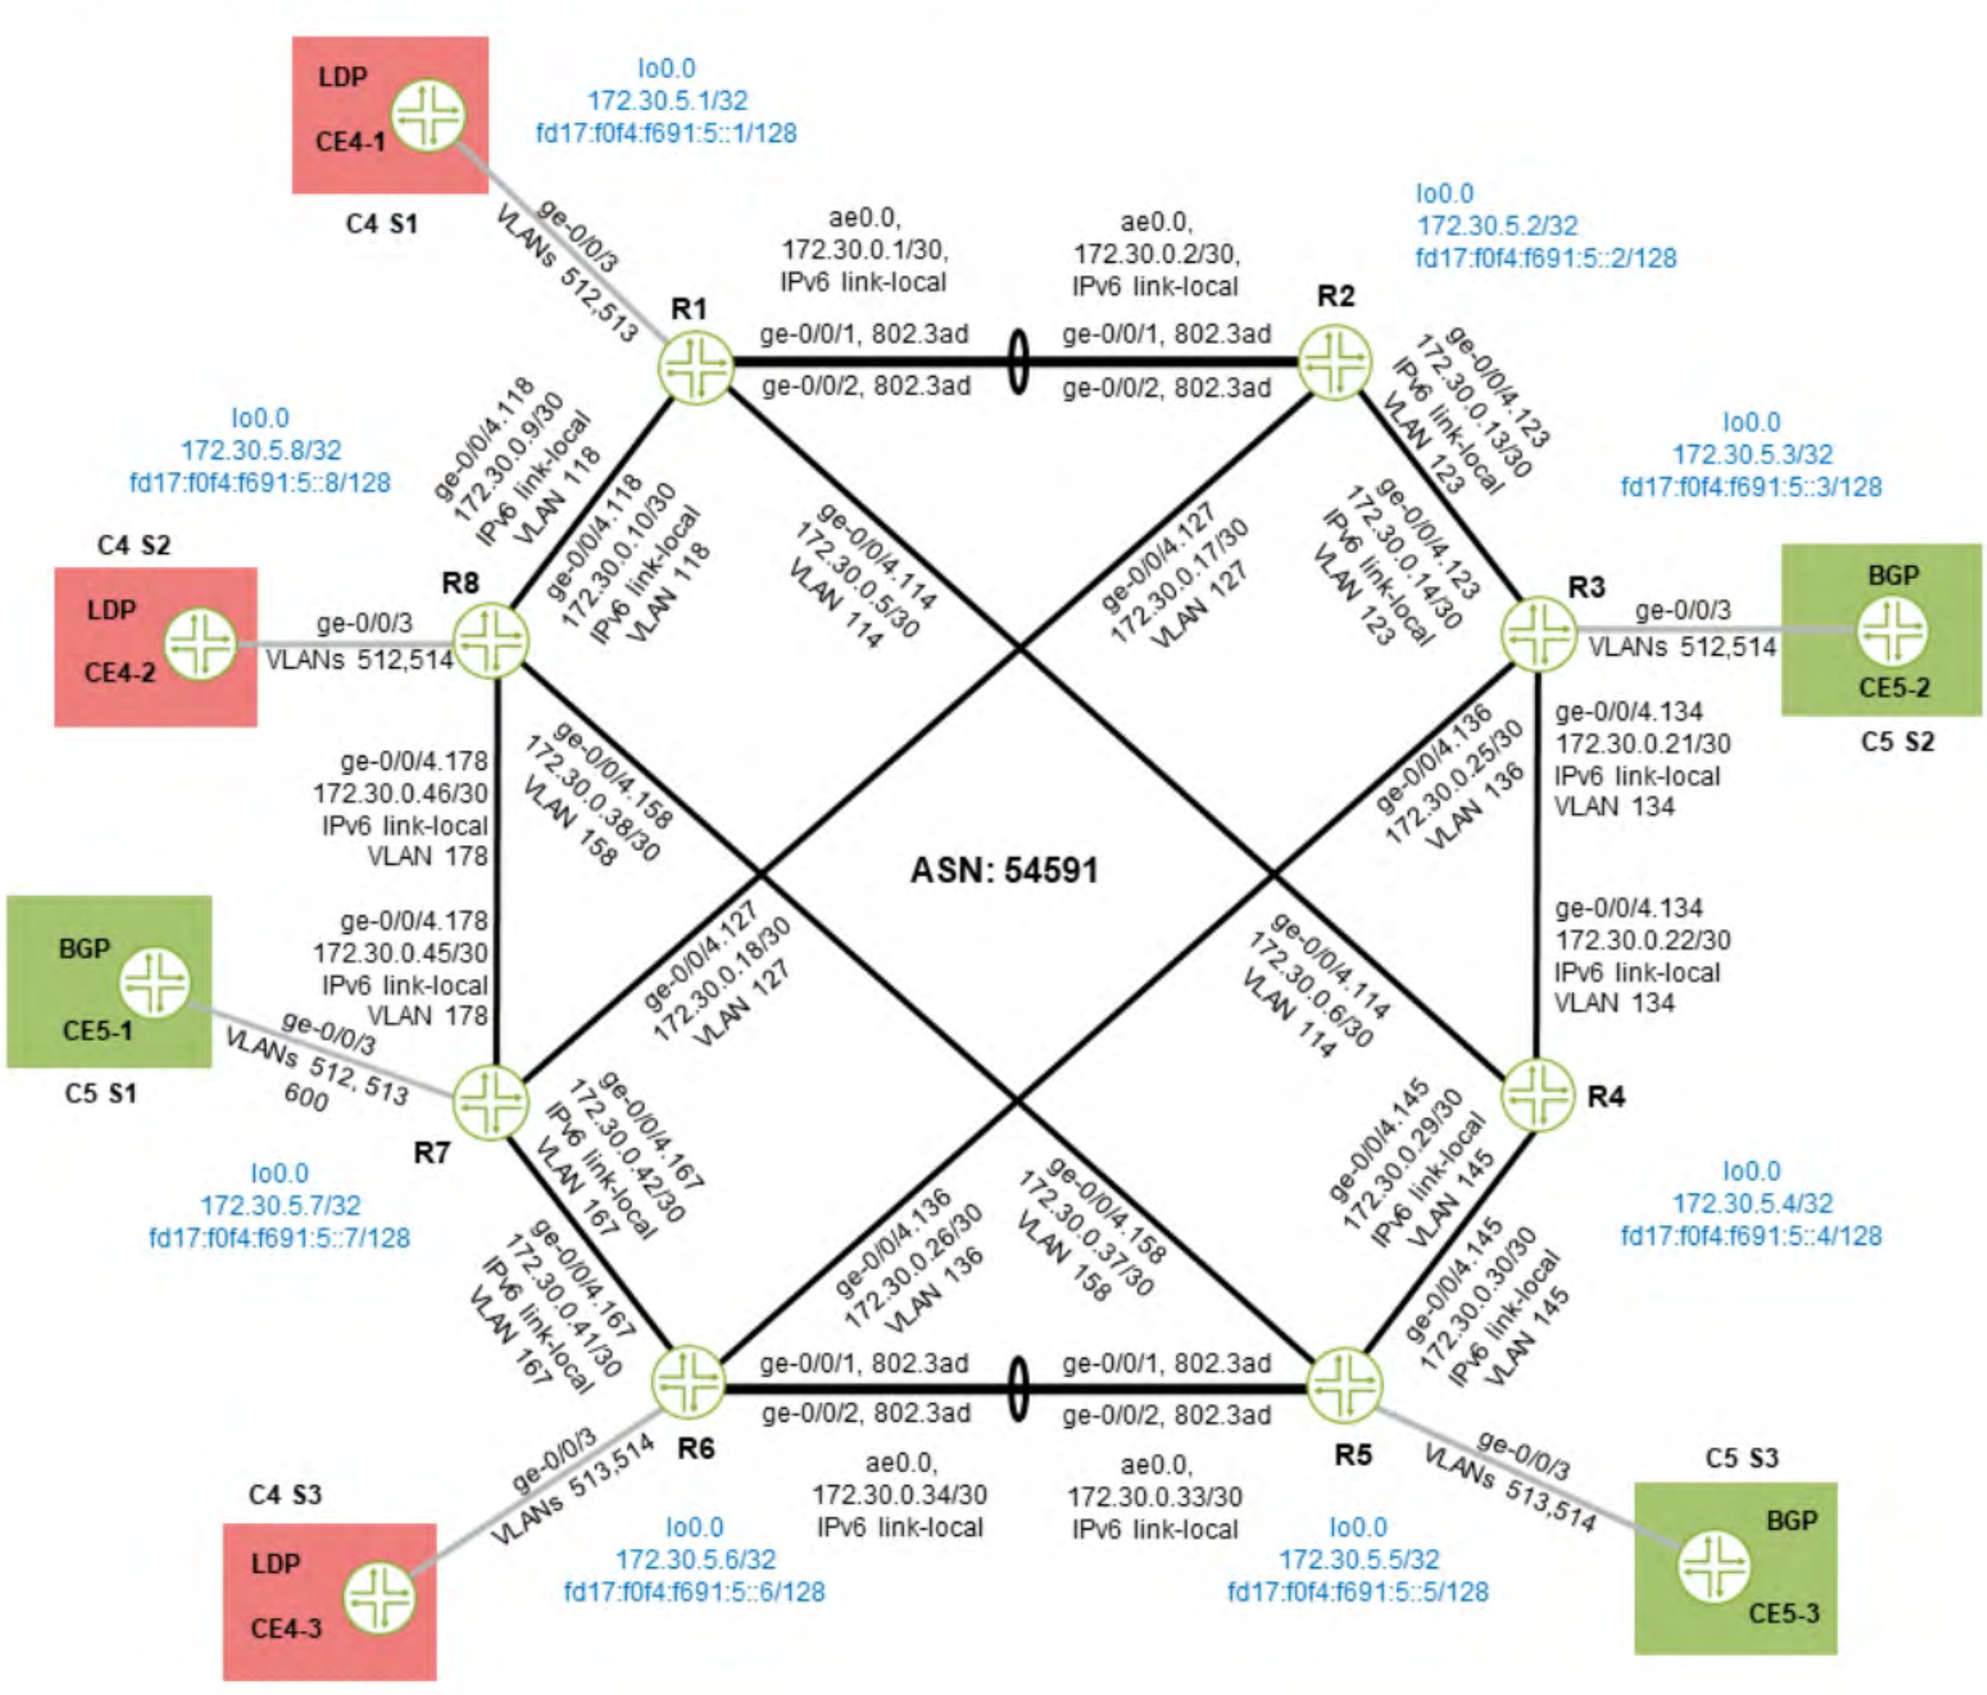


Figure 9

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure L2VPN as shown in Figure 12. Table 21 specifies the L2VPN details. Configure customer VLANs as shown in Table 22.

Customer	Site	Router	L2VPN signaling	CE facing interface
C4	S1	CE4-1	LDP	ge-0/0/3
	S2	CE4-2	LDP	ge-0/0/3
	S3	CE4-3	LDP	ge-0/0/3
C5	S1	CE5-1	BGP	ge-0/0/3
	S2	CE5-2	BGP	ge-0/0/3
	S3	CE5-3	BGP	ge-0/0/3
	S4	CE5-4	BGP	ge-0/0/3

Table 10

Customer	VLAN	Connection
C4	512	S1-S2
	513	S1-S3
	514	S2-S3
C5	512	S1-S2
	513	S1-S3
	514	S2-S3
	600	S1-S4

Table 11

2) Make sure that both customers' sites are fully meshed. The connection table is shown in Table 22.

Task 2: VPLS Configuration

In this task, you configure LDP- and BGP-signaled VPLS services, VPLS and L2VPN interworking, LDP and BGP VPLS interworking, and Internet access to VPLS customers.

© 2020 Juniper Networks, Inc. All rights reserved.

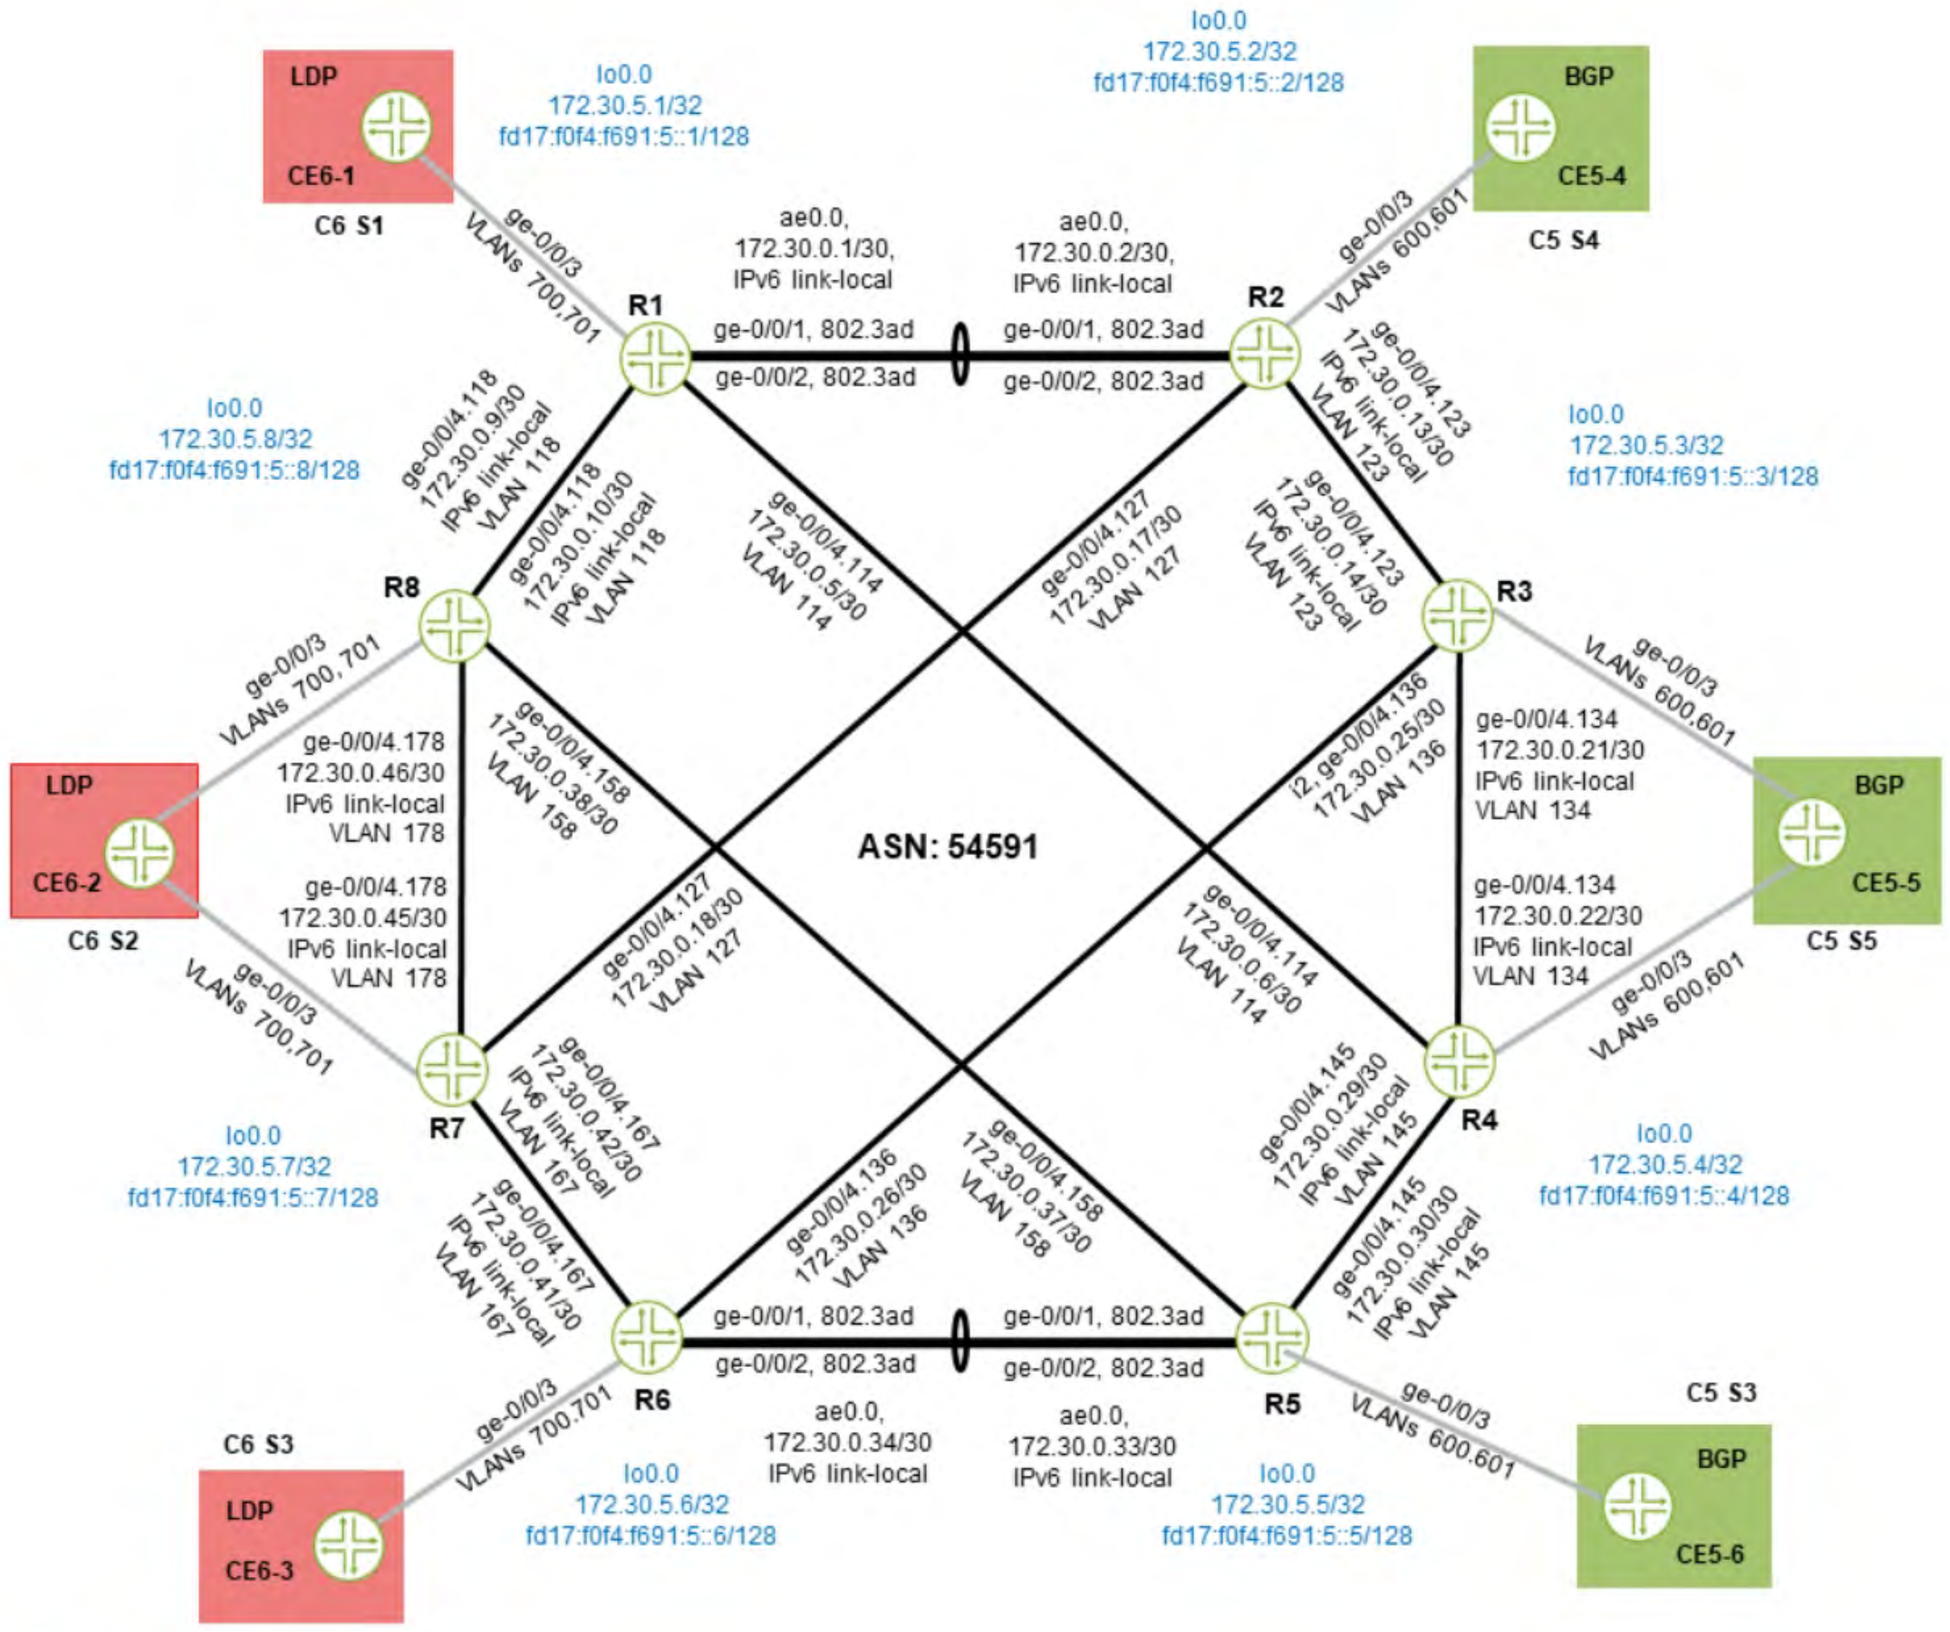


Figure 10

1) Configure VPLS as shown in Figure 13. Table 23 specifies the VPLS details. Configure customer VLANs as shown in Table 23.

Customer	Site	Router	VPLS signaling	CE facing interface
C5	S4	CE5-4	BGP	ge-0/0/3
	S5	CE5-5	BGP	ge-0/0/3
	S6	CE5-6	BGP	ge-0/0/3
C6	S1	CE6-1	LDP	ge-0/0/3
	S2	CE6-2	LDP	ge-0/0/3

© 2020 Juniper Networks, Inc. All rights reserved.

S3	CE6-3	LDP	ge-0/0/3

Table 12

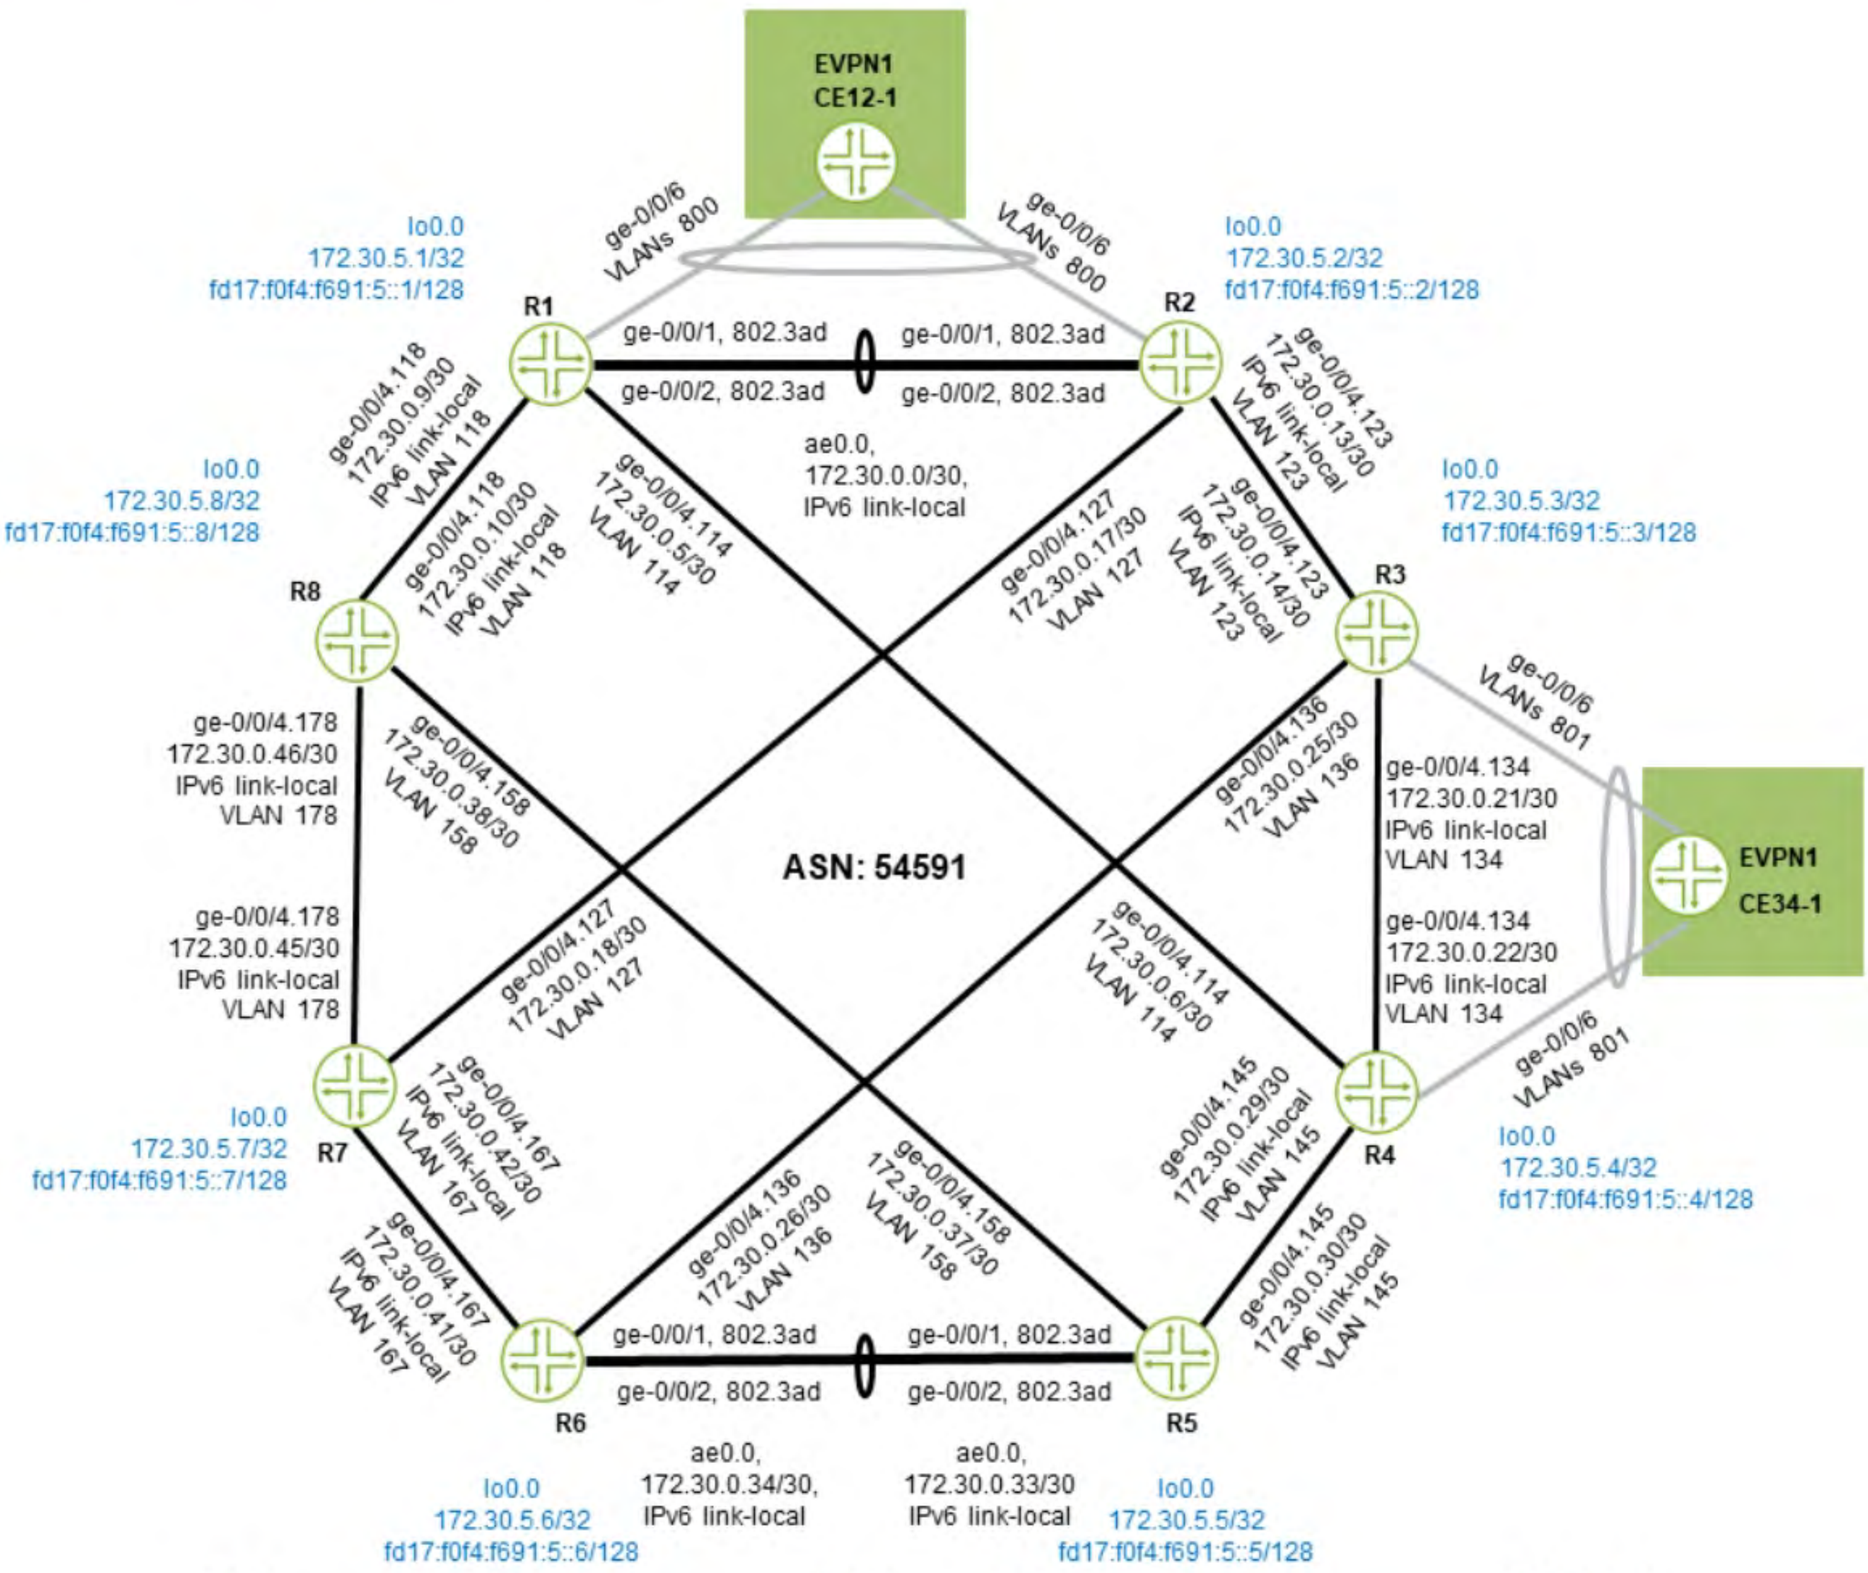
Customer	VLAN
C5	600
	601
C6	700
	701

Table 24

3) No L2 switching loops are allowed anywhere in the customers' VPLS networks. You cannot use Spanning Tree protocol for loop prevention.
4) Make sure that customer C6 dual-homed site S2 connection to R8 is the primary one. Configure the customer VPLS so that if the primary connection is active it is always preferred by other PE routers.
5) Customer C5 requires that you provide interworking between the customer's L2VPN and VPLS networks. Configure L2VPN and VPLS interworking at R2 such as CE5-1 is connected to VPLS VLAN 600.
6) Make sure that customer C5 MAC table size is limited to 200 entries per site, and customer C6 MAC table size is limited to 100 entries per site. Make sure that if customer C6 MAC table limit is reached, packets are dropped.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 3: EVPN Configuration



CE Router	CE L3 address	PE router	CE facing interface	LACP	VLAN ID	Default gateway
CE12-1	192.168.2.12/24	R1	ae1.800	Enabled	800	192.168.2.254/24
		R2				
CE34-1	192.168.2.34/24	R3	ae1.801		801	
		R4				

Table 13

Note: The CE devices are preconfigured on VR-DEVICE2.

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure an EVPN service to enable communication between CE12-1 and CE34-1. Enable active-active multihoming for unicast traffic. Use the information as provided in Table 25.

   Note: The CE routers are preconfigured on VR-DEVICE2.

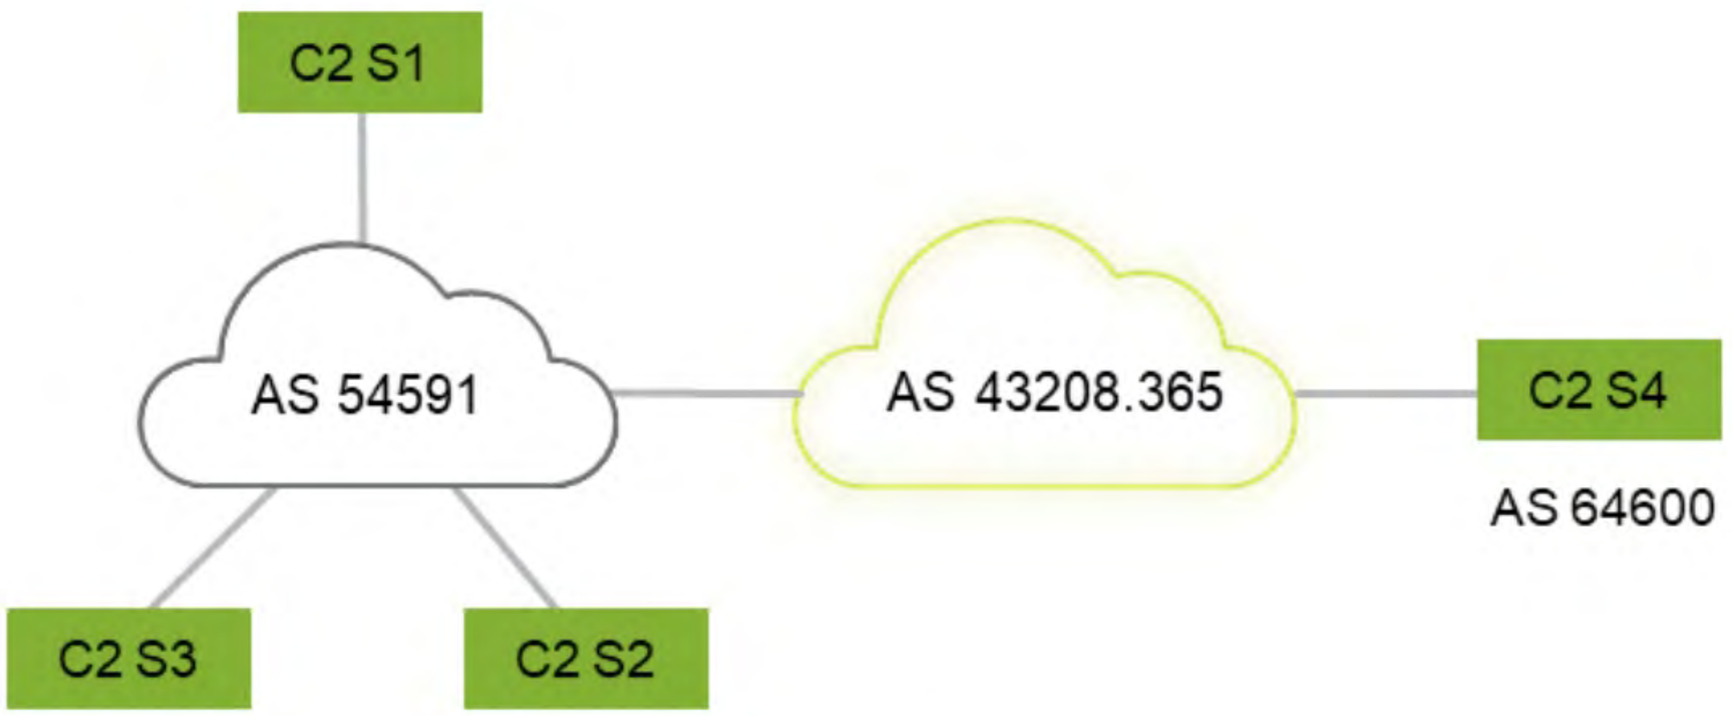
© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 7: Inter-Provider VPN Configuration

In this chapter, you practice with configuring inter-provider VPNs. The tasks include inter-provider VPN option B and option C.

Task 1: Inter-Provider VPN Option B

In this task, you configure inter-provider VPN option B.



Inter-Provider VPN Option B

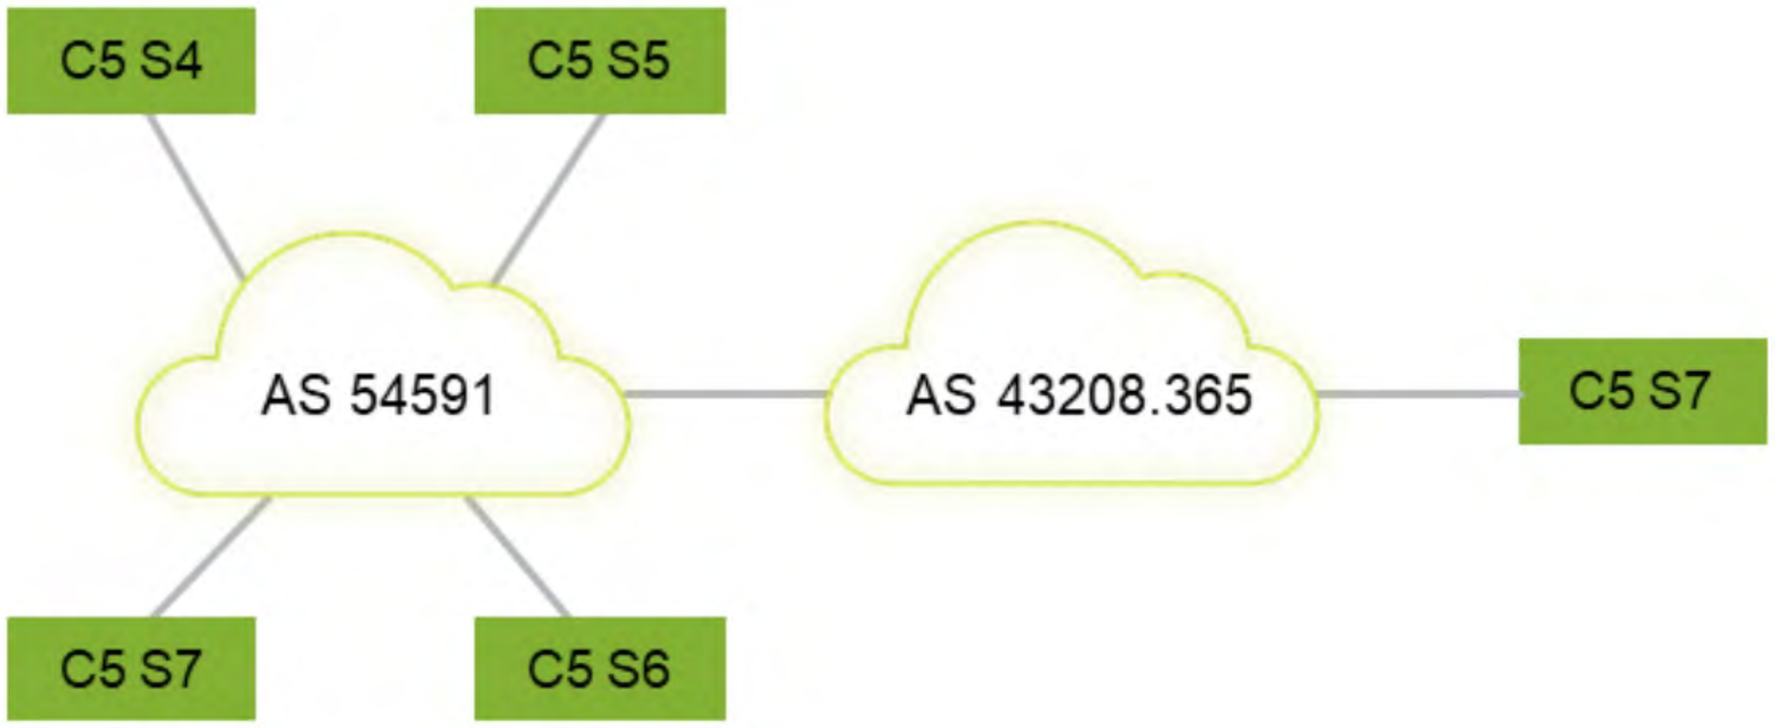
Figure 11

1) Customer C2 has a remote site S4 in the neighboring AS 43208.365 as shown in Figure 14. Configure your network to connect the remote site to the customer L3VPN using inter-provider VPN option B.
2) The remote site must be a spoke site in the customer hub-and-spoke VPN structure. Find out what VPN target is used by the remote site S4 PE router by using router monitoring tools and make sure that you advertise the customer VPN routes to the neighboring AS using the same community value.

Task 2: Inter-Provider VPN Option C

In this task, you configure inter-provider VPN option C.

© 2020 Juniper Networks, Inc. All rights reserved.



Inter-Provider VPN Option C

Figure 12

1) Customer C5 has a remote site S7 in the neighboring AS 43208.365 as shown in Figure 15. Configure your network to connect the remote site to the customer VPLS using inter-provider VPN option C.
2) The remote site S7 PE router IP address is 172.17.47.3. Find out what VPN target is used by the remote site S7 PE router by using router monitoring tools and make sure that you advertise the customer VPLS routes to the neighboring AS using the same community value.

© 2020 Juniper Networks, Inc. All rights reserved.

Chapter 8: Class of Service

This chapter is focused on Class of Service applications. You will configure MF and BA classifiers, policers, forwarding classes, queues and schedulers, rewrite markers, and RED drop profiles.

Task 1: Forwarding Classes, Queues, and Schedulers

In this task, you configure your network to support 4 DiffServ model Behavior Aggregates: VPN, VPN priority, best effort, and network control.

1) Configure Forwarding Classes and map them to the outgoing Queues as indicated in Table 26.

Forwarding Class	Queue	Scheduler
best-effort	0	be-sc-q0
vpn	1	vpn-sc-q1
vpn-priority	2	vpn-pri-sc-q2
nc	3	nc-sc-q3

Table 14

2) Configure Schedulers with parameters shown in Table 27 and map them to the Forwarding Classes as indicated in Table 26.

Scheduler	Parameter	Value
be-sc-q0	Priority	low
	Transmit rate	remainder
	Buffer size	remainder
	Drop profile LP any	high-drop
vpn-sc-q1	Priority	medium-low
	Transmit rate	20%
	Buffer size	20%
	Drop profile LP low	low-drop
	Drop profile LP high	high-drop
vpn-pri-sc-q2	Priority	medium-high
	Transmit rate	10%
	Buffer size	5 msec

© 2020 Juniper Networks, Inc. All rights reserved.

nc-sc-q3	Priority	high
	Transmit rate	5%
	Buffer size	5%

Table 15

3) Configure a Drop Profile called low-drop. Have a router to automatically build a smooth graph line based on the data points defined in Table 28.

Fill Level	Drop Probability
25	5
50	15
75	40

Table 28

4) Configure a Drop Profile called high-drop. Have a router to automatically build a smooth graph line based on the data points defined in Table 29.

Fill Level	Drop Probability
25	10
50	30
75	65

Table 16

5) Apply the schedulers to all your routers' core-facing interfaces. Make sure that the schedulers are applied at the interface logical unit level.

Task 2: Classification, Policing, and Marking

In this task, you configure packet classification, rate limiting and marking. You also map customer traffic to the respective DiffServ-enabled MPLS LSPs.

1) Configure the PE routers servicing customer C3 sites to classify packets received on the customer-facing interfaces using the MF classifier. The classification criteria are listed in Table 30.

© 2020 Juniper Networks, Inc. All rights reserved.

Traffic Type	Criteria	Forwarding Class
VPN regular	DSCP 0b000000	vpn
VPN priority	Any other DSCP value	vpn-priority

Table 17

2) Map the customer C3 VPN traffic to LSPs K and L, and VPN priority traffic to LSPs I and J.
3) Make sure that traffic entering LSPs I and J is limited to the LSP bandwidth value. The excess traffic must be dropped.
4) Make sure that traffic entering LSPs K and L is limited to the LSP bandwidth value. The excess traffic must have loss priority set to high.
5) Configure all routers to mark the packet CoS fields on the packets transmitted on the core-facing interfaces as shown in Table 31. Make sure that the CoS codes are configured as code point aliases.
6) Make sure that PE routers servicing customer C3 sites mark both IPv6 and MPLS packet headers' CoS fields.

Forwarding Class	Loss Priority	DSCP Value	EXP Value
best-effort	any	0b000000	0b000
vpn	low	0b001010	0b010
	high	0b001100	0b011
vpn-priority	any	0b101110	0b101
nc	any	0b110000	

Table 18

7) Configure all your routers to classify incoming traffic on all core-facing interfaces with BA classifiers using EXP bits value for MPLS packets and DSCP bits for IPv4 packets as specified in Table 31.

© 2020 Juniper Networks, Inc. All rights reserved.

Super Lab 1

This lab is designed to combine many of the solutions discussed in the previous chapters into a complete practice lab exam. This lab is similar to a real JNCIE-SP lab exam experience. We recommend attempting this lab if you are able to complete the tasks in the previous chapters without any help.

The figure below shows the network topology used for this chapter.

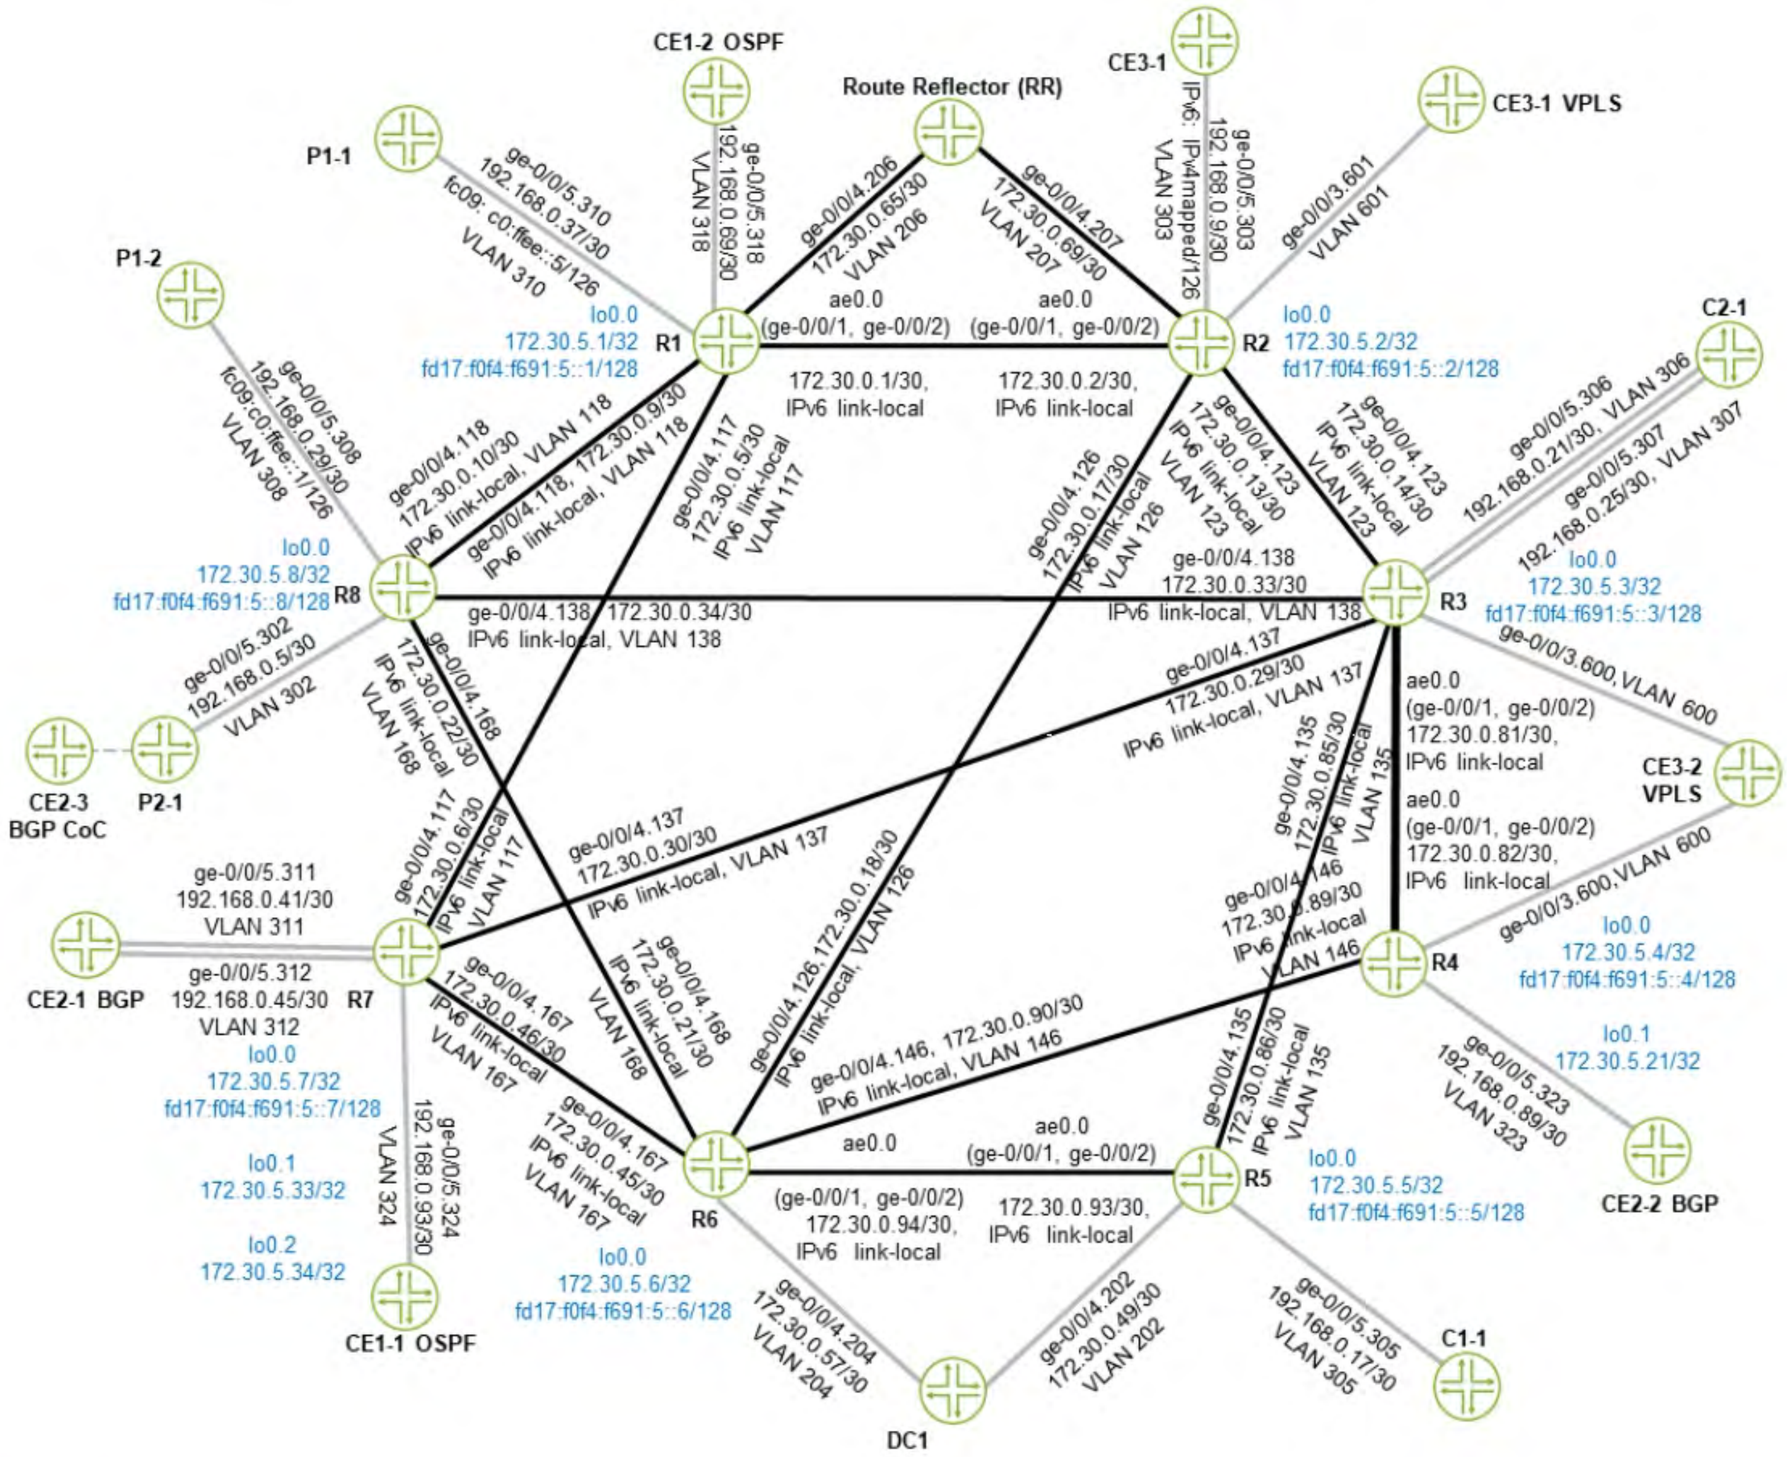


Figure 13

Note: You are not allowed to use static routes in any of the tasks in this chapter unless indicated explicitly otherwise.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 1: Initial System Configuration

The task objectives: Configure the initial system settings on all your devices. Ensure that your configuration meets the following criteria.

Use root password *root123* in every router. Please do not change the root password on our devices to prevent unnecessary password recovery.

1) Configure the host names in the routers according to Table 32.

Router	Router Type	Host Name
vmx1	MX	R1
vmx2	MX	R2
vmx3	MX	R3
vmx4	MX	R4
vmx5	MX	R5
vmx6	MX	R6
vmx7	MX	R7
vmx8	MX	R8

Table 19

2) Configure OoB management interfaces on each device with the appropriate IP addresses. The devices and their respective IP addresses are listed in Table 33. Set the interface description.

Device	OoB Interface Name	OoB Interface IP Address
R1	fxp0	10.10.1.1/24
R2	fxp0	10.10.1.2/24
R3	fxp0	10.10.1.3/24
R4	fxp0	10.10.1.4/24
R5	fxp0	10.10.1.5/24
R6	fxp0	10.10.1.6/24
R7	fxp0	10.10.1.7/24
R8	fxp0	10.10.1.8/24

Table 20

3) Enable each device to accept management connections for the SSH, Telnet, and FTP services only.
4) Configure static route to remote management network 10.10.10/24 with the next-hop 10.10.1.254 on all your devices. Make sure the network is never

© 2020 Juniper Networks, Inc. All rights reserved.

redistributed to any dynamic routing protocol. Ensure that the devices are reachable while RPD is not running.

5) Configure the S1 server as the DNS server.

Note: Server S1 is a virtual NTP/FTP/SNMP/Syslog/RADIUS/DNS proxy server. The server is reachable at 10.10.1.100 IP address.

6) Set the time zone to Europe/Amsterdam on all your devices.
7) Ensure that all your devices synchronize their time with the NTP server S1. Configure the devices to synchronize time with the S1 at boot time.
8) Configure the authentication method that first tries to authenticate users on RADIUS server and then if not successful with local password. Use S1 as the RADIUS server. Configure the RADIUS server with retry attempts 1 and timeout 2 seconds. Use workbook as the RADIUS shared secret.
9) Create on every device a new user lab, with the password lab123, that will have super user privileges. From this point on configure your devices using user lab account.
10) Configure additional users on all the devices as defined in Table 34.

Username	Password	Privileges
noc	noc123	Class “operator” permissions. Additionally, this user can read and modify SNMP configuration, execute system maintenance commands but not allowed to execute “start shell” command
tac	tac123	Class “super-user” permissions. Additionally, this user cannot execute the “clear”, “configure” or “edit” commands

Table 21

11) Configure Syslog settings on all your devices as indicated in Table 35.

Receiver	Message Type
File “jncie-sp-messages”	All info level messages
File “firewall.log”	All firewall filter messages
Syslog server S1	Configuration changes
User “noc”	All warning level messages

Table 22

© 2020 Juniper Networks, Inc. All rights reserved.

12) Configure SNMP v2 for read-only access using a community workbook. Make sure that SNMP server S1 is the only server allowed to access the device with this community.
13) Configure SNMP v2 to send traps to the SNMP server S1 for routing, link, and chassis events.
14) Configure an IPv4 firewall filter allowing any protocol packets sourced from 10.10.1/24 10.10.10/24 management networks, and 172.17/16, 172.30/16, 172.31/16 and 192.168/16 operative networks. Configure the firewall filter to discard all other packets, increment a named counter and send notifications to syslog.
15) Apply the firewall filter to protect the Routing Engine.
16) Set all your devices to archive configuration periodically every 24 hours to the FTP server S1 using username lab and password lab123.
17) Download op script called “show-interfaces.slax”, commit script called “interface-mask-check.slax” and event script called “ospf_adjacency_flapping.slax” from the FTP server S1 to all your routers.

Note: These are example scripts written by Juniper Networks and available in public domain.

a) Enable the scripts.

Task 2: Building the Network

The task objectives: Configure network interfaces on all your devices. Provide basic network connectivity. Ensure that your configuration meets the following criteria.

1) Build the network by configuring interfaces as indicated in Table 37. Aggregated Ethernet interfaces are listed in Table 36.
2) Enable LACP continuity checking on the Aggregated Ethernet interfaces.
3) Set all the interfaces descriptions.

Router	Aggregated Ethernet	Interfaces
R1	ae0	ge-0/0/1
		ge-0/0/2
R2	ae0	ge-0/0/1
		ge-0/0/2
R3	ae0	ge-0/0/1
		ge-0/0/2
R4	ae0	ge-0/0/1

© 2020 Juniper Networks, Inc. All rights reserved.

		ge-0/0/2
R5	ae0	ge-0/0/1
		ge-0/0/2
R6	ae0	ge-0/0/1
		ge-0/0/2

Table 23

Note: The interface unit numbers match the VLAN tags.

Router	Interface Name	IP Address	IPv6 Address
R1	ae0.0	172.30.0.1/30	link-local
	ge-0/0/4.117	172.30.0.5/30	link-local
	ge-0/0/4.118	172.30.0.9/30	link-local
	ge-0/0/4.206	172.30.0.65/30	
	ge-0/0/5.318	192.168.0.69/30	
	ge-0/0/5.310	192.168.0.37/30	fc09:c0:ffee::5/126
	lo0.0	172.30.5.1/32	fd17:f0f4:f691:5::1/128
R2	ae0.0	172.30.0.2/30	link-local
	ge-0/0/4.126	172.30.0.17/30	link-local
	ge-0/0/4.123	172.30.0.13/30	link-local
	ge-0/0/4.207	172.30.0.69/30	
	ge-0/0/5.303	192.168.0.9/30	IPv4 compatible/126
	ge-0/0/3.601		
	lo0.0	172.30.5.2/32	fd17:f0f4:f691:5::2/128
R3	ge-0/0/4.123	172.30.0.14/30	link-local
	ge-0/0/4.138	172.30.0.33/30	link-local
	ge-0/0/4.137	172.30.0.29/30	link-local
	ge-0/0/4.135	172.30.0.85/30	link-local
	ae0.0	172.30.0.81/30	link-local
	ge-0/0/5.306	192.168.0.21/30	
	ge-0/0/5.307	192.168.0.25/30	
	ge-0/0/3.600		
	lo0.0	172.30.5.3/32	fd17:f0f4:f691:5::3/128
R4	ge-0/0/4.146	172.30.0.89/30	link-local
	ae0.0	172.30.0.82/30	link-local
	ge-0/0/3.600		
	ge-0/0/5.323	192.168.0.89/30	
	lo0.0	172.30.5.4/32	fd17:f0f4:f691:5::4/128

© 2020 Juniper Networks, Inc. All rights reserved.

	lo0.1	172.30.5.21/32	
R5	ge-0/0/4.135	172.30.0.86/30	link-local
	ae0.0	172.30.0.93/30	link-local
	ge-0/0/5.305	192.168.0.17/30	
	ge-0/0/4.202	172.30.0.49/30	
	lo0.0	172.30.5.5/32	fd17:f0f4:f691:5::5/128
R6	ge-0/0/4.126	172.30.0.18/30	link-local
	ge-0/0/4.146	172.30.0.90/30	link-local
	ae0.0	172.30.0.94/30	link-local
	ge-0/0/4.167	172.30.0.45/30	link-local
	ge-0/0/4.168	172.30.0.21/30	link-local
	ge-0/0/4.204	172.30.0.57/30	
	lo0.0	172.30.5.6/32	fd17:f0f4:f691:5::6/128
R7	ge-0/0/4.117	172.30.0.6/30	link-local
	ge-0/0/4.137	172.30.0.30/30	link-local
	ge-0/0/4.167	172.30.0.46/30	link-local
	ge-0/0/5.311	192.168.0.41/30	
	ge-0/0/5.312	192.168.0.45/30	
	ge-0/0/5.324	192.168.0.93/30	
	lo0.0	172.30.5.7/32	fd17:f0f4:f691:5::7/128
	lo0.1	172.30.5.33/32	
	lo0.2	172.30.5.34/32	
R8	ge-0/0/4.118	172.30.0.10/30	link-local
	ge-0/0/4.138	172.30.0.34/30	link-local
	ge-0/0/4.168	172.30.0.22/30	link-local
	ge-0/0/5.308	192.168.0.29/30	fc09:c0:ffee::1/126
	ge-0/0/5.302	192.168.0.5/30	
	lo0.0	172.30.5.8/32	fd17:f0f4:f691:5::8/128

Table 24

Task 3: IGP Configuration

The task objectives: Enable OSPFv3 routing in your AS. Enable RIP – OSPFv3 redistribution. Provide intra-domain connectivity. Ensure that your configuration meets the following criteria.

1) Configure OSPFv2 and OSPFv3 in your network according to the Table 38 specifications. Make sure that OSPF is not running on the OoB management interface and on the AS external interfaces.

© 2020 Juniper Networks, Inc. All rights reserved.

Note: Both OSPFv2 and OSPFv3 are referred to as OSPF in the subsequent tasks.

Router	Interface		Area
R1	ae0.0		0.0.0.0
	ge-0/0/4.117		0.0.0.0
	ge-0/0/4.118		0.0.0.0
	lo0.0		0.0.0.0
R2	ae0.0		0.0.0.0
	ge-0/0/4.126		0.0.0.0
	ge-0/0/4.123		0.0.0.0
	lo0.0		0.0.0.0
R3	ge-0/0/4.123		0.0.0.0
	ge-0/0/4.138		0.0.0.0
	ge-0/0/4.137		0.0.0.0
	ge-0/0/4.135		0.0.0.1
	ae0.0		0.0.0.1
	lo0.0		0.0.0.0
R4	ge-0/0/4.146		0.0.0.1
	ae0.0		0.0.0.1
	lo0.0		0.0.0.1
R5	ge-0/0/4.135		0.0.0.1
	ae0.0		0.0.0.1
	lo0.0		0.0.0.1
R6	ge-0/0/4.126		0.0.0.0
	ge-0/0/4.146		0.0.0.1
	ae0.0		0.0.0.1
	ge-0/0/4.167		0.0.0.0
	ge-0/0/4.168		0.0.0.0
	lo0.0		0.0.0.0
R7	ge-0/0/4.117		0.0.0.0
	ge-0/0/4.137		0.0.0.0
	ge-0/0/4.167		0.0.0.0
	lo0.0		0.0.0.0
R8	ge-0/0/4.118		0.0.0.0
	ge-0/0/4.138		0.0.0.0
	ge-0/0/4.168		0.0.0.0
	lo0.0		0.0.0.0

Table 25

© 2020 Juniper Networks, Inc. All rights reserved.

2) Configure OSPFv2 only on R1 and R2 as shown in Table 39. Enable OSPFv2 on route reflector.

Router	Interface	Area
R1	ge-0/0/4.206	0.0.0.0
R2	ge-0/0/4.207	0.0.0.0

Table 26

3) Make sure that router ID is configured explicitly on all routers.
4) Make sure that you do not have Type 2 LSAs in your domain.
5) Make sure that Area 1 LSDB does not have any of the OSPF Type 4 or Type 5 LSAs.
6) Make sure that routers in Area 1 will not be isolated in case of a single link or ABR failure.
7) Configure Area 1 OSPF internal IPv4 routes tightest possible summarization to the backbone area.
8) Configure all routers to automatically calculate metrics reflecting interfaces' bandwidth.
9) Make sure that all OSPF adjacencies are in Full state and connectivity is provided among all routers' loopback interfaces for both IPv4 and IPv6 families.
10) Make sure that connectivity is provided between all routers' loopback interfaces and route reflector loopback interface address 172.30.5.41. Any of the R1 or R2 failure must not result in loss of route reflector loopback reachability.
11) Enable RIP on R5 (ge-0/0/4.202) and R6 (ge-0/0/4.204) interfaces.
12) Redistribute the default route into RIP. Make sure that the R6 default route advertisement is preferred by DC1.
13) Redistribute RIP routes into OSPF.
14) Any OSPF ASBR failure must not result in RIP routes disappearing from OSPF or the default route disappearing from RIP.
15) Configure Area 1 OSPF external IPv4 routes tightest possible summarization to the backbone area. Make sure that the more specific external routes do not appear in the backbone area.
16) Any OSPF ABR failure must not result in RIP summary route disappearing from OSPF backbone area.
17) Make sure that R5 and R6 use optimal routing to reach OSPF destinations outside Area 1.
18) No routing loops are allowed anywhere.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 4: BGP Configuration

The task objectives: Configure BGP network including IBGP sessions with route reflector and EBGP sessions with multiple peers and customers. Configure routing policies to handle IPv4 and IPv6 routing exchanges. Ensure that your configuration meets the following criteria.

1) Configure IBGP with route reflection. There must be two clusters and any client might be a member of one cluster only. Your AS number is 54591.
2) Clients can only have IBGP sessions with the route reflector.
3) You cannot use native IPv6 IBGP sessions anywhere.
4) Make sure that IBGP sessions use loopback interface peering.
5) Configure MD5 authentication for all IBGP sessions.
6) Ensure that all IBGP sessions state changes are logged to syslog.
7) Configure EBGP sessions as shown in Table 40.

Device	Peer Router	Peer AS	Peer IPv4 Address	Peer IPv6 Address
R1	P1-1	1679.12483	192.168.0.38	fc09:c0:ffee::6
R2	C3-1	64514	192.168.0.10	IPv4 compatible
R3	C2-1	64513	172.31.31.1	
R5	C1-1	64512	192.168.0.18	
R8	P1-2	1679.12483	192.168.0.30	fc09:c0:ffee::2
	P2-1	43208.365	192.168.0.6	

Table 27

8) Make sure that no more than 20 prefixes are accepted from any customer. If this limit is exceeded the session should be torn down and remain down for 5 minutes.
9) You might not establish native IPv6 EBGP session with customer C3 but you must enable IPv6 routing support.
10) Use loopback interface peering for R3 to C2-1 session. Make sure that a single interface failure will not break the EBGP session down. You can use static routing at this step.
11) All routes received from any customer should be damped in case of flapping. C1 routes must be damped more aggressively.
12) Make sure that the private AS numbers do not appear in the AS Path of any routes advertised to any EBGP peer.

© 2020 Juniper Networks, Inc. All rights reserved.

13) Configure the EBGP sessions with P1 and P2 peers to send keepalive messages once in 10 seconds.
14) Ensure that all EBGP sessions state changes are logged to syslog.
15) Make sure that any customer IPv4 routes are advertised to all EBGP peers.
16) Make sure that routes received from P1 neighbors are not advertised to P2 neighbors and vice versa.
17) Do not accept any IPv4 prefixes with AS Path length longer than 5 hops from P2 peers.
18) Do not advertise any external BGP routes to customer C1. Advertise the default route instead.
19) If a route is learned directly from a customer, it should always be preferred to the same route learned from any other peer.
20) Do not accept IPv4 routes that have a mask shorter than /8 or longer than /24 from anywhere. You may accept routes with mask /32 originated in AS 43208.365.
21) Do not accept the 0.0.0.0 route with any mask length from any of the peers or customers.
22) Do not accept any IPv6 routes that are not originated in their AS from P1 neighbors.
23) Use two standard communities to identify IPv4 routes received from either a customer or a peer. None of these communities may be seen outside of your AS.
24) Advertise a single summary IPv4 route that aggregates your AS local routes including the RIP routes to all your EBGP peers.
25) Make sure that IPv6 routes advertised to P1 neighbors are not advertised further outside of their AS.
26) Make sure that R1 is the preferred point both for inbound and outbound IPv4 traffic for P1 AS.
27) Make sure that if a customer advertises an IPv4 route with a community of "<Customer AS>:666" the traffic to that destination is black-holed.
28) No unresolved IPv4 or IPv6 routes are allowed anywhere.

Task 5: MPLS Configuration

The task objectives: Configure backbone MPLS network including configuration of LDP- and RSVP-signaled LSPs, traffic engineering, traffic protection and optimization, and LDP tunneling. Ensure that your configuration meets the following criteria.

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure LDP interfaces as shown in Table 41. Enable LDP on route reflector.

Router	Interface
R1	ge-0/0/4.206
R2	ge-0/0/4.207
R3	ge-0/0/4.135
	ae0.0
R4	ge-0/0/4.146
	ae0.0
R5	ge-0/0/4.135
	ae0.0
R6	ge-0/0/4.146
	ae0.0

Table 28

2) Configure MD5 authentication for all LDP sessions.
3) Configure OSPF to track the LDP operational status on all LDP-enabled interfaces.
4) Make sure that LDP LSPs show the same metrics as the IGP paths they follow.
5) Configure RSVP interfaces as shown in Table 42. Enable RSVP message aggregation.
6) Configure link administrative groups as shown in Table 42.

Router	Interface	Color
R1	ae0.0	green
	ge-0/0/4.117	blue
	ge-0/0/4.118	purple
R2	ae0.0	green
	ge-0/0/4.126	blue
	ge-0/0/4.123	purple
R3	ge-0/0/4.123	purple
	ge-0/0/4.138	blue
	ge-0/0/4.137	green
R6	ge-0/0/4.126	blue
	ge-0/0/4.167	purple
	ge-0/0/4.168	green
R7	ge-0/0/4.117	blue
	ge-0/0/4.137	green

© 2020 Juniper Networks, Inc. All rights reserved.

R8	ge-0/0/4.167	purple
	ge-0/0/4.118	purple
	ge-0/0/4.138	blue
	ge-0/0/4.168	green

Table 29

7) Configure all RSVP-enabled interfaces except the Aggregated Ethernet bundles to allow bandwidth reservation with 20% oversubscription.
8) Configure full mesh of RSVP sessions among all routers except R4, R5 and route reflector.
9) Configure MD5 authentication for all RSVP sessions.
10) Enable RSVP path MTU discovery for all RSVP sessions.
11) Make sure that LSPs originated at R1, R2, R3 use only links belonging to “green” or “blue” administrative groups.
12) Make sure that LSPs originated at R6, R7, R8 use only links belonging to “purple” or “blue” administrative groups.
13) Configure an additional LSP from R2 to R1 and an LSP from R2 to R8. The additional LSPs may not use administrative group constraint.
14) Make sure that the two LSPs from R2 to R1 and the two LSPs from R2 to R8 do not use the same physical link anywhere on the path to the egress nodes.
15) Configure all LSPs except those from R2 to R1 and from R2 to R8 to reserve 100Mbps of bandwidth.
16) Configure the LSPs from R2 to R1 and to R8 to automatically adjust bandwidth once in 24 hours based on the average bandwidth usage. Make sure that the LSPs are signaled with not less than 50Mbps and not more than 100Mbps.
17) Configure LSPs originated at R3 and R6 to ensure that they have higher priority for bandwidth reservation than the remaining LSPs, including the P2MP LSPs. Make sure that the remaining P2P LSPs have lower priority than that of P2MP LSPs.
18) Configure LDP tunnels to establish MPLS LSPs between R4, R5 and route reflector. Make sure that a single link or node failure will not result in these LSPs break down.
19) Make sure that IPv4 and IPv6 traffic from C3 to P1 are mapped to different LSPs.
20) Configure a backup protection path for all RSVP-signaled LSPs. Make sure that for the LSPs originated at R3 and R6 the protection path is established immediately.

© 2020 Juniper Networks, Inc. All rights reserved.

21) Make sure that bandwidth is shared between the main path and protection path for the LSPs originated at R3 and R6.
22) Configure the LSPs originated at R3 and R6 to use fast reroute protection. Make sure that bandwidth is inherited by the detour path, but administrative groups are not.
23) Configure the remaining LSPs to use link protection.
24) Enable IPv6 over MPLS tunneling in your AS.

Task 6: VPN Configuration

The task objectives: Implement L3VPN infrastructure including customers running either OSPF or BGP, hub-and-spoke topologies, customer internet access, multicasting in VPNs and inter-provider VPNs. Implement VPLS and EVPN services including dual-homed customer sites and VLAN normalization. Ensure that your configuration meets the following criteria.

1) Configure L3VPN as shown in Table 43.

Customer	Site	Router	PE-CE Protocol	Protocol details
CE1	S1	CE1-1	OSPF	Area 0.0.0.0
	S2	CE1-2	OSPF	Area 0.0.0.0
CE2	S1	CE2-1	BGP	AS 64600
	S2	CE2-2	BGP	AS 64600
	S3	CE2-3	BGP	AS 64600

Table 30

2) Make sure that all PE routers receive only the routes with those targets that they specifically request for.
3) Customer CE1 has a backdoor OSPF connection and wants to use your MPLS network as a backup path between the customer sites. Make sure that in the customer VPN all remote site OSPF routes always appear as external routes.
4) Customer CE2 requires that the customer site S1 is used as a central transit site for all traffic exchanges among all the customer sites in a hub-and-spoke fashion.
5) Make sure that PE-CE link subnets in customer CE2 VPN are advertised to the customer remote VPN sites.
6) Allow route exchange between customer CE1 site S1 and customer CE2 site S1 at R7. Make sure that the routes exchanged between the local VRFs are not

© 2020 Juniper Networks, Inc. All rights reserved.

advertised to any of the remote customer sites. You may not use RIB groups in this step.

7) Customer CE2 must be provided with Internet access at the customer site S1 using single customer-facing VRF interface. Other customer CE2 sites in your AS should be able to reach the Internet via the central site. Static route is permissible in this step.
8) Configure NG MVPN in customer CE2 VPN in your AS. Customer sites S1 and S2 can both act either as a sender site or a receiver site. Make sure that P2MP LDP-signaled LSP is used as the PMSI.
9) Customer CE2 outsources its RP to your network. Make sure that your PE routers act as the customer RPs. Use 172.30.5.253 as the RP address.
10) Make sure that the customer CE2 sites join only source-based multicast distribution trees.
11) Customer CE2 has a remote site S3 in the neighboring AS 43208.365. Configure your network to connect the remote site to the customer VPN using inter-provider VPN option C.
12) The remote customer CE2 site PE router IP address is 172.17.47.2 Find out what VPN target is used by the customer CE2 remote site PE router by using router monitoring tools.
13) Configure customer CE3 VPLS as shown in Table 44. The customer uses VLANs 600 and 601.

Customer	Site	Router	VPLS signaling	CE facing interface
CE3	S1	CE3-1	BGP	ge-0/0/3.601
	S2	CE3-2	BGP	ge-0/0/3.600

Table 31

14) No L2 switching loops are allowed anywhere in the customer VPLS network. You may not use Spanning Tree protocol for loop prevention.
15) Configure customer CE3 VLAN normalization.
16) Make sure that customer CE2 MAC table size is limited to 100 entries per interface on all PE routers. Make sure that if the limit is reached, packets are dropped.
17) Configure an EVPN service to enable communication between CE12-1 and CE34-1. Enable active-active multihoming for unicast traffic. Use the information as provided in table above.

© 2020 Juniper Networks, Inc. All rights reserved.

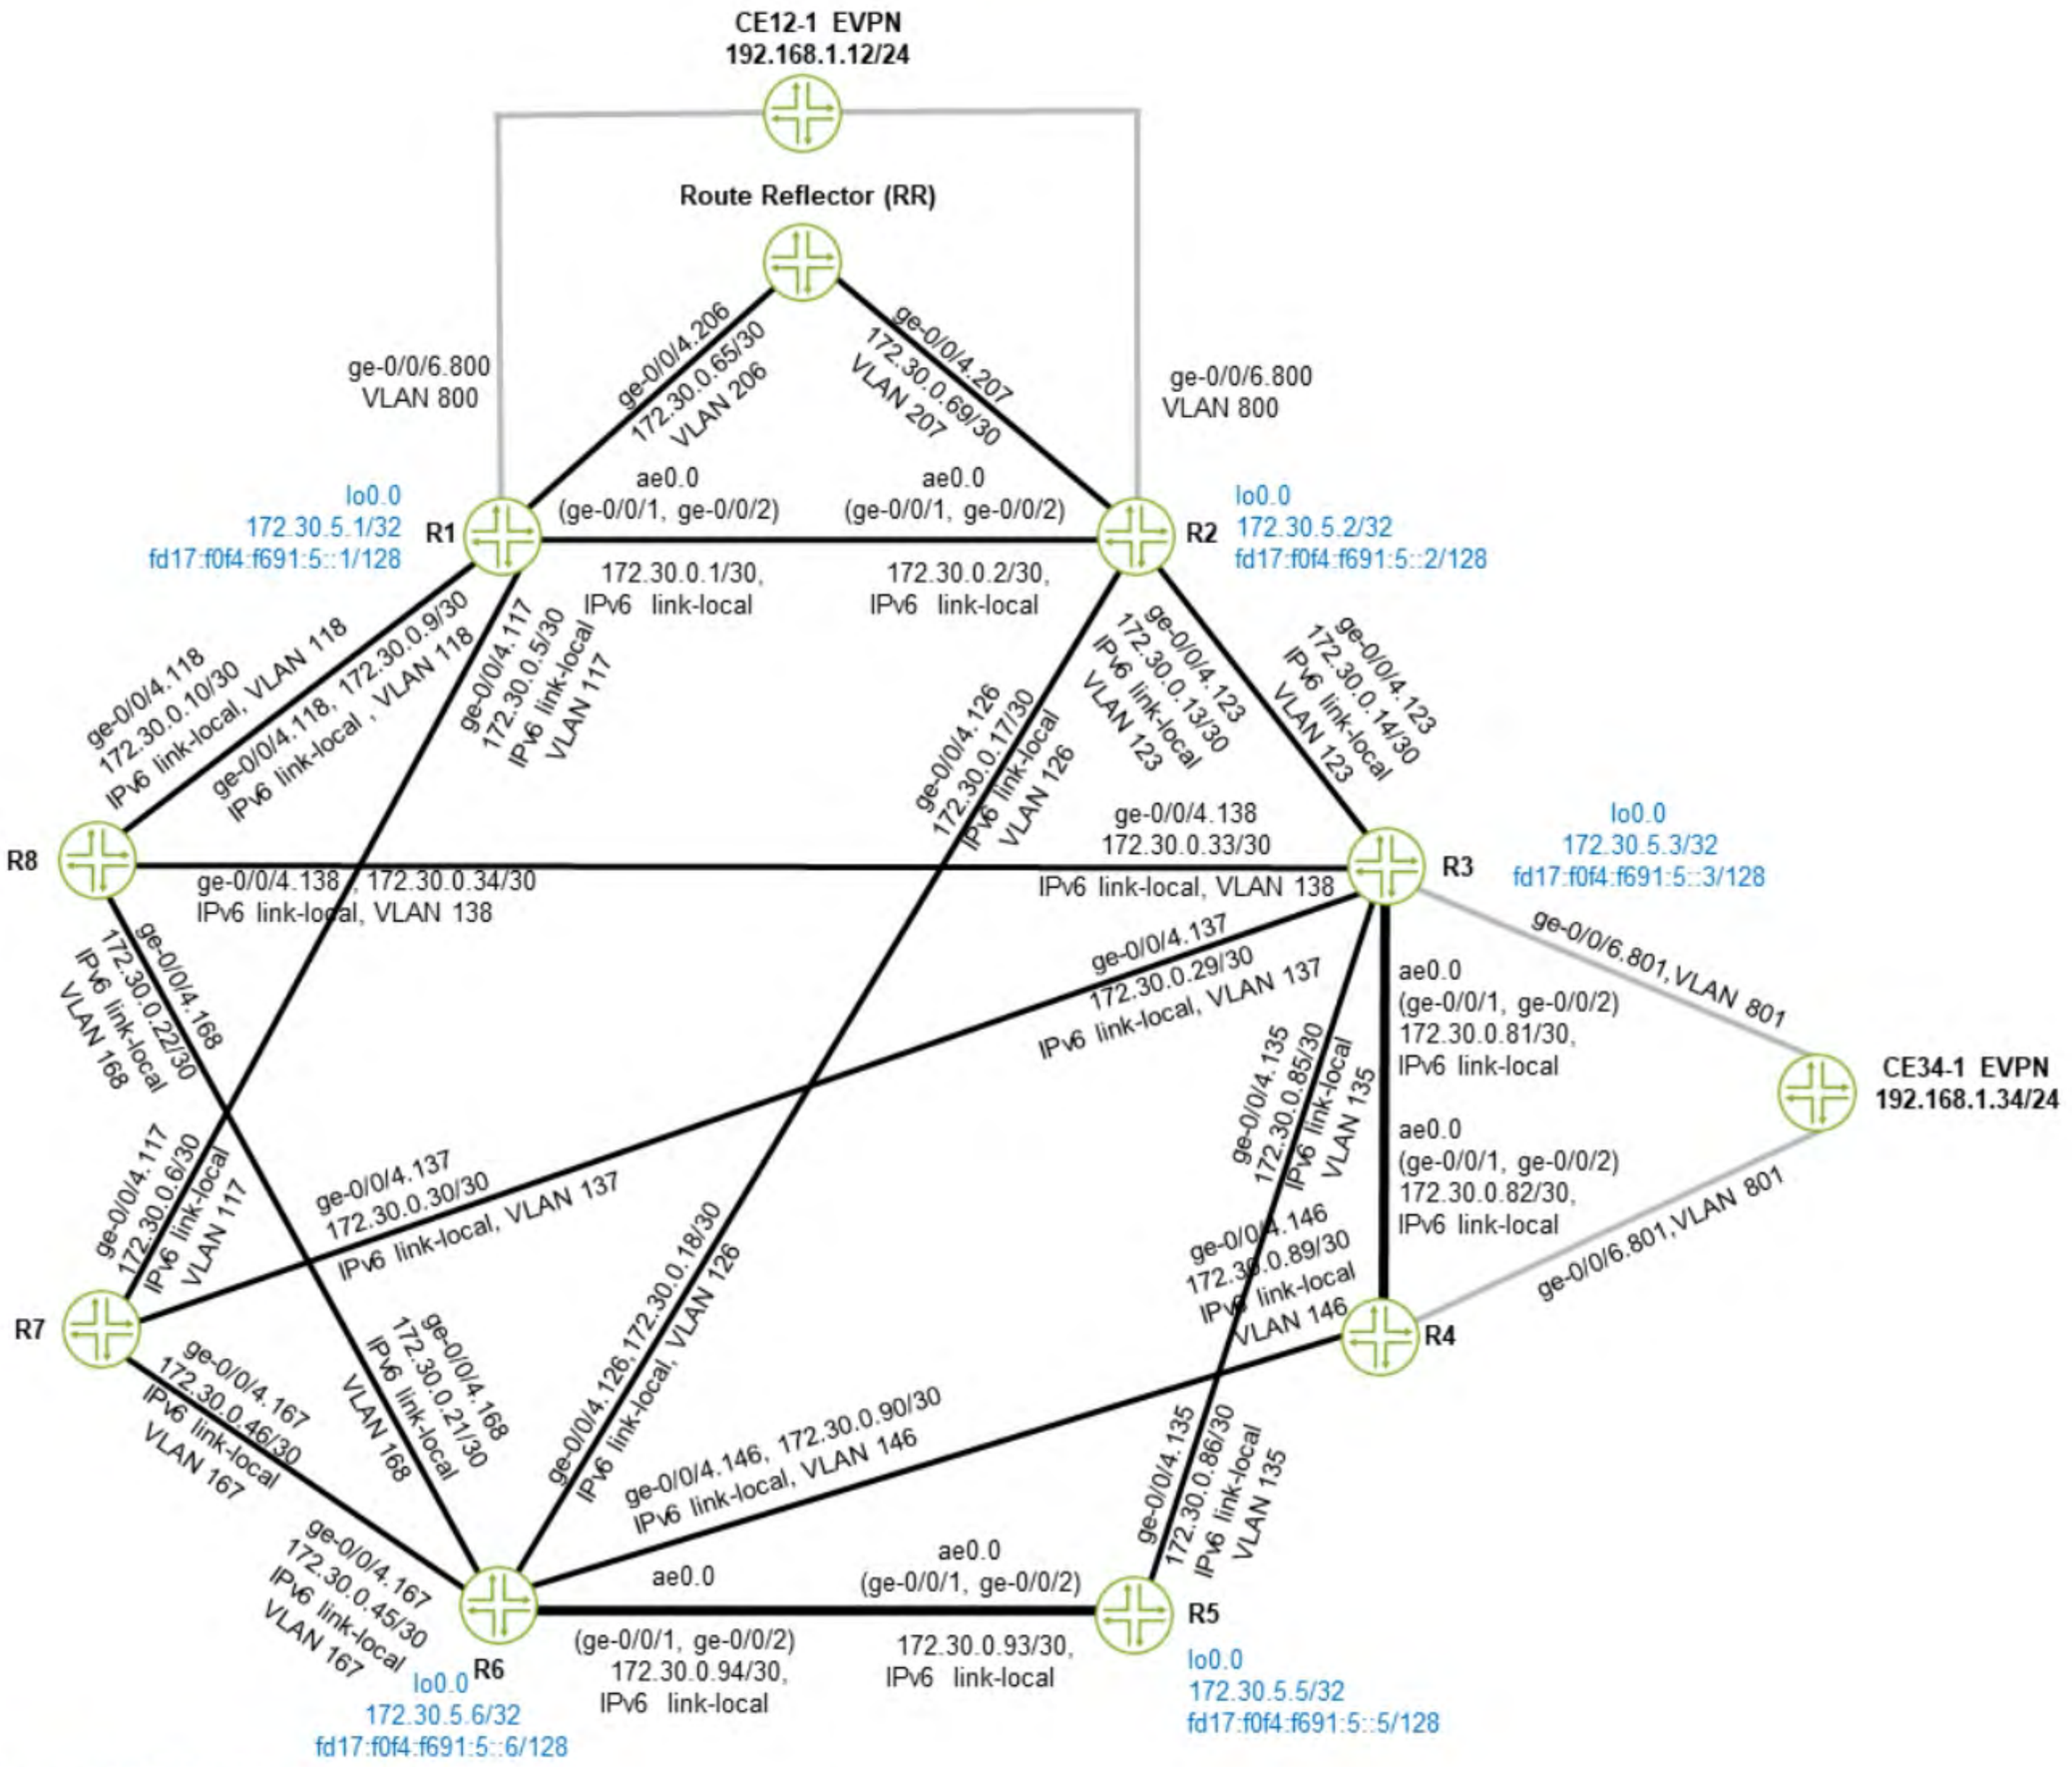


Figure 14 EVPN topology

CE Router	CE L3 address	PE router	CE facing interface	LACP	VLAN ID	Default gateway
CE12-1	192.168.1.12/24	R1	ae1.800	Enabled	800	192.168.1.254/24
		R2				
CE34-1	192.168.1.34/24	R3	ae1.801		801	
		R4				

Figure 15 EVPN configuration requirements

Note: The CE devices are preconfigured on VR-DEVICE2.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 7: Class of Service Configuration

The task objectives: Configure CoS aware network including classifiers, policers, forwarding classes, schedulers and rewrite markers.

1) Configure the PE routers servicing L3VPN customers to classify packets received on the customer-facing interfaces using the Multi-Field classifier as specified in Table 45.

Traffic Type	Criteria	Forwarding Class
VPN regular	DSCP 0b000000	l3vpn
VPN priority	DSCP 0b101110	l3vpn-priority

Table 32

2) Configure the PE routers servicing VPLS customers to classify packets received on the customer-facing interfaces using the Multi-Field classifier so that all received packets are assigned to “l2vpn” forwarding class.
3) Make sure that traffic entering PE routers from L3VPN customers and classified as l3vpn-priority does not exceed 25Mbps with allowed bursts up to 15KB, the excess traffic must be dropped.
4) Make sure that traffic entering PE routers from VPLS customers does not exceed 50Mbps with allowed bursts up to 62KB, the excess traffic must have drop priority increased.
5) Configure forwarding classes and map them to the outgoing queues as shown in Table 46.

Forwarding Class	Queue	Scheduler
be	0	be-sc
l3vpn	1	l3vpn-sc
l2vpn	2	l3vpn-pri-sc
l3vpn-priority	3	l2vpn-sc
nc	4	nc-sc

Table 33

© 2020 Juniper Networks, Inc. All rights reserved.

6) Configure schedulers with parameters shown in Table 47.

Scheduler	Parameter	Value
be-sc	Priority	low
	Transmit rate	remainder
	Buffer size	remainder
	Drop profile LP any	high-drop
l3vpn-sc	Priority	medium-low
	Transmit rate	20%
	Buffer size	20%
l2vpn-sc	Priority	medium-high
	Transmit rate	20%
	Buffer size	20%
	Drop profile LP low	low-drop
	Drop profile LP high	high-drop
l3vpn-pri-sc	Priority	high
	Transmit rate	10%
	Buffer size	5 msec
nc-sc	Priority	high
	Transmit rate	5%
	Buffer size	5%

Table 34

7) Configure drop profiles as shown in Table 48. Have a router to automatically build a smooth graph line based on the defined data points.

Drop Profile	Fill Level	Drop Probability
low-drop	25	5
	50	15
	75	40
high-drop	25	10
	50	30
	75	65

Table 35

8) Apply the schedulers to all your routers' core-facing interfaces.
9) Configure all routers to mark the packets' CoS fields on the packets transmitted on the core-facing interfaces as shown in Table 49. Make sure that all PE and BGP ASBR routers mark both IPv4 and MPLS packet headers' CoS fields.

© 2020 Juniper Networks, Inc. All rights reserved.

10) Configure all your routers to classify incoming traffic on all core-facing interfaces with Behavior Aggregate classifiers using EXP bits value for MPLS packets and DSCP bits for IPv4 packets.

Forwarding Class	Loss Priority	DSCP Value	EXP Value
be	low	0b000000	0b000
l3vpn	low	0b001000	0b001
l2vpn	low	0b001010	0b010
	high	0b001011	0b011
l3vpn-priority	low	0b101110	0b101
nc	low	0b110000	

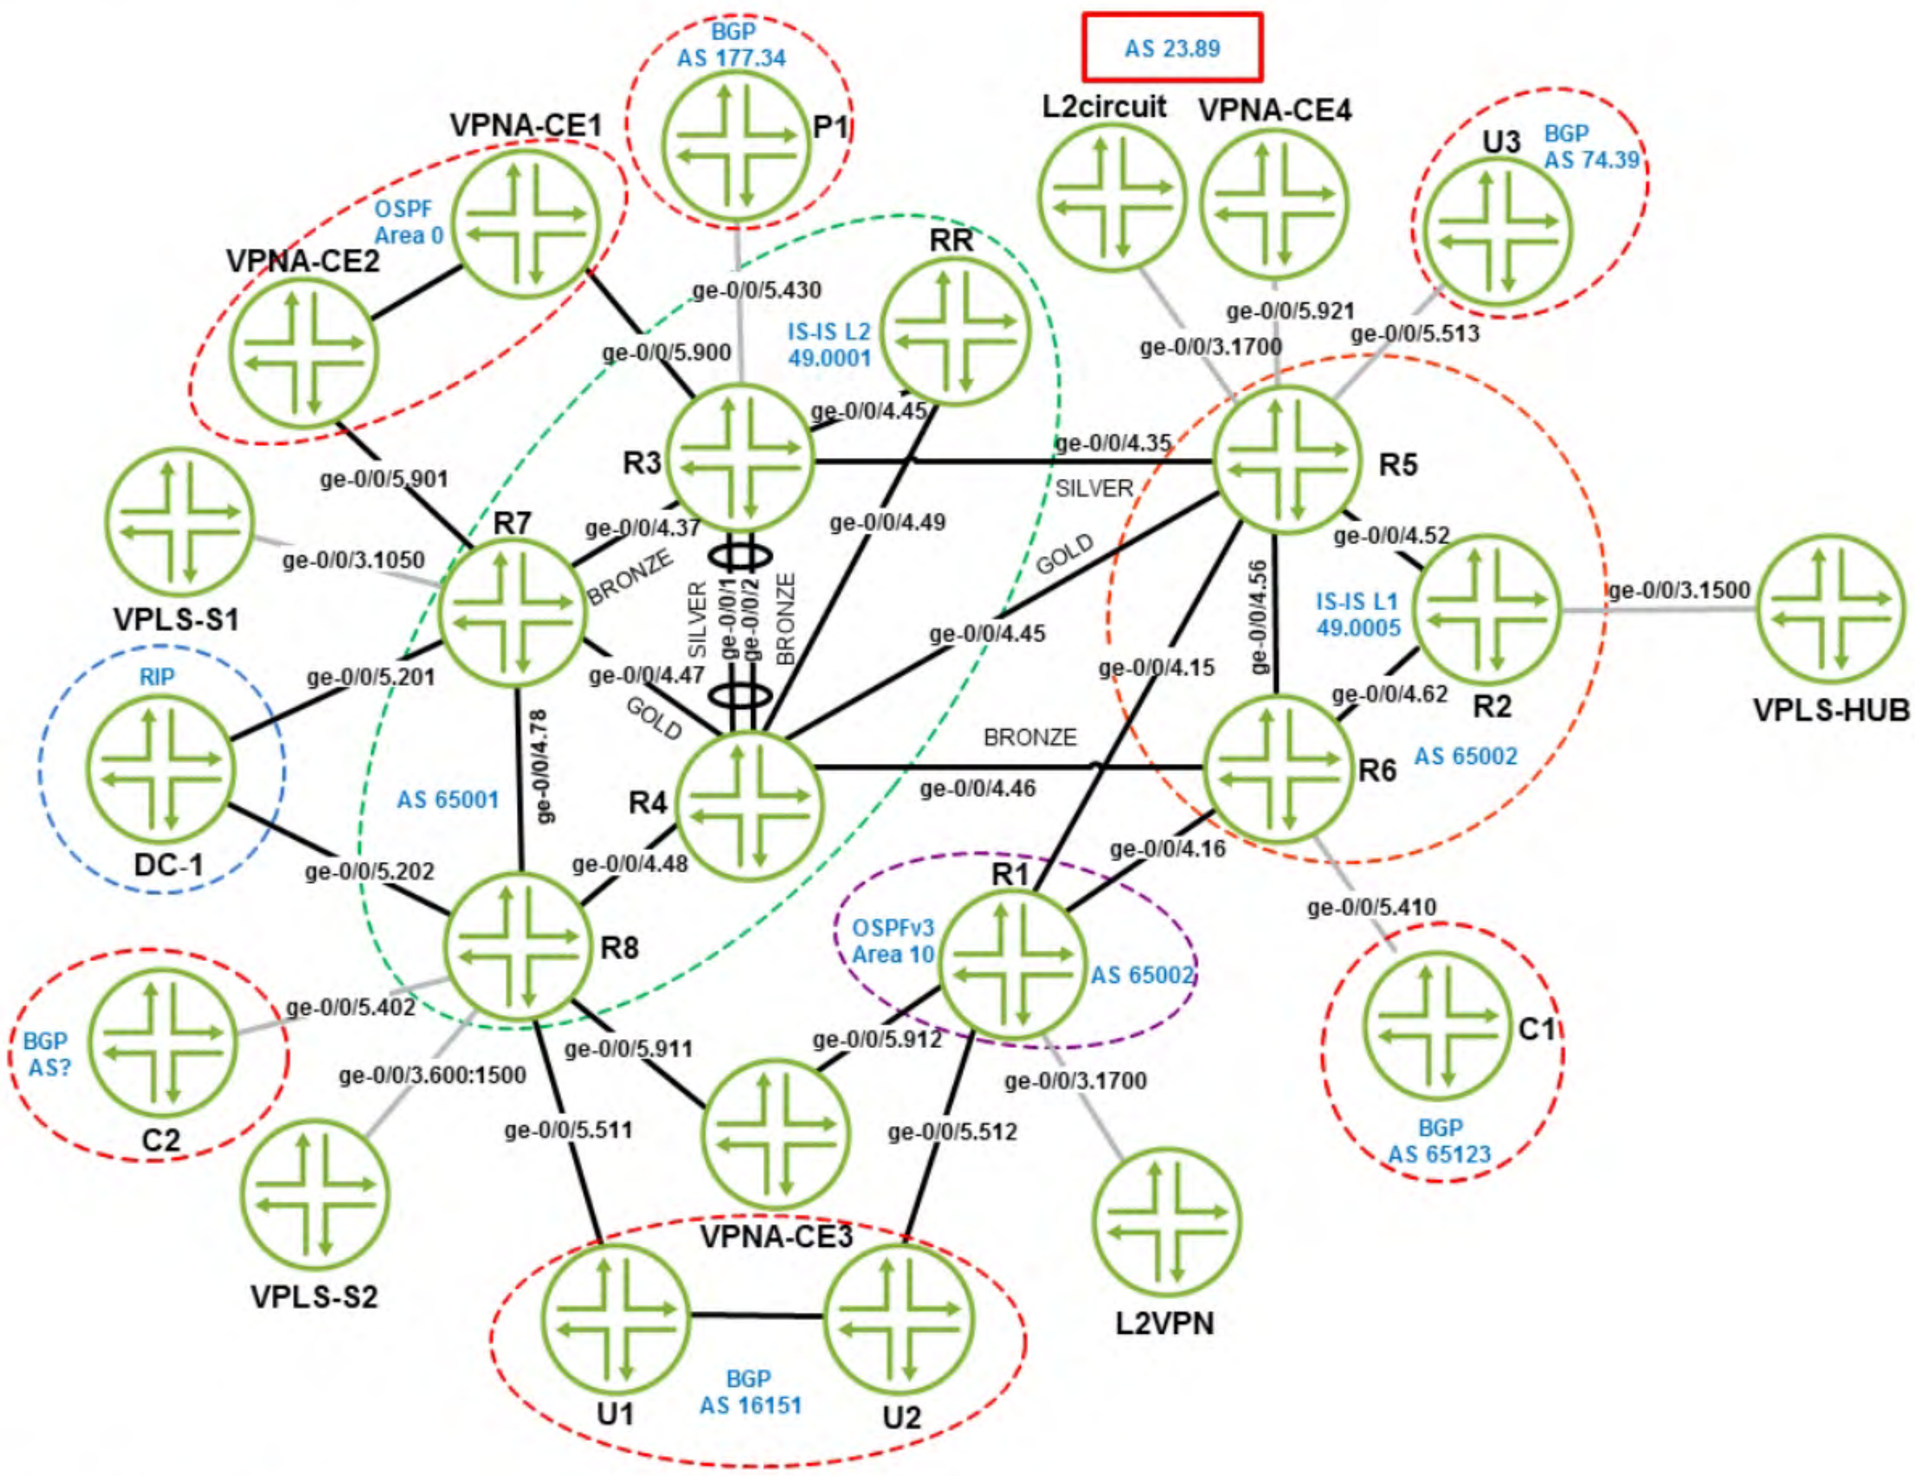
Table 36

© 2020 Juniper Networks, Inc. All rights reserved.

Super Lab 2

This lab is designed to combine many of the solutions discussed in the previous chapters into a complete practice lab exam. This lab is similar to a real JNCIE-SP lab exam experience. We recommend attempting this lab if you are able to complete the tasks in the previous chapters without any help.

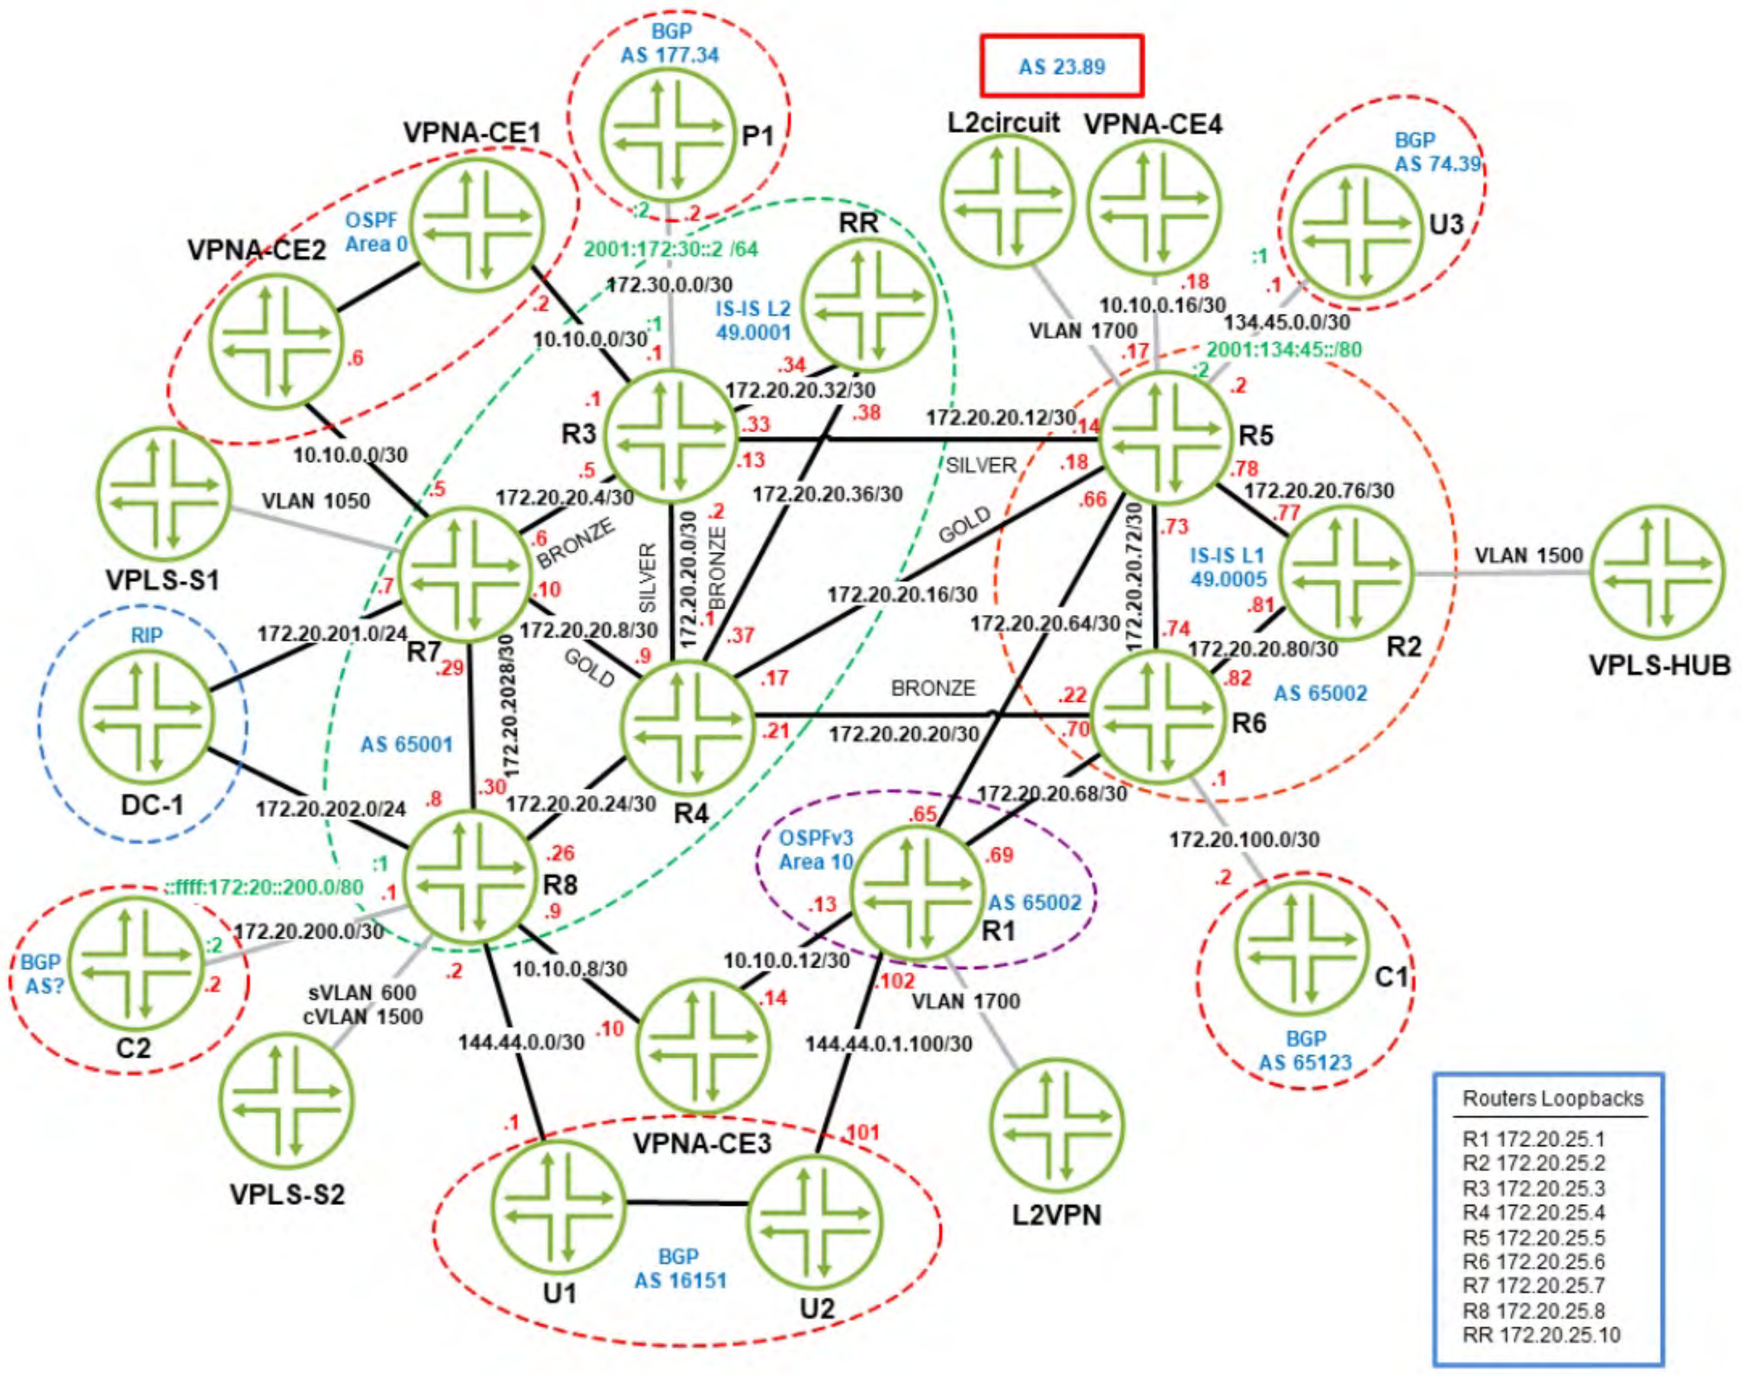
The physical diagram as shown below, visualizes the links between the routers. You are going to configure the devices with a darker color (PE and P). The other devices are virtualized on the VR-device, which configuration should not be changed. Use only the interfaces and VLANs on the diagram for connectivity between the routers.



Physical Diagram

© 2020 Juniper Networks, Inc. All rights reserved.

The logical diagram displays all the information needed for configuring the logical connectivity between the devices.



Logical Diagram

Like the real JNCIE-SP exam, some or all devices have partial configuration. Before you start with solving the Tasks, load the initial configurations provided with the workbook to all 10 physical devices. Use user *lab* password *lab123* for administrative access to every router.

Task 1: System Features

The task objectives: Configure the system related settings on all your devices. Ensure that your configuration meets the following criteria.

© 2020 Juniper Networks, Inc. All rights reserved.

1) Console access is preconfigured in the initial configuration on all routers. Your task is to increase the security of the console port. Ensure that console sessions logout automatically when the console cable is disconnected.
2) In addition to SSH protocol, enable and limit Telnet connections to 2 at a time.
3) Configure all routers to be accessible for file transfer using the FTP protocol. Limit simultaneous connections to only 2.
4) Configure all devices to authenticate users attempting to login via RADIUS server located on IP address 10.10.10.1, using secret workbook. The server is configured to accept requests only from the loopback address of the routers.
5) All user login attempts should fall back to the local used database if not configured on the RADIUS server.
6) Ensure that the reachability to 10.10.10.1/32 is provided even if the rpd process is not yet running. 10.10.1.200 is the gateway for the management segment.
7) Configure on all routers users noc1 and noc2, in such way that they do not fall back to the local database for authentication if the RADIUS server is not accessible. The users must be restricted to only view the configuration and issue operational commands with read-only access. Due to limited expertise of the users, ensure they are not able to request changes in operational mode, like software changes and power-off the devices.
8) In addition to that, restrict the use of those accounts only during the weekdays starting from 8am. to 6:30pm.
9) Configure user support for cases where additional control over the devices is required. This user should authenticate using password ops123 if the RADIUS server is inaccessible. The privileges assigned to the user must provide access to the network, reset processes, request changes, clear counters and view the configuration. However, you are not allowed to configure custom class for that specific user.
10) During login ensure a banner is displayed warning the users that unauthorized access is not allowed. Use the example below as shown and make sure it is displayed to the user before they are logged in.

##
WARNING: Unauthorized access is prohibited!
All sessions are monitored and recorded!
##

11) Configure all devices to send the copy of the configuration at each successful commit to an FTP server on address 10.10.10.1. Routers should copy the files in folder /tmp/upload/configurations using username jncie and password workbook.

© 2020 Juniper Networks, Inc. All rights reserved.

12) Administrators logged in with username support should be informed of any configuration change on the console.
13) Ensure that all possible messages are logged in the messages syslog file. Limit the message files to not more than 5 files, less than 10MB each.
14) Configure the interactive-commands syslog file to be transferred once per 12 hours to the FTP server located at address 10.10.10.1. Routers should copy the files in folder /tmp/upload/syslog using username jncie and password workbook. To overwrite the default archive size, increase the interactive-commands syslog file size to 20MB.
15) Configure the aggregated Ethernet interface connecting R3 and R4. Enable active monitoring of each physical interface link. You must make sure that the bundle always supports a capacity of minimum 1.5Gbps. The aggregated Ethernet interface must not forward traffic if the above requirement is not met.
16) Enable support for switching labeled packets, iso and IPv6 traffic on all logical interfaces (except management interface) on all routers except the RR. You must use the most efficient way configuration-wise to accomplish this task.
17) Enable graceful-restart on all routers for all routing protocols except BGP.
18) You must ensure that the control plane of all routers is protected. Configure a firewall filter using the following criteria:

- Allow main IPv4 protocols that are already configured and would be required by the subsequent tasks: SSH, BFD, OSPF, RIP, PIM, LDP, RSVP, RADIUS, ping and traceroute.
- Allow Telnet access from 172.20.20.0/24 and 172.20.25.0/24 subnets. To limit the size of the firewall filter, ensure that any additional changes to the allowed subnets does not require a change of the firewall filter.
- Allow only active FTP to the routers.
- You must accept BGP sessions only from configured neighbors. Any additional neighbors must be automatically accepted.
- Allow main IPv6 protocols that would be required by the subsequent tasks: OSPFv3, BGP, RIPng, ping and traceroute.
- Silently discard all other traffic received by the routers. Log the packets discarded by the filter. The log information must be available after a reboot

Task 2: Troubleshooting and Configuring IGP

1) Ensure all IS-IS adjacencies are in UP state. Make sure the following criteria are met:

© 2020 Juniper Networks, Inc. All rights reserved.

- All adjacencies that are configured with authentication must authenticate successfully. The system ID should be based on the loopback IP address.
- Any single interface should have only one adjacency in the appropriate level as shown in the diagram.
- Interfaces between area 49.0001 and 49.0005 should not elect DIS.
- Loopback interface addresses must be present as internal routes in the respected areas. R2 loopback should appear as an external address.

2) Configure R1 to exchange OSPF routing information with routers R5 and R6. Make sure that IPv4 and IPv6 loopback IP addresses between the OSPF domain and IS-IS domains are mutually reachable.
3) Advertise a default route to R1 so that it can reach the rest of the network. Ensure that a default route is advertised only when the interfaces to Level 2 area are operational. In such a way, that if the R6 interface to R4 fails the default route is not advertised anymore to R1.
4) Configure R7 and R8 to exchange RIP and RIPng routes from DC-1. A summary route representing all IPv4 and IPv6 loopback addresses should be advertised to DC-1 only when the connection to R3 and R4 is present. You are not allowed to amend the routing-options hierarchy on R7 and R8 to accomplish this task.
5) The routes received from DC-1 must be present on all IS-IS enabled routers.
6) R7 must be the gateway by default for the traffic to and from the DC-1 data center.
7) Make sure that R5 and R6 advertise only the Level 1 loopback addresses to Level 2 together with the R1 loopback address.
8) R2 must learn all RIP and RIPng prefixes and all the loopback addresses of the rest of the routers. Limit R2 to send traffic to destinations not explicitly present in the routing table.
9) The physical design of the network provides two end-to-end symmetrical paths from router R2 to router R7. To better utilize the combined bandwidth, make sure that IPv4 routes prefers the south path, through routers R4 and R6. Whereas the IPv6 traffic takes the north path, through routers R3 and R5. This requirement should only concern the traffic originated from routers R2 and R7.
10) Make sure that R5 and R6 receive routes with a metric higher than 63 from Level 2 routers R3 and R4.
11) Ensure that R7 and R8 will advertise overload bit set if DC-1 advertises more than 50 prefixes.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 3: Troubleshooting and Configuring BGP

1) Configure the necessary connectivity to the dedicated route-reflector. It already has partial configuration; however, you must ensure that appropriate IGP configuration is in place for providing connectivity to and from the route-reflector.
2) Design the IBGP network consisting of two confederations. AS65001 should include all L2 IS-IS routers, where AS65002 must include the rest of the routers. Configure redundant IBGP sessions between the two confederations.
3) Routers in confederations 65001 should use the dedicated route-reflector to peer with each other. The routers in the other confederation should peer in full mesh fashion.
4) Advertise an IPv4 summary route 172.20.25.0/28 from routers R3 and R4. The summary route should be advertised to all routers in your network.
5) Routers R3 and R4 should also advertise an IPv6 summary route 2001:fedc::/64 representing all IPv6 loopback addresses to the routers in AS65002. However, the summary must not be advertised to R1 and R2. Native IPv6 IBGP sessions are not allowed on the route-reflector.
6) Make sure that IBGP session on the route-reflector are not initiated on port 179.
7) Configure an EBGP session with the Upstream (U1, U2 and U3) peers. U1 and U2 are part of the same autonomous system, whereas U3 is in a different AS. However, there is connectivity and route exchange between the two autonomous-systems. Peers U1 and U2 provide only IPv4 connectivity. Whereas peer U3 provides IPv6 service also, make sure that you configure two separate sessions.
8) Configure an EBGP session with Provider (P1) peer. You should configure a single session capable of receiving IPv4 and IPv6 prefixes. You are not allowed to use the knob accept-remote-nexthop.
9) Customer (C1) is using IPv4 and IPv6 services provided by your network. Two separate BGP sessions must be used. However, due to hardware limitation at the customer side the IPv6 interface can use only link-local addresses. You must discover the peer IPv6 link-local address and establish EBGP session to the customer.
10) Customer (C2) is also using IPv4 and IPv6 services over two separate EBGP sessions. The autonomous-system number of the customer is not provided. You must discover the EBGP AS and configure the EBGP session using md5 authentication password jncie.
11) All IPv4 and IPv6 prefixes received from Upstream, Partner and Customer peers must be active and reachable on all devices in your AS. Unless they are filtered by policy in some of the subsequent tasks.

© 2020 Juniper Networks, Inc. All rights reserved.

12) Advertise the IBGP IPv4 summary route to all external peers except to P1. You are not allowed to use route policy to accomplish this.
13) R1 is chosen as backup gateway to AS 16151 therefore the routes advertised from peer U1 are prepended to enforce this in the outbound direction. Make sure that the inbound direction R1 is also less preferred than the path to router R8.
14) Make sure that router R8 is the preferred exit gateway for the same prefixes received from all Upstream peers. This requirement should not affect R8 routing table. In case the connectivity to U1 is lost, R1 and R5 should prefer their local connections to the upstreams.
15) Prefixes advertised by C1 should be uniquely marked in AS65001 only. Also, C1 routes must not be advertised to any EBGP peer except to peer C2. In reverse direction, peer C1 should only receive the C2 routes and the local summary routes.
16) Only advertise a summary route representing all C2 IPv4 routes to the internal routers and make sure it is advertised only to the Upstream neighbors and C1. Ensure all the routes received from C2 as originated from its own autonomous-system. In reverse direction, C2 must receive the C1 routes and Upstream routes.
17) The P1 routes received by R3 should not be advertised to any other peer. P1 should receive all RIP routes received from DC-1. Make sure that if DC-1 advertises more routes they are advertised automatically to P1.
18) Do not accept IPv4 routes that are part of RFC1918 and any default routes received from any of the Upstream peers. Prefixes with an as-path attribute longer than 10 should also be rejected.
19) All peerings with Upstream peers must enforce more aggressive measures against flapping prefixes with a mask length equal or greater than 20. These routes must be suppressed after two flap occurrences and stay maximum suppressed for twice the default time. Prefixes with a mask length of 10 or less must not be subject to damping.
20) No hidden (except filtered by applied policies) or unresolved routes should be present in AS23.89.
21) No static routes are allowed if not explicitly permitted by a particular task.
22) Ensure optimal routing in the network.

Task 4: MPLS Configuration

1) Configure RSVP on all routers except R2 and R8. The aggregated interface between R3 and R4 must not allow more than 1Gbps to be utilized by the LSPs.

© 2020 Juniper Networks, Inc. All rights reserved.

2) Ensure the RSVP messages between R1, R5 and R6 are authenticated with password key jncie-sp.
3) Configure the appropriate colors of the links according to the diagram.
4) Configure the LSPs using the criteria below:

- Configure bi-directional LSPs between R1 and R3, named r1-to-r3 and r3-to-r1. Make sure the shortest path between the endpoints is used. Configure automatic protection of the link between routers R3 and R5.
- Configure bi-directional LSPs between R1 and R7, named r1-to-r7 and r7-to-r1. Use two diverse paths, one that goes through R5 and the other through R6. By default, the path through R6 must be used as active. Make sure both paths are signaled and ready for use at any the time. The LSPs should not revert to the path that has failed. The two paths must be able to share the same bandwidth.
- Configure LSP from R5 to R3, named r5-to-r3, with secondary path built over links with colors GOLD and BRONZE, but avoiding links with color SILVER. The opposite direction LSP, named r3-to-r5, should have backup path with not more than two hops.
- Configure LSPs from R3 to R7 and in the opposite direction from R7 to R3, named r3-to-r7 and r3-to-r7. Configure link-protection to automatically protect the LSPs, but make sure R5 and R6 are not used as transit nodes.

5) Configure all RSVP enable routers to periodically optimize the path that the LSPs take through the network. The period should be set to two hours.
6) R8 and R2 only support the LDP protocol for label distribution. Connect the LDP enabled routers with rest of the routers, except router R1, in such way that they can effectively provide MPLS VPN service at a later stage.
7) Configure all the needed LSPs to accomplish redundant communication.
8) Ensure that all the MPLS enabled devices do not reveal the nodes inside the core.

Task 5: IPv6 Configuration

1) Ensure that Customers C1 and C2 can communicate using IPv6. Configure BGP to exchange the Customer routes using a single session for IPv4 and IPv6.
2) The local summary IPv6 route should also be advertised to C1 and C2. However, neither P1 nor U3 should be able to communicate with the Customers using IPv6.
3) To protect your customers, tunnel the IPv6 traffic between peers P1 and U3 in LSP.

© 2020 Juniper Networks, Inc. All rights reserved.

4) You are not allowed to add additional BGP families besides the ones enabled so far. Ensure that the transport label is not stripped off from the label stack, one hop before the egress PE.

Task 6: MPLS VPNs Configuration

1) Configure a L3VPN connecting the sites: VPNA-CE1, VPNA-CE2, VPNA-CE3 and VPNA-CE4. Sites CE1 and CE2 are using the OSPF protocol for PE-CE routing protocol, whereas the rest of the sites use EBGP with AS 64888. Use your local AS number for creating the route-distinguisher and route-target communities. Make sure you receive all prefixes from each site and share them between the sites.
2) Site CE1 and CE2 have a backdoor connection. Make sure that the two sites prefer to send traffic through the L3VPN connection. The backdoor link has a configured metric of 50.
3) Site CE3 is multihomed to both R1 and R8. Make sure the path through R8 is the preferred for sending and receiving traffic to CE3. You must use already applied policy to the route-reflector on R8 to complete the task.
4) Site C4 is configured to peer with AS 23.890 by mistake. You must connect the site to the MPLS network. However, you are not allowed to use local-as command to accomplish that. Ensure the mistaken AS number is not added to the AS path attribute on the routes send to the other sites.
5) Make sure the EBGP sites, CE3 and CE4, exchange routes through the core successfully without using as-override command. The two sites are configured to tolerate AS loops.
6) Use static routing to resolve any hidden routes in RIB inet.3 on the route-reflector.
7) Configure VPNA to support multicast transmission using NG-MVPN protocol and BGP for the control protocol with RSVP multicast data plane. The sources are located on the CE1 and CE2 routers. They provide source specific multicast service for group 239.255.255.3 from source IP addresses, 10.10.0.2 and 10.10.0.6.
8) Only R1 must be used for multicast traffic received by site CE3. Static routing is configured on CE3 to prefer R1 as a gateway to the sources.
9) Designate manually the receiving sites and ensure that they cannot be the root of a provider-tunnel. Configure the routers to choose the upstream multicast PE router based on the unicast route.
10) Ensure that the multicast PEs are using a dedicated tunnel interface for the multicast traffic and assign different route-target to the mvpn routes compared to the unicast vpn routes.

© 2020 Juniper Networks, Inc. All rights reserved.

11) Configure a VPLS connecting the VPLS-HUB site to the other two VPLS sites. Make sure you use BGP for VPN signaling and auto discovery.
12) The two VPLS spoke sites should be restricted to communicate only with HUB site. Because they are not allowed to communicate with each other, ensure that vpls connection between them is not available.
13) According to the diagram, the two VPLS spoke sites are configured with different VLANs than the HUB. You must normalize the VLAN used in the VPN to VLAN 1500. Do not make changes under the interface level on router R7 to achieve the requirement
14) Limit the mac-address table to not more than 20 hosts on the Spoke sites. Drop traffic is this limit is exceeded. Disable the mac learning function on the HUB site.
15) Configure interconnectivity between the L2VPN site and L2circuit site, using router R4 as stitching point. L2VPN site is using BGP signaled VPN whereas the L2circuit is using LDP based VPN.
16) Dedicate router R4 to provide connectivity between both L2 VPN sites. You are not allowed to use the lt-0/0/0 interface to accomplish the task.
17) Configure an additional LSP between R4 and R5 avoiding links with GOLD color. Make sure the L2circuit VPN is using that particular LSP for forwarding traffic.

© 2020 Juniper Networks, Inc. All rights reserved.

Appendices

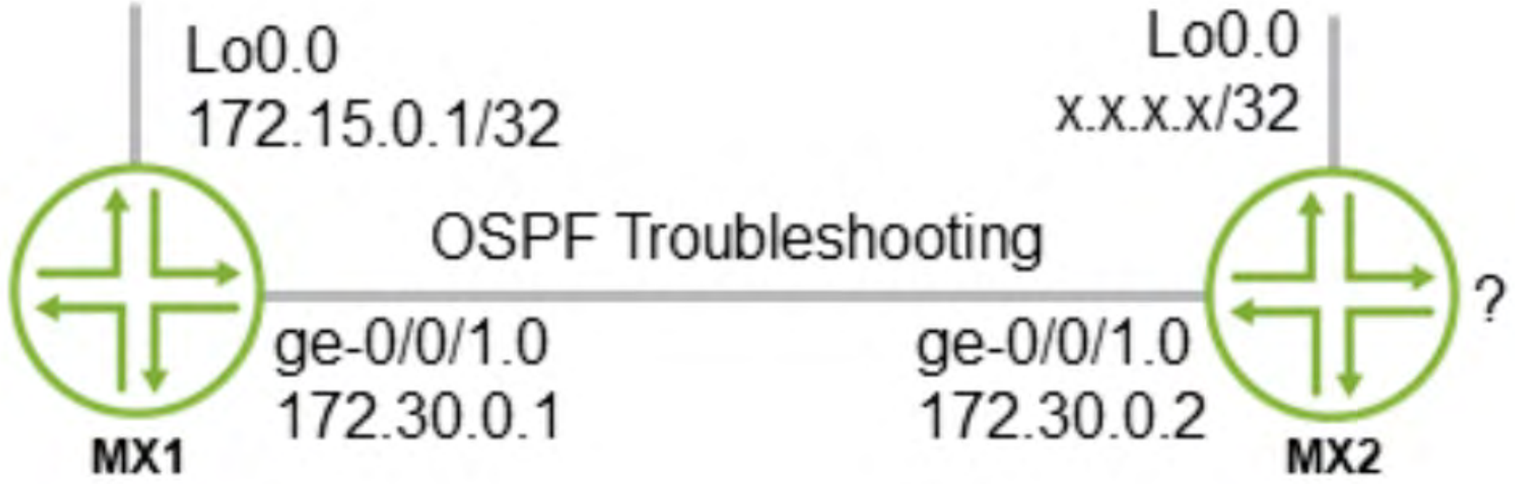
The following section provides the detailed explanation of the previous chapters and practice exam tasks. Here you will find details about the configuration and verification steps required to complete each section.

© 2020 Juniper Networks, Inc. All rights reserved.

Appendix - Additional Theory

OSPF Adjacency Troubleshooting

In this section, we demonstrate how to troubleshoot an OSPF neighbor adjacency using traceoptions.



There are two MX devices in the above topology. Assume MX1 is under our administrative control and MX2 is not. MX2 has been preconfigured with OSPF, but we do not have access to this device. Our goal is to establish an OSPF adjacency with MX2. The initial OSPF configuration for MX1 is very basic. Interface ge-0/0/1.0 and loopback 0.0 are both participating in the OSPF backbone area (0.0.0.0).

MX1's initial configuration:

```
interfaces {
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.30.0.1/30;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.30.15.1/32 {
                    primary;
                    preferred;
                }
            }
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-0/0/1.0;
        }
    }
}
```

1) Verify if you have IP connectivity to MX2.

```
root@R1# run ping 172.30.0.2
PING 172.30.0.2 (172.30.0.2): 56 data bytes
64 bytes from 172.30.0.2: icmp_seq=0 ttl=64 time=21.819 ms
64 bytes from 172.30.0.2: icmp_seq=1 ttl=64 time=1.226 ms
```

Super!

2) Verify is you have an OSPF adjacency with MX2 on interface ge-0/0/1.0.

```
root@R1# run show ospf neighbor interface ge-0/0/1.0
[edit]
```

Unfortunately, we do not have an adjacency with MX2. This means we must troubleshoot if MX2 has OSPF configured and try to determine its settings.

3) Enable OSPF traceoptions on MX1 and verify traceoptions output.

```
root@R1# set protocols ospf traceoptions file ospf
root@R1# set protocols ospf traceoptions flag all

root@R1# run monitor start ospf

[edit]
root@R1#

*** ospf ***
Apr  4 10:18:31.441041 OSPF packet ignored: area mismatch (0.0.0.99) from
172.30.0.2 on intf ge-0/0/1.0 area 0.0.0.0
Apr  4 10:18:31.441119 OSPF rcvd Hello 172.30.0.2 -> 224.0.0.5 (ge-0/0/1.0 IFL 70
area 0.0.0.0)
Apr  4 10:18:31.441189   Version 2, length 44, ID 172.30.15.2, area 0.0.0.99
Apr  4 10:18:31.441256   checksum 0x2fc8, authtype 0
Apr  4 10:18:31.441310   mask 255.255.255.252, hello_ivl 2, opts 0x12, prio 128
Apr  4 10:18:31.441424   dead_ivl 8, DR 0.0.0.0, BDR 0.0.0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

We can determine the following from the output related to OSPF adjacency formation:

- MX2 is sending OSPF packets to MX1.
- MX2 interface ge-0/0/1.0 participates in OSPF area 99.
- MX2 does not have authentication configured (auth type 0)
- MX2 interface ge-0/0/1.0 has an OSPF hello interval of 2 and dead interval of 8.

4) Change MX1's OSPF configuration to reflect MX2' settings.

```
root@R1# rename protocols ospf area 0 to area 99
root@R1# set protocols ospf area 0.0.0.99 interface ge-0/0/1.0 hello-interval 2
root@R1# set protocols ospf area 0.0.0.99 interface ge-0/0/1.0 dead-interval 8 *
```

* By default, if the dead-interval is not configured OSPF assumes a dead interval of 4 x the hello interval. In other words, although we did configure the dead-interval in our example, it is actually not needed.

5) Verify OSPF adjacency with MX2 (Venus).

```
root@R1# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.1       ge-0/0/1.0             Init      172.30.15.2      128     6
```

Now we see OSPF in the "init" state. This usually means that we have received an OSPF hello packet, but the other end (MX2) did not receive or at least did not accept our OSPF hello packet. Let's clear our OSPF process and check the traceoptions output if we missed an important clue. Its looks like we missed something.

6) Clear the ospf process and verify traceoptions output on MX1.

```
root@R1# run clear ospf neighbor

Apr  4 14:35:49.687959 OSPF rcvd Hello 172.30.0.2 -> 224.0.0.5 (ge-0/0/1.0 IFL 70
area 0.0.0.99)
Apr  4 14:35:49.688020   Version 2, length 44, ID 172.30.15.2, area 0.0.0.99
Apr  4 14:35:49.688084   checksum 0x0, authtype 0
Apr  4 14:35:49.688140   mask 255.255.255.252, hello_ivl 2, opts 0x12, prio 128
Apr  4 14:35:49.688191   dead_ivl 8, DR 0.0.0.0, BDR 0.0.0.0
```

Interface ge-0/0/1.0 on MX2 has been configured with the link type set to "p2p", since no DR/BDR election is desired on ge-0/0/1.0 interface. It is a direct connection between the devices. Let's change the OSPF interface type to "p2p" on our ge-0/0/1.0 interface.

© 2020 Juniper Networks, Inc. All rights reserved.

7) Change OSPF interface type to p2p on ge-0/0/1.0 and verify OSPF neighborship.

```
root@R1#  set protocols ospf area 0.0.0.99 interface ge-0/0/1.0 interface-type p2p

root@R1# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.2       ge-0/0/1.0             Exchange  172.30.15.2      128  6
```

The OSPF neighborship with MX2 is in "Exchange" state, this means that at least both OSPF routers have seen each other's hello packets. OSPF "Exchange" state is usually related to MTU issue's or other layer 2 issues. We can rule out the latter one, since we were able to ping MX2.

8) Verify OSPF traceoptions output on MX1 to verify if there is an MTU issue.

```
Apr  4 14:55:24.717198 OSPF rcvd DbD 172.30.0.2 -> 224.0.0.5 (ge-0/0/1.0 IFL 70
area 0.0.0.99)
Apr  4 14:55:24.717267   Version 2, length 32, ID 172.30.15.2, area 0.0.0.99
Apr  4 14:55:24.717317   checksum 0x0, authtype 0
Apr  4 14:55:24.717386   options 0x52, i 1, m 1, ms 1, r 0, seq 0xac159be3, mtu
9178
```

9) Check our local IP MTU on interface ge-0/0/1.0.

```
root@R1# run show interfaces ge-0/0/1.0 | match MTU
    Protocol inet, MTU: 1500
```

It seems there is an IP MTU mismatch between MX1 and MX2. MX2 appears to have set the IP MTU to 9178 (jumbo) on interface ge-0/0/1.0.

10) Change the IP mtu on interface ge-0/0/1.0 to 9178 and verify OSPF neighborship.

There are two ways to change the IP MTU. We can change the interface MTU to 9192 or change the IP MTU. Please note that the interface MTU is 14 bytes more than the IP MTU due to encapsulation overhead. Note: if the interfaces used vlan-tagging the difference between the IP MTU and interface MTU is 18 instead of 14 bytes. This is because of the additional 4 bytes for the vlan tag.

```
root@R1# set interfaces ge-0/0/1 mtu 9192

or

root@R1# set interfaces ge-0/0/1.0 family inet mtu 9178
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
root@R1# commit
commit complete
```

11) Verify if the OSPF adjacency is established

```
root@R1# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.2       ge-0/0/1.0             Full      172.30.15.2      128     7
```

Finally, our OSPF neighborship is in FULL state.

12) Verify if we receive OSPF routes from MX2.

```
root@R1# run show route table inet.0 protocol ospf

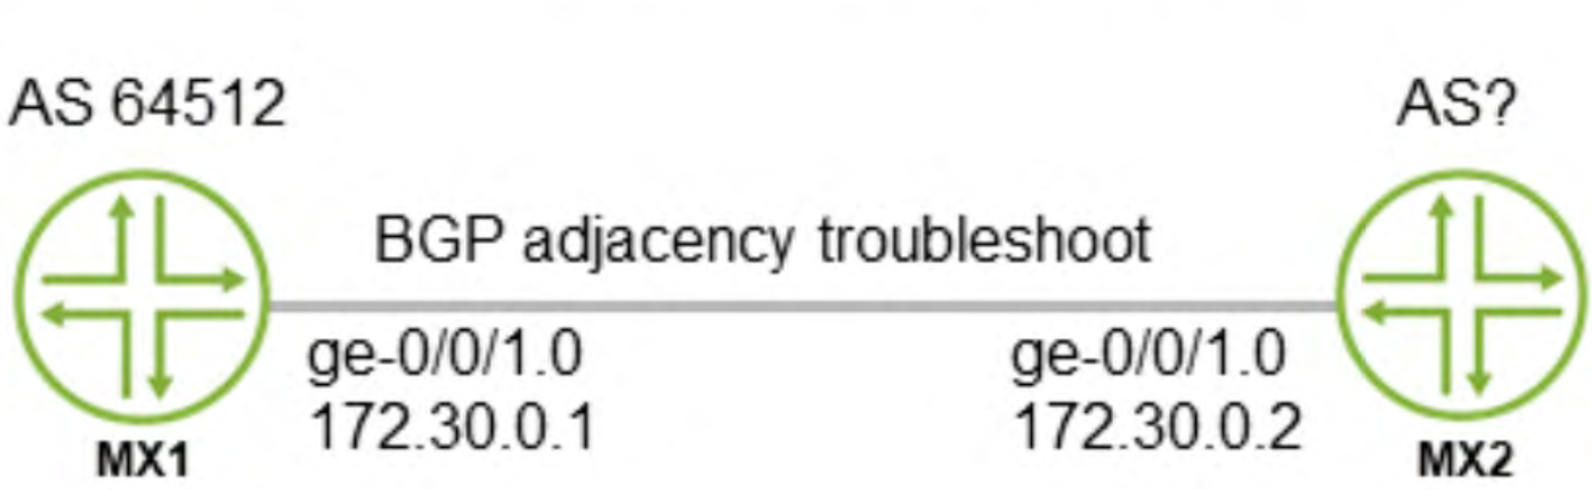
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.15.2/32     *[OSPF/10] 00:12:42, metric 1
                    > to 172.30.0.2 via ge-0/0/1.0
224.0.0.5/32       *[OSPF/10] 06:52:18, metric 1
                      MultiRecv
```

That's it. We have managed to get the OSPF adjacency up without access to MX2! Note that it's also possible to use the "**monitor traffic interface x/y/z extensive**" command to "debug" OSPF adjacencies.

BGP Adjacency Troubleshooting

In this section, we will troubleshoot an EBGP adjacency issue using traceoptions.



There are two MX devices in the above topology. Assume MX1 is under our administrative control and MX2 is not. MX2 has been preconfigured with an EBGP session towards MX1,

© 2020 Juniper Networks, Inc. All rights reserved.

but we do not have access to this device, and we do not know MX2 autonomous system number. Our goal is to establish an EBGP adjacency with MX2.

MX1 initial configuration.

```
interfaces {
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.30.0.1/30;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.30.15.1/32 {
                    primary;
                    preferred;
                }
            }
        }
    }
}
```

Let's verify if we have layer 3 connectivity to MX2.

```
[edit]
root@R1# run ping 172.30.0.2
PING 172.30.0.2 (172.30.0.2): 56 data bytes
64 bytes from 172.30.0.2: icmp_seq=0 ttl=64 time=25.389 ms
64 bytes from 172.30.0.2: icmp_seq=1 ttl=64 time=1.278 ms
```

It appears we have layer 3 connectivity to MX2, so that's good. This means that MX2 can reach MX1 and hence also able to send BGP open messages to MX1 which we can monitor using traceoptions.

Enable BGP traceoptions on MX1 to see if we can retrieve MX2 autonomous system number and configure MX2 as EBGP neighbor with a fake peer-as number.

```
root@R1# show protocols bgp
traceoptions {
    file bgp;
    flag open;
}
group ebgp {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    neighbor 172.30.0.2 {
        peer-as 1;
    }
}
```

Check BGP adjacency with MX2.

```
root@R1# run show bgp summary
Groups: 1 Peers: 1 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.0.2            64555          7          6       0       2          23
Active
```

As expected, our neighborship with MX2 is not established.

Enable BGP traceoptions to see if we can retrieve MX2 AS number.

```
root@R1# run monitor start bgp

Feb  4 20:08:41.342020 bgp_process_open:2822: NOTIFICATION sent to 172.30.0.2
(External AS 1): code 2 (Open Message Error) subcode 2 (bad peer AS number),
Reason: peer 172.30.0.2 (External AS 1) claims 64555, 1 configured
```

We can determine from the traceoptions output that MX2 AS number is "64555".

Reconfigure the peer-as statement

```
root@R1# show protocols bgp
traceoptions {
    file bgp;
    flag open;
}
group ebgp {
    neighbor 172.30.0.2 {
        peer-as 64555;
    }
}
```

Check the BGP peering with MX2 again!

```
root@R1# run show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.0.2             64555         33         33       0       2        3:36
0/0/0/0                0/0/0/0

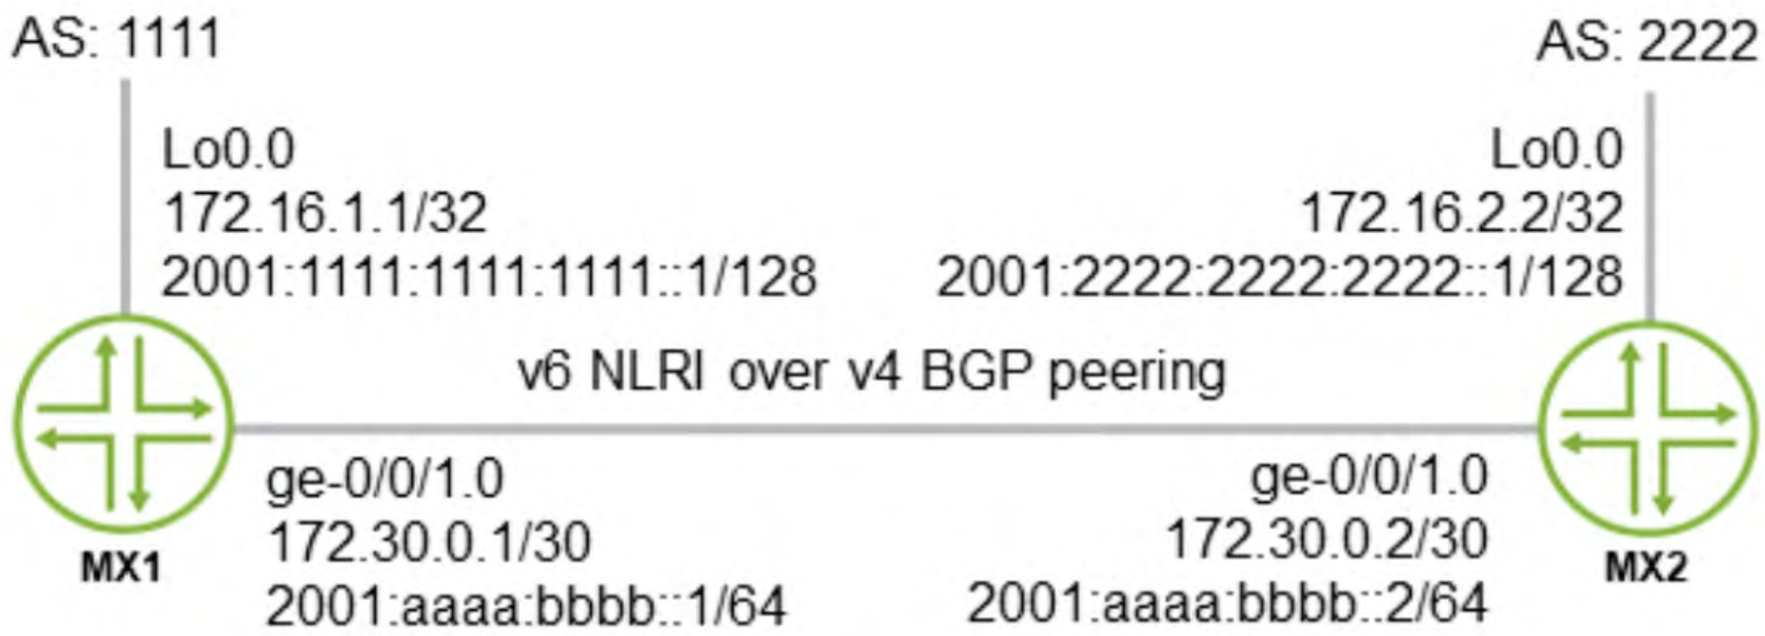

root@R1# run show bgp neighbor 172.30.0.2
Peer: 172.30.0.2+179 AS 64555  Local: 172.30.0.1+49402 AS 64512
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: Cease
Holdtime: 90 Preference: 170
Number of flaps: 2
Last flap event: RecvNotify
Error: 'Cease' Sent: 1 Recv: 1
Peer ID: 172.30.0.2      Local ID: 173.30.15.1      Active Holdtime: 30
Keepalive Interval: 10         Peer index: 0
BFD: disabled, down
Local Interface: ge-0/0/1.0
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer supports 4 byte AS extension (peer-as 64555)
Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:              0
  Received prefixes:            0
  Accepted prefixes:            0
  Suppressed due to damping:    0
  Advertised prefixes:          0
Last traffic (seconds): Received 7    Sent 6    Checked 12
Input messages:  Total 37     Updates 2       Refreshes 0     Octets 753
Output messages: Total 37     Updates 0       Refreshes 0     Octets 829
Output Queue[0]: 0
Trace options: open
Trace file: /var/log/bgp size 0 files 10
```

The BGP peering is established!

© 2020 Juniper Networks, Inc. All rights reserved.

BGP IPV6 NLRI over IPV4 Peering

In the following example, we demonstrate how to configure V6 NLRI exchange over IPv4 BGP peerings.



In the above topology, there are two routers: MX1 is an ASBR for BGP Autonomous System (AS): 1111 and MX2 is the ASBR for BGP AS: 2222. There is an ipv4 EBGP peering configured between MX1 and MX2. This ipv4 EBGP peering is also used to exchange IPv6 NLRI. Each device will announce its loopback IP address (v4 and v6) to the other ASBR.

MX1 initial configuration:

```
root@R1#show interfaces
ge-0/0/1 {
unit 0 {
family inet {
address 172.30.0.1/30;
}
family inet6 {
address 2001:aaaa:bbbb::1/64;
}
}
}
lo0 {
unit 0 {
family inet {
address 172.16.1.1/32 {
primary;
preferred;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
}
}
family inet6 {
address 2001:1111:1111:1111::1/128;
}
}
}

root@R1# show protocols bgp
group ebgp {
type external;
family inet {
unicast;
}
family inet6 {
unicast;
}
export myloopback;
neighbor 172.30.0.2 {
peer-as 2222;
}
}

root@R1# show policy-options policy-statement myloopback
from interface lo0.0;
then accept;

root@R1# show routing-options
autonomous-system 1111;
```

Please note that we configured an IPv4 neighborship with MX2 for IPv4 NLRI (family inet unicast) and IPv6 NLRI (family inet6 unicast). As you can see, we did not configure a native IPv6 peering with MX2!

Verify if our BGP peering with MX2 is in the **Established** state.

```
root@R1# run show bgp neighbor 172.30.0.2

Peer: 172.30.0.2+49898 AS 2222 Local: 172.30.0.1+179 AS 1111
Type: External    State: Established    Flags: <Sync>
←
Last State: OpenConfirm   Last Event: RecvKeepAlive
Last Error: Cease
Export: [ myloopback ]
Options: <Preference AddressFamily PeerAS Refresh>
Address families configured: inet-unicast inet6-unicast
Holdtime: 90 Preference: 170
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Number of flaps: 2
Last flap event: Stop
Error: 'Cease' Sent: 3 Recv: 0
Peer ID: 172.16.2.2      Local ID: 172.16.1.1        Active Holdtime: 90
Keepalive Interval: 30         Peer index: 0
BFD: disabled, down
Local Interface: ge-0/0/1.0
NLRI for restart configured on peer: inet-unicast inet6-unicast
NLRI advertised by peer: inet-unicast inet6-unicast
NLRI for this session: inet-unicast inet6-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Peer supports 4 byte AS extension (peer-as 2222)
Peer does not support Addpath
Table inet.0 Bit: 10000
RIB State: BGP restart is complete
Send state: in sync
Active prefixes:              1
Received prefixes:            1
Accepted prefixes:            1
…
<output omitted>
```

As you can see the EBGP peering with MX2 is in the `established` state. We also notice that the NLRI received and used for this session is: `inet-unicast` and `inet6-unicast`.

This is because we and the remote ASBR configured the "`family inet unicast`" and "`family inet6 unicast`" NLRI's under the EBGP peer-group.

MX2 has been configured the same say as MX1 and announces its ipv4 and ipv6 loopback addresses into EBGP.

```
root@MX2# run show route advertising-protocol bgp 172.30.0.1

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.2.2/32           Self                                    I

inet6.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix                  Nexthop              MED     Lclpref    AS path
2001:2222:2222:2222::1/128
*                         Self
```


© 2020 Juniper Networks, Inc. All rights reserved.

So far so good. Let's verify if we receive the ipv4 and ipv6 loopback addresses from MX2.

```
root@R1# run show route receive-protocol bgp 172.30.0.2 extensive

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
* 172.16.2.2/32 (1 entry, 1 announced)
Accepted
Nexthop: 172.30.0.2
AS path: 2222 I

inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

We can confirm that we receive the ipv4 prefix from MX2 with a next-hop of 172.30.0.2. Unfortunately, we do not receive the ipv6 prefix from MX2.

Configure BGP traceoptions and verify if an issue is reported.

```
root@R1# set protocols bgp  traceoptions file bgp
root@R1# set protocols bgp  traceoptions flag route
root@R1# set protocols bgp  traceoptions flag
root@R1# commit

May 20 20:02:37.357005 bgp_nexthop_sanity: peer 172.30.0.2 (External AS 2222) next
hop ::ffff:172.30.0.2 unexpectedly remote, ignoring routes in this update
```

We can tell that the BGP sanity check did not accept a prefix as the next-hop is an IPv4 mapped IPv6 address. MX1 does not have a valid route installed forwards ::ffff:172.30.0.2 so this is as expected.

Recall that we have configured IPv6 NLRI over an IPv4 BGP session. When exchanging IPv6 NLRI over an IPv4 MP-BGP peering session Junos will encode the BGP next-hop in IPv4–mapped format. Below is an example of an IPv4-mapped address (RFC 3513).

::ffff:172.16.1.1

If an IPv4-mapped IPv6 address is used as the BGP next hop, this means that this address must be reachable for the learned prefixes to be accepted. An ASBR, by default, will not accept a next-hop which is not directly connected.

In the following section we will demonstrate two scenarios how to configure IPv6 NLRI exchange over and IPv4 peering.

Option 1:

© 2020 Juniper Networks, Inc. All rights reserved.

The easiest way to ensure that MX1 accepts and installs the IPv6 prefixes is to configure an ipv4-mapped address on MX1 and MX2 so that the next-hop is reachable. Please note that this solution requires that you can also configure the remote EBGP peer.

```
root@R1# show interfaces ge-0/0/1.0 family inet6
address ::ffff:172.30.0.1/127;

and

root@MX2# show interfaces ge-0/0/1.0 family inet6
address ::ffff:172.30.0.2/127;
```

We also must enable V4 mapped packet processing in Junos.

```
root@R1#set system allow-v4mapped-packets
```

Check if we receive the ipv4 loopback address from MX2.

```
root@R1# run show route receive-protocol bgp 172.30.0.2 extensive

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
* 172.16.2.2/32 (1 entry, 1 announced)
Accepted
Nexthop: 172.30.0.2
AS path: 2222 I

inet6.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)

* 2001:2222:2222:2222::1/128 (1 entry, 1 announced)
Accepted
Nexthop: ::ffff:172.30.0.2
AS path: 2222 I
```

The above output shows that MX received and installed the IPv4 prefix.

Now the most important part. Verify if MX1 accepts and installs the IPv6 prefix.

```
root@R1#run show route table inet6.0 extensive 2001:2222:2222:2222::1/128

inet6.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
2001:2222:2222:2222::1/128 (1 entry, 1 announced)
TSI:
KRT in-kernel 2001:2222:2222:2222::1/128 -> {::ffff:172.30.0.2}
*BGP    Preference: 170/-101
Next hop type: Router, Next hop index: 574
Address: 0x155c860
Next-hop reference count: 3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Source: 172.30.0.2
Next hop: ::ffff:172.30.0.2 via ge-0/0/1.0, selected

State: <Active Ext>
Local AS:  1111 Peer AS:  2222
Age: 1:54
Task: BGP_2222.172.30.0.2+179
Announcement bits (2): 0-KRT 2-Resolve tree 2
AS path: 2222 I
Accepted
Localpref: 100
Router ID: 172.16.2.2
```

Yes, the IPv6 prefix is installed in the inet6.0 table. We've seen that when we configure an IPv4 mapped IPv6 address the next-hop is resolved, hence the route is learned on MX1 and installed in the inet6 routing-table.

Let's remove the IPv4 mapped IPv6 addresses we configured previously and try the second option as explained in the beginning of this section.

```
Root@R1#delete interfaces ge-0/0/1.0 family inet6 address ::ffff:172.30.0.1/126
root@R1# delete system allow-v4mapped-packets
Root@R1#commit

and

Root@MX2#delete interfaces ge-0/0/1.0 family inet6 address ::ffff:172.30.0.2/126
root@MX2# delete system allow-v4mapped-packets
Root@MX2#commit
```

Confirm that indeed the IPv6 prefix has disappeared.

```
root@R1# run show route receive-protocol bgp 172.30.0.2

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.2.2/32           172.30.0.2                             2222 I

inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

root@R1# run show route table inet6 hidden

inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
[edit]
```

Ok, we are back at the original issue where MX1 will not accept and thus not install the IPv6 prefix.

© 2020 Juniper Networks, Inc. All rights reserved.

The second option is to not use IPv4 mapped IPv6 addresses on the links between MX1 and MX2. To make this work we must ensure that:

- MX1 accepts the ::ffff:172.30.0.2 prefix.
- MX1 rewrites the next hop to an ipv6 address that is usable. In our case this will be the native ipv6 address of MX2 on ge-0/0/1.0.

To have MX1 accept next-hop values that are not directly connected, we can use the "accept-remote-nexthop" command.

```
root@R1# set protocols bgp group ebgp accept-remote-nexthop
root@R1# commit
```

MX1 is certainly still not able to install the prefix in the inet6 routing table as we did not rewrite the next-hop **::ffff:172.30.0.2** to a native ipv6 address yet. At least we should see the prefix learned from MX2 but is hidden because the next-hop :ffff:172.30.0.2 is not reachable.

```
root@R1# run show route table inet6.0 hidden extensive

inet6.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
2001:2222:2222:2222::1/128 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Next hop type: Unusable
                Address: 0x113bc8c
                Next-hop reference count: 1
                State: <Hidden Ext>
                Local AS:  1111 Peer AS:  2222
                Age: 43
                Task: BGP_2222.172.30.0.2+179
                AS path: 2222 I
                Accepted
                Localpref: 100
                Router ID: 172.16.2.2
                Indirect next hops: 1
                        Protocol next hop: ::ffff:172.30.0.2
                        Indirect next hop: 0 -
```

Now we must make sure to rewrite the next-hop. We are going to use an BGP import policy to address this issue.

```
root@R1# show policy-options policy-statement fixnexthop
from protocol bgp;
then {
    next-hop 2001:aaaa:bbbb::2;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
root@R1# set protocols bgp group ebgp import fixnexthop

[edit]
root@R1# commit
```

In our BGP traceoptions output we notice that the next-hop changed!

```
May 20 19:59:59.149372 CHANGE   2001:2222:2222:2222::1/128  nhid 0 gw
2001:aaaa:bbbb::2 BGP      pref 170/-101 metric  ge-0/0/1.0 <Active Ext>  as 2222
```

We see that the prefix is now received and accepted by the BGP sanity check.

Verify if the IPv6 prefix is now correctly installed!

```
root@R1# run show route table inet6 extensive 2001:2222:2222:2222::1/128

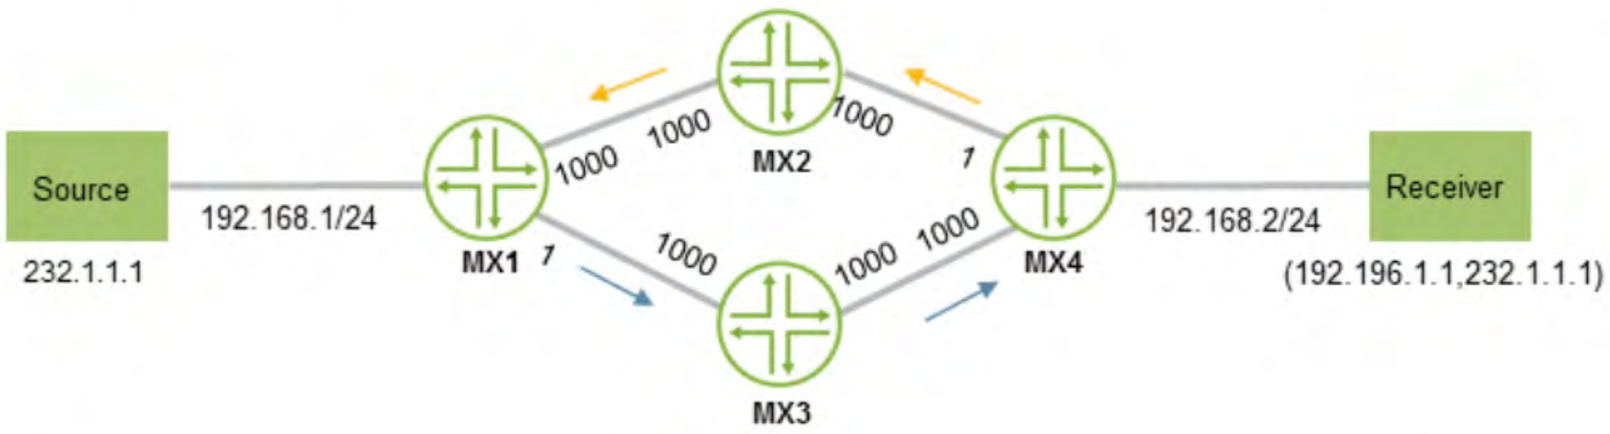
inet6.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
2001:2222:2222:2222::1/128 (1 entry, 1 announced)
TSI:
KRT in-kernel 2001:2222:2222:2222::1/128 -> {2001:aaaa:bbbb::2}
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 568
                Address: 0x155c860
                Next-hop reference count: 3
                Source: 172.30.0.2
                Next hop: 2001:aaaa:bbbb::2 via ge-0/0/1.0, selected
                State: <Active Ext>
                Local AS:  1111 Peer AS:  2222
                Age: 8:48
                Task: BGP_2222.172.30.0.2+179
                Announcement bits (2): 0-KRT 2-Resolve tree 2
                AS path: 2222 I
                Accepted
                Localpref: 100
                Router ID: 172.16.2.2
```

The “accept-remote-nexthop” command together with the “fixnexthop” policy ensured that the IPv6 prefix is installed in the inet6.0 table.

© 2020 Juniper Networks, Inc. All rights reserved.

Troubleshooting: Multicast Traffic Engineering Using RIB-groups

In the following scenario, we will troubleshoot a multicast RPF issue with given restrictions.



A multicast receiver attached to MX4 would like to join source specific multicast (SSM) group 232.1.1.1 send by multicast source 192.168.1.1. Assume the following requirements:

- **Unicast traffic from MX1 to MX4 should always transit MX3.**
- **Unicast traffic from MX4 to MX1 should always transit MX2.**

To meet the unicast flow requirement the IGP metrics for prefixes in the inet.0 table are tuned on MX1 and MX4 (metric 1). For some reason the multicast traffic is not received by the receiver attached to MX4.

Verify the PIM signaling in our network on MX4 and MX1:

```
root@MX4# run show pim join inet 232.1.1.1
Instance: PIM.master Family:   INET
R = Rendezvous Point Tree,   S = Sparse,  W = Wildcard

Group:  232.1.1.1
          Source: 192.168.1.1
          Flags: sparse
          Upstream interface: unknown  (no nexthop)


root@R1# run show pim source inet 192.168.1.1
Instance:      PIM.master  Family:    iNET

Source 192.168.1.1
          Prefix 192.168.1.0/24
          Upstream interface ge-0/0/1.0
          Upstream neighbor 192.168.1.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

We can determine that MX4 has a reverse path forwarding (RPF) failure for multicast group 232.1.1.1.

```
root@MX4# run show multicast route group 232.1.1.1 extensive
Family:     INET

Group:  232.1.1.1
          Source:  192.168.1.1/32
          Upstream interface: ge-0/0/0.0
          Downstream interface list:
                    ge-0/0/1.0
          Session description: Source specific multicast
          Statistics:  0 kBps,   0 pps,   0 packets



root@MX4#  run show multicast usage
Group                              Sources  Packets
Bytes
232.1.1.1                            1              0
0

Prefix                                  /len  Groups  Packets
Bytes
192.168.1.1                        /32    1            0
0
```

It seems that no multicast traffic is flowing through our network.

Verify the RPF table on MX1 and MX4.

```
root@R1#  run show multicast rpf 192.168.2.1
Multicast  RPF  table:  inet.0  ,   32 entries

192.168.2.0/24
          Protocol:  OSPF
          Interface:  ge-0/0/3.0     ←This is the interface connected to MX3


root@MX4#  run show multicast rpf 192.168.1.1
Multicast  RPF  table:  inet.0  ,   34 entries

192.168.1.0/24
          Protocol:  OSPF
          Interface:  ge-0/0/2.0     ←This is the interface connected to MX2
```


© 2020 Juniper Networks, Inc. All rights reserved.

When a multicast packet enters an interface, the router will check the reverse path for the packet. The reverse path for the multicast packet must be on the same interface as where the multicast packet arrived on (symmetrical forwarding). If this check fails, the packet is dropped. Multicast RPF check is needed to break possible multicast loops in the network.

The above RPF output clearly shows that there is an RPF failure in this network. Due to the requirement that unicast traffic from MX1 to MX4 must transit MX3 and traffic from MX4 to MX1 must transit MX2 the IGP (OSPF) metrics in the inet.0 table have been changed in our network (see topology diagram). This is fine; however, it introduces an RPF failure in this scenario. It also means that we cannot modify the inet.0 table to fix the RPF failure as this would break our unicast flow requirement.

Recall that Junos has a dedicated table for multicast RPF lookups, the inet.2 table. If we ensure that Protocol Independent Multicast (PIM) uses the inet.2 table for RPF checks, we can manipulate multicast RPF check without breaking the unicast routing requirement.

Create two rib-groups. The first rib-group "`myrpffix`" imports the inet.0 and inet.2 table. The second rib-group only imports the inet.2 table.

```
root@MX4# show routing-options rib-groups
myrpffix {
    import-rib [ inet.0 inet.2 ];
}
fullrpf {
    import-rib inet.2;
}
```

Create a static route in the inet.2 table to ensure that MX4 uses MX3 as the next-hop for prefix 192.16.1.0/24 and passes the RPF check.

```
root@MX4# set routing-options rib inet.2 static route 192.168.1.0/24 next-hop <R3
interface>
```

Ensure that the "interface routes" are used in "`myrpffix`" rib-group. This is needed as the next-hop for the previously created static route in inet.2 must be resolvable.

```
root@MX4# set routing-options interface-routes rib-group myrpffix
```

Ensure that the protocol independent multicast (PIM) protocol uses the `fullrpf` rib-group (inet.2 table) to perform RPF checks.

```
root@MX4# set protocols pim rib-group fullrpf
```


© 2020 Juniper Networks, Inc. All rights reserved.

That's it! We ensured that PIM uses the inet.2 table for RPF check. The inet.2 table has a static route configured to fix the next-hop. Since we use the inet.2 table and not the inet.0 table we did not break our unicast flow requirement.

Advanced Firewall Filtering

To protect the control plane for Junos devices you typically apply a firewall filter the loopback interface for the address families that require protection. For the JNCIE exam it might be that you need to apply filtering for IPv4 control plane and IPV6 control plane directed traffic.

For the following scenario our goal is to create a firewall filter "term" which allows only BGP traffic from our current peers and our solution must also ensure to automatically add new peers when they are added in our BGP peer groups.

It's simple to create a firewall filter rule and match each configured BGP peer as listed in our peer-groups. Unfortunately, this method does not solve the requirement to also add future peers automatically. Fortunately, with Junos you can create dynamic prefix lists with the "apply-path" feature.

The apply-path feature makes it possible to dynamically update a prefix list based on matching certain parts in the configuration. For example, you can match all configured DNS servers or all configured BGP peers. This also ensures that there is no need to constantly update a prefix-list when new BGP peers are added. Further is will reduce the possibility of errors or network outages due to a typo in a manually configured prefix-list.

Let's get started. The following output shows our configured BGP peer groups.

```
lab@MX# show protocols bgp group ibgp
type internal;
local-address 192.168.1.1;
family inet {
    unicast;
}
neighbor 192.168.1.2;

lab@MX# show protocols bgp group ebgp
type external;
neighbor 172.16.1.1 {
    export [ myexport ];
    peer-as 2222;
}
```

There are two BGP peer groups configured. One for IBGP, one for EBGP.

© 2020 Juniper Networks, Inc. All rights reserved.

Instead of a regular prefix list we use a prefix list with the `apply-path` feature to ensure that new BGP peers are automatically added to our prefix list, when configured under the bgp peer-group hierarchy.

```
[edit policy-options]
lab@MX# show
prefix-list bgp-peers {
    apply-path "protocols bgp group <*> neighbor <*>";
```

We created a prefix-list called “bgp-peers”. The apply-path statement matches ALL groups <*> and all neighbors <*> under the “protocols bgp group” hierarchy.

You can verify if the apply-path prefix-list is working as expected with the “`display inheritance`” appended to the “show policy prefix-list” command.

```
lab@MX# show policy-options prefix-list bgp-peers | display inheritance
##
## apply-path was expanded to:
##     192.168.1.2/32;
##     172.16.1.1/32;
##
apply-path "protocols bgp group <*> neighbor <*>";
```

Our dynamic prefix-list is working!

You can apply the prefix-list “bgp-peers” just like any other prefix-list in a firewall filter term:

```
lab@MX# show firewall family inet
filter protect-re {
    term allow-bgp {
        from {
            source-prefix-list {
                bgp-peers;
            }
            protocol tcp;
            port bgp;
        }
        then accept;
    }
}
```

That’s it. In the above example we used the “apply-path” feature for adding BGP peers to our source-prefix-list.

© 2020 Juniper Networks, Inc. All rights reserved.

Another great Junos feature is "`apply-flags omit`". With this feature it is possible to remove extensive configuration listings from the "show configuration" command. You can apply "apply-flags omit" almost everywhere in the Junos configuration hierarchy.

In the following example, we demonstrate the usage of the "apply-flags omit" feature for firewall filters. Imagine a very long firewall filter (in our case it contains just one term). For day-to-day operation, you do not want to be bothered with endless pages of firewall filters.

Configure the "`apply-flags omit`" statement for the re-protect firewall filter.

```
lab@MX# set firewall family inet filter re-protect apply-flags omit
lab@MX# commit
```

Verify our re-protect filter.

```
lab@MX# show firewall family inet
filter re-protect { /* OMITTED */ };

lab@MX#
```

As you can see the details of our firewall filter "`re-protect`" are now omitted from our configuration.

There are two ways to show the firewall filter details. You can use the "`display omit`" or "`display set`" statements when showing the configuration.

```
lab@MX# show firewall family inet | display omit
filter re-protect {
    apply-flags omit;
    term allow-bgp {
        from {
            source-prefix-list {
                bgp-peers;
            }
            protocol tcp;
            port bgp;
        }
        then accept;
    }
}

or

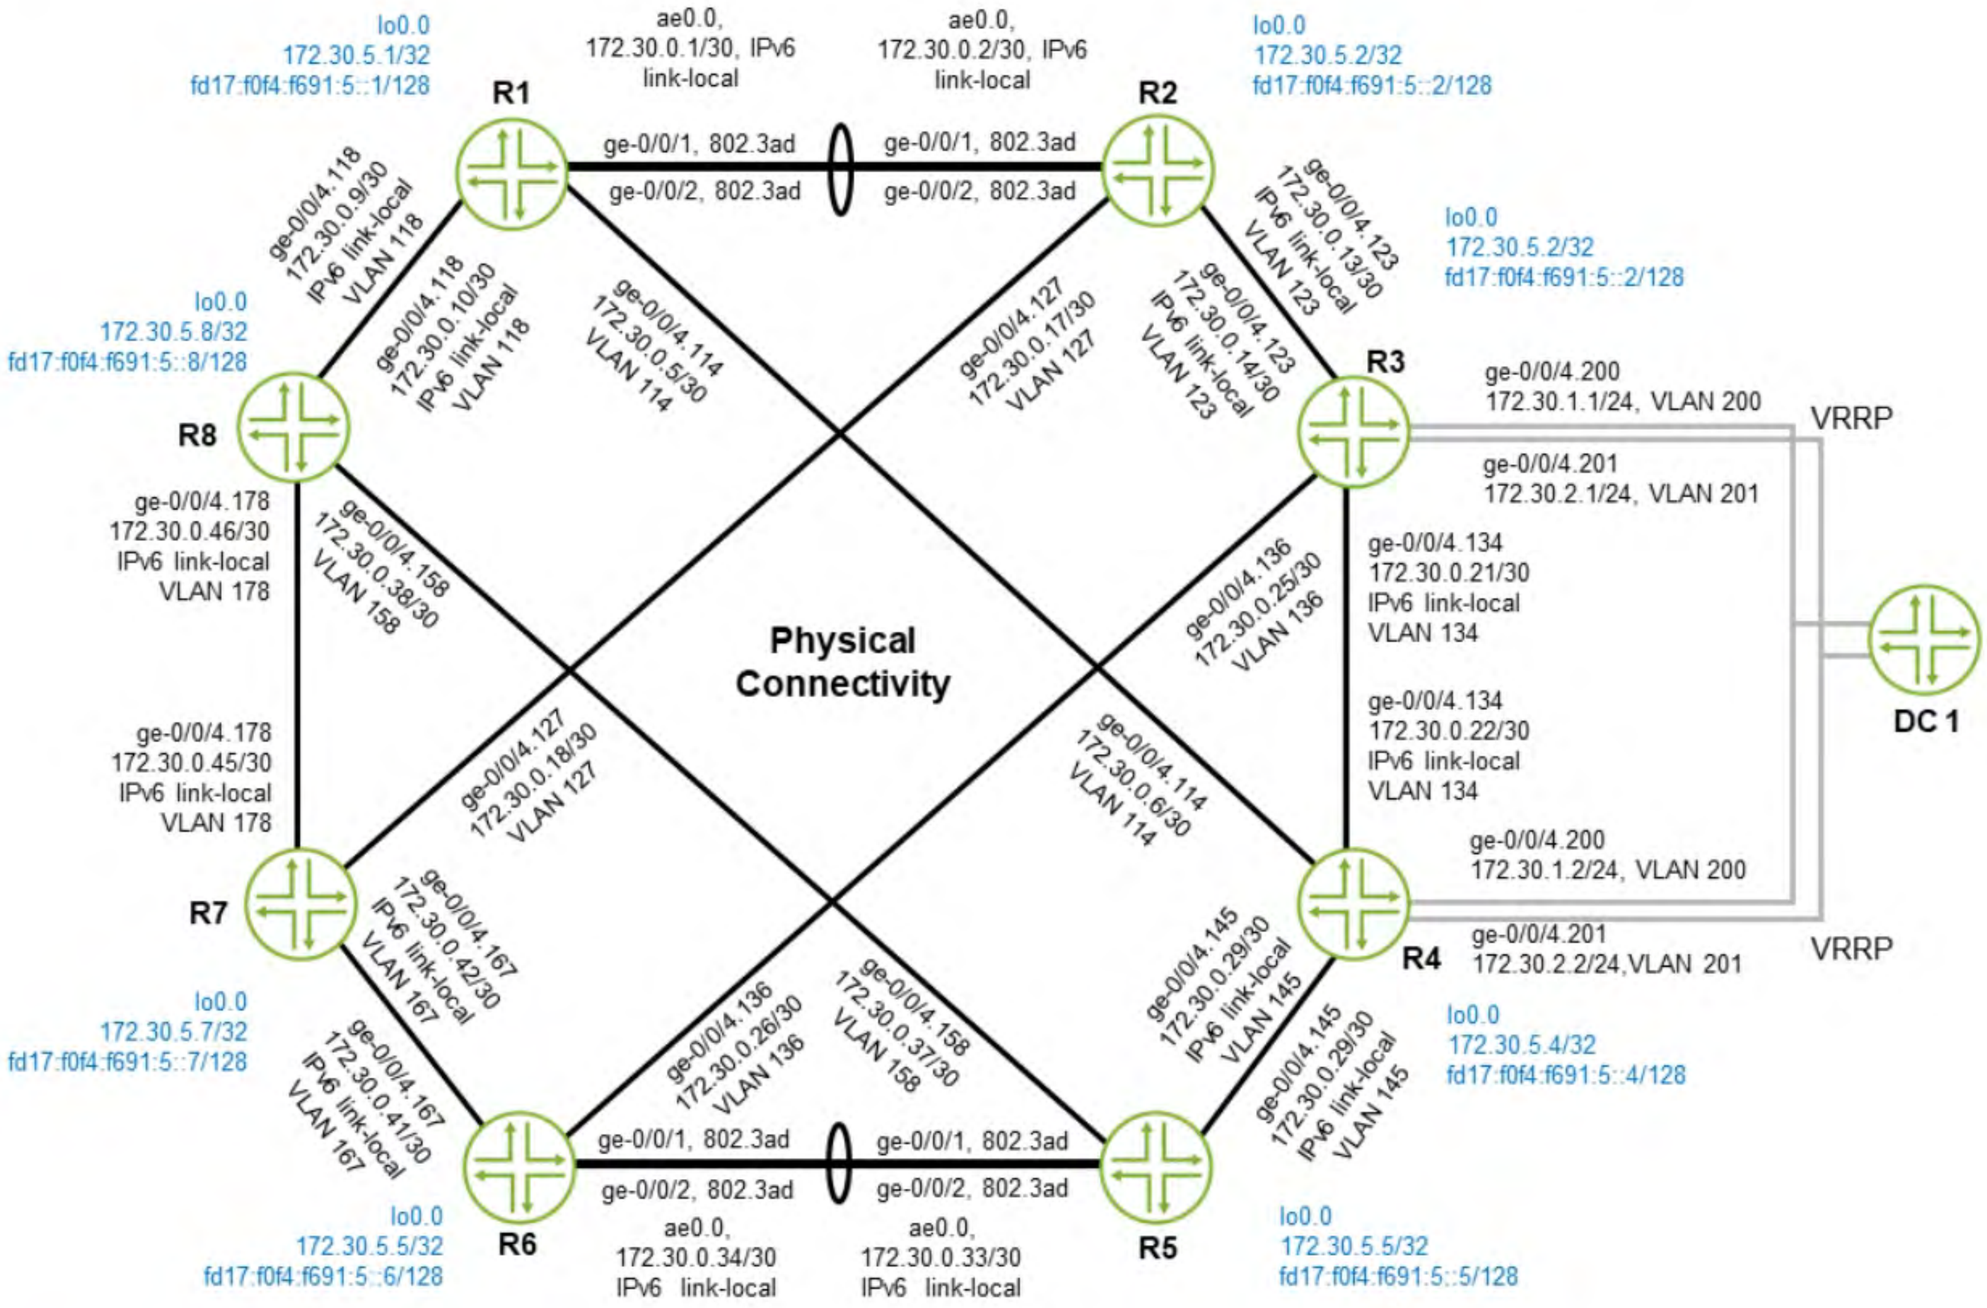
lab@MX# show firewall family inet | display set
set firewall family inet filter re-protect apply-flags omit
set firewall family inet filter re-protect term allow-bgp from source-prefix-list
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
bgp-peers
set firewall family inet filter re-protect term allow-bgp from protocol tcp
set firewall family inet filter re-protect term allow-bgp from port bgp
set firewall family inet filter re-protect term allow-bgp then accept
```

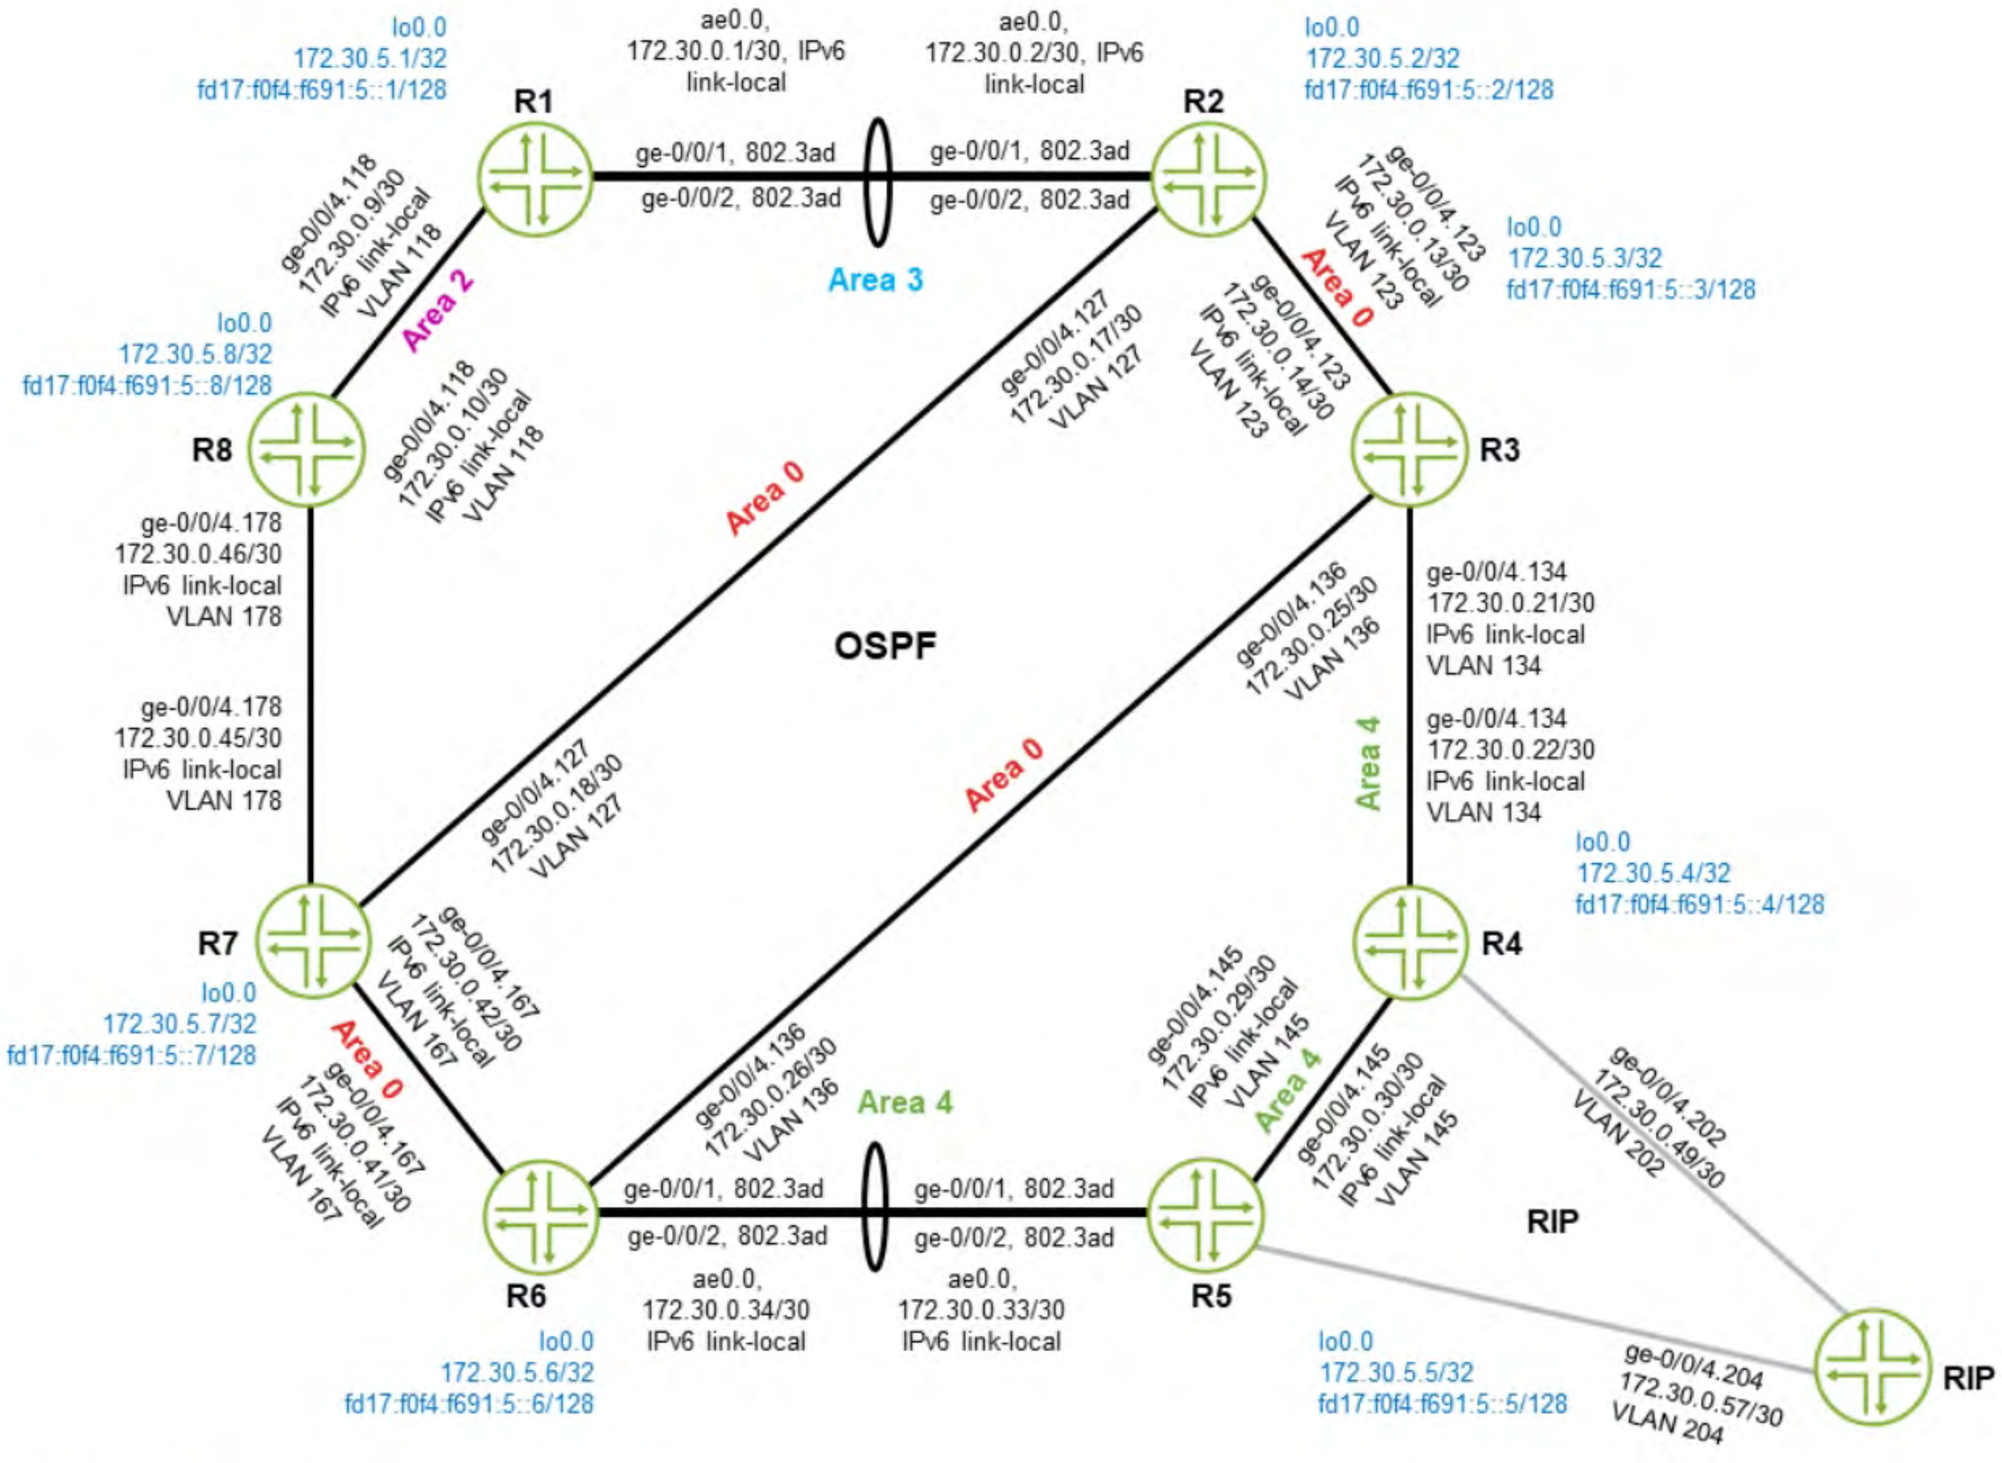
Appendix - Topology Diagrams

In this appendix you will find the chapters topology diagrams in full size format.

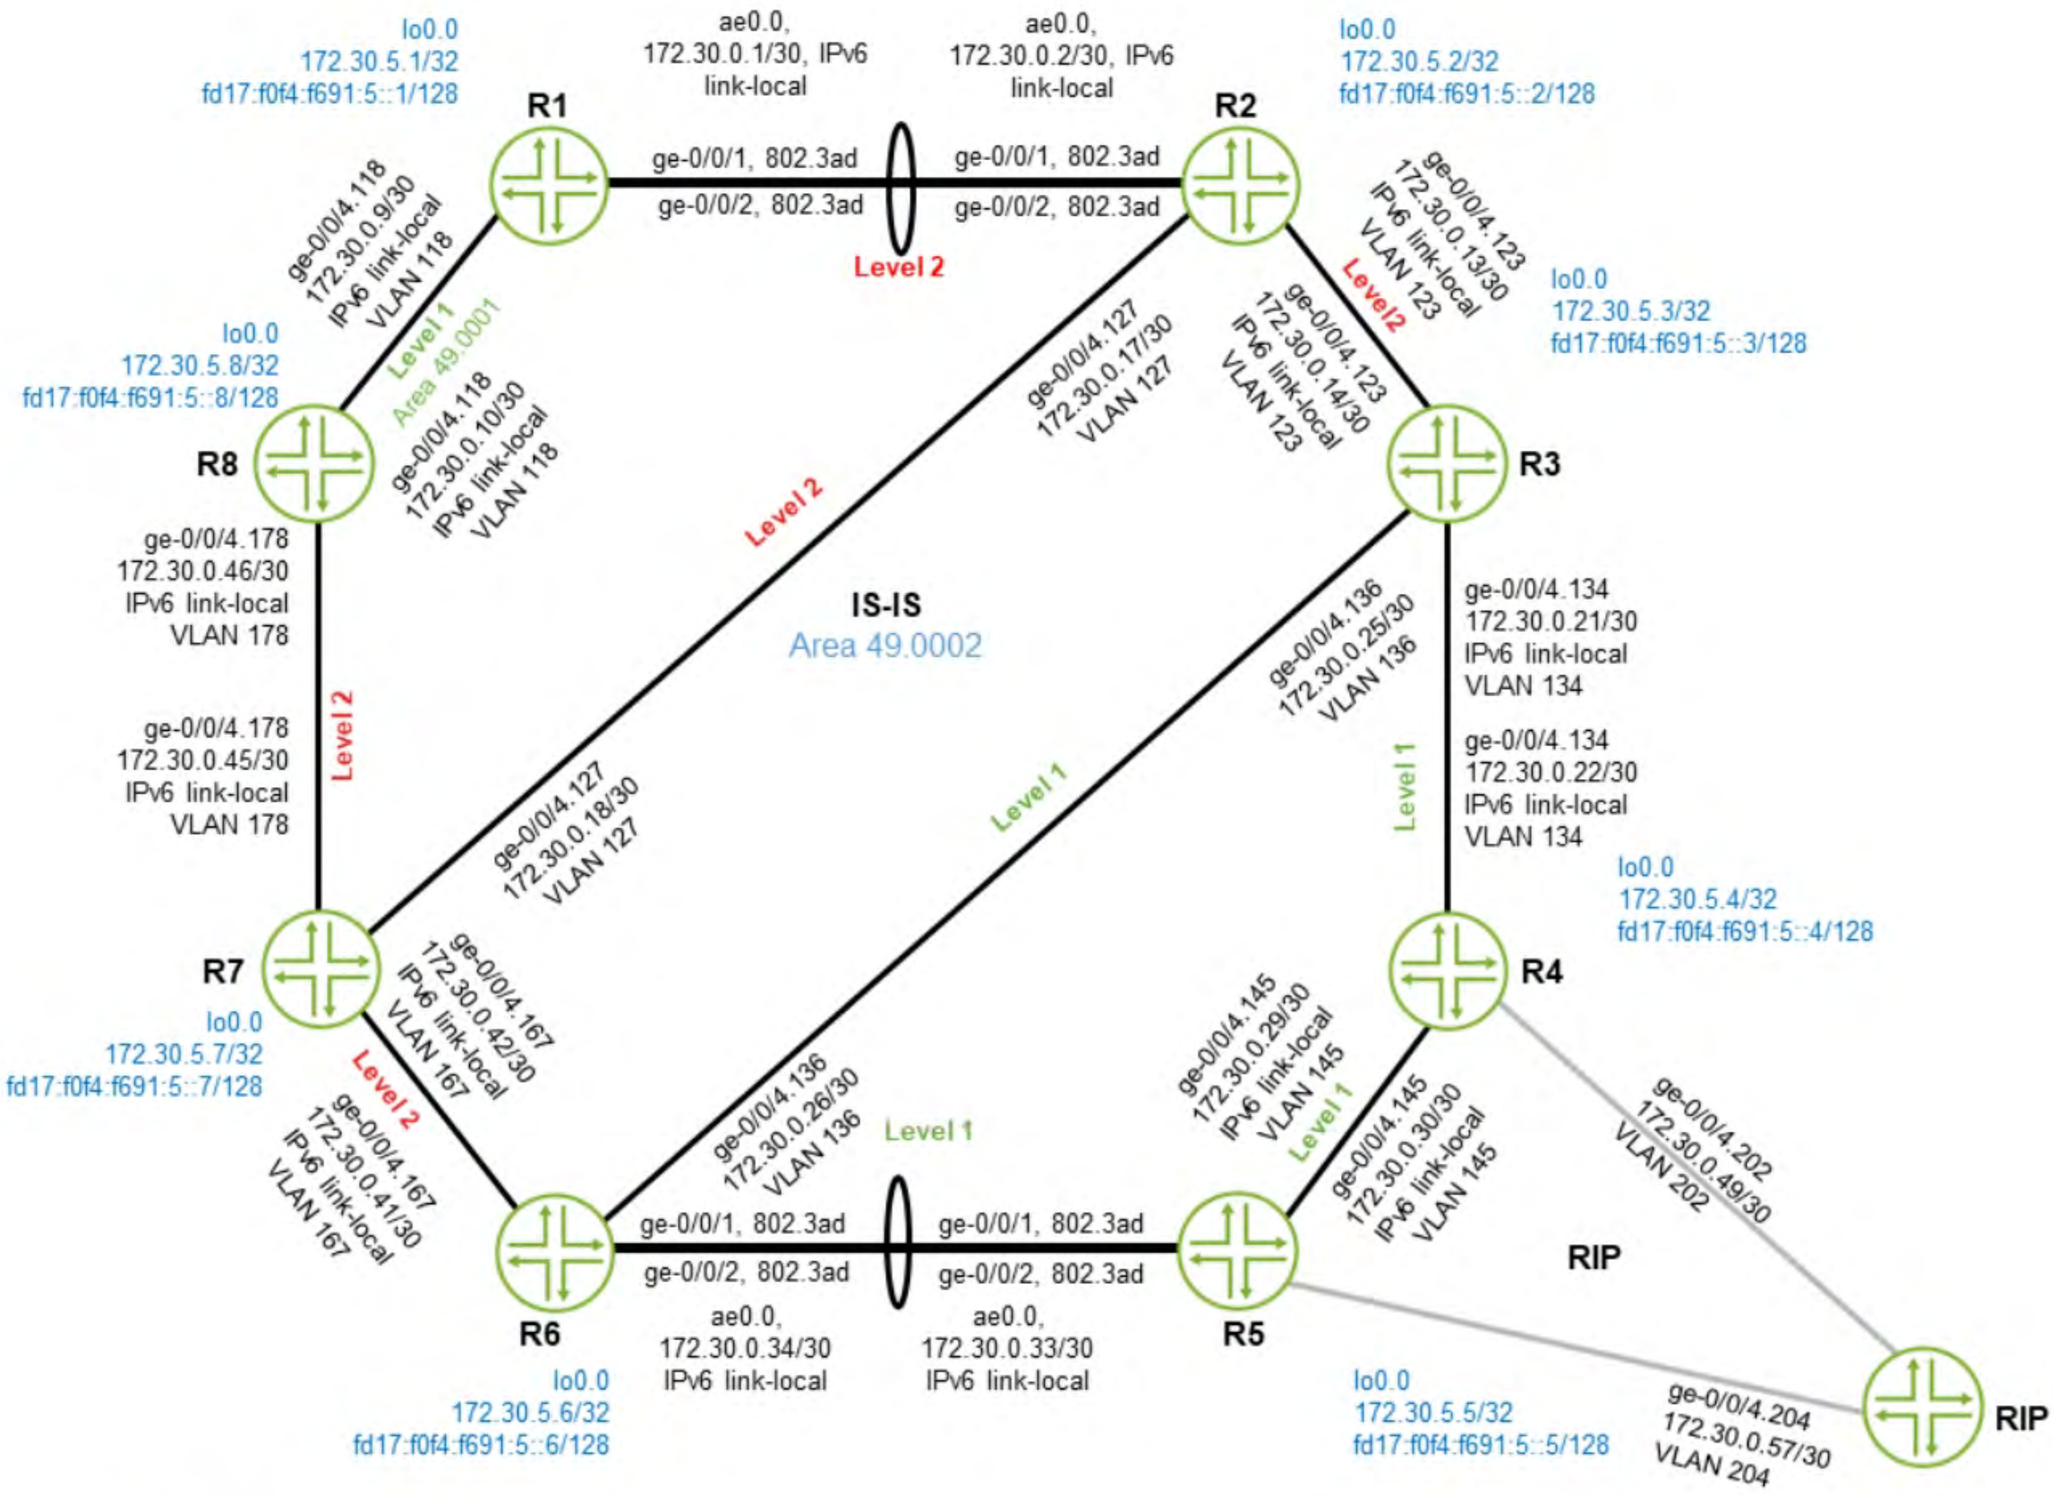


Chapter 1 - task 4

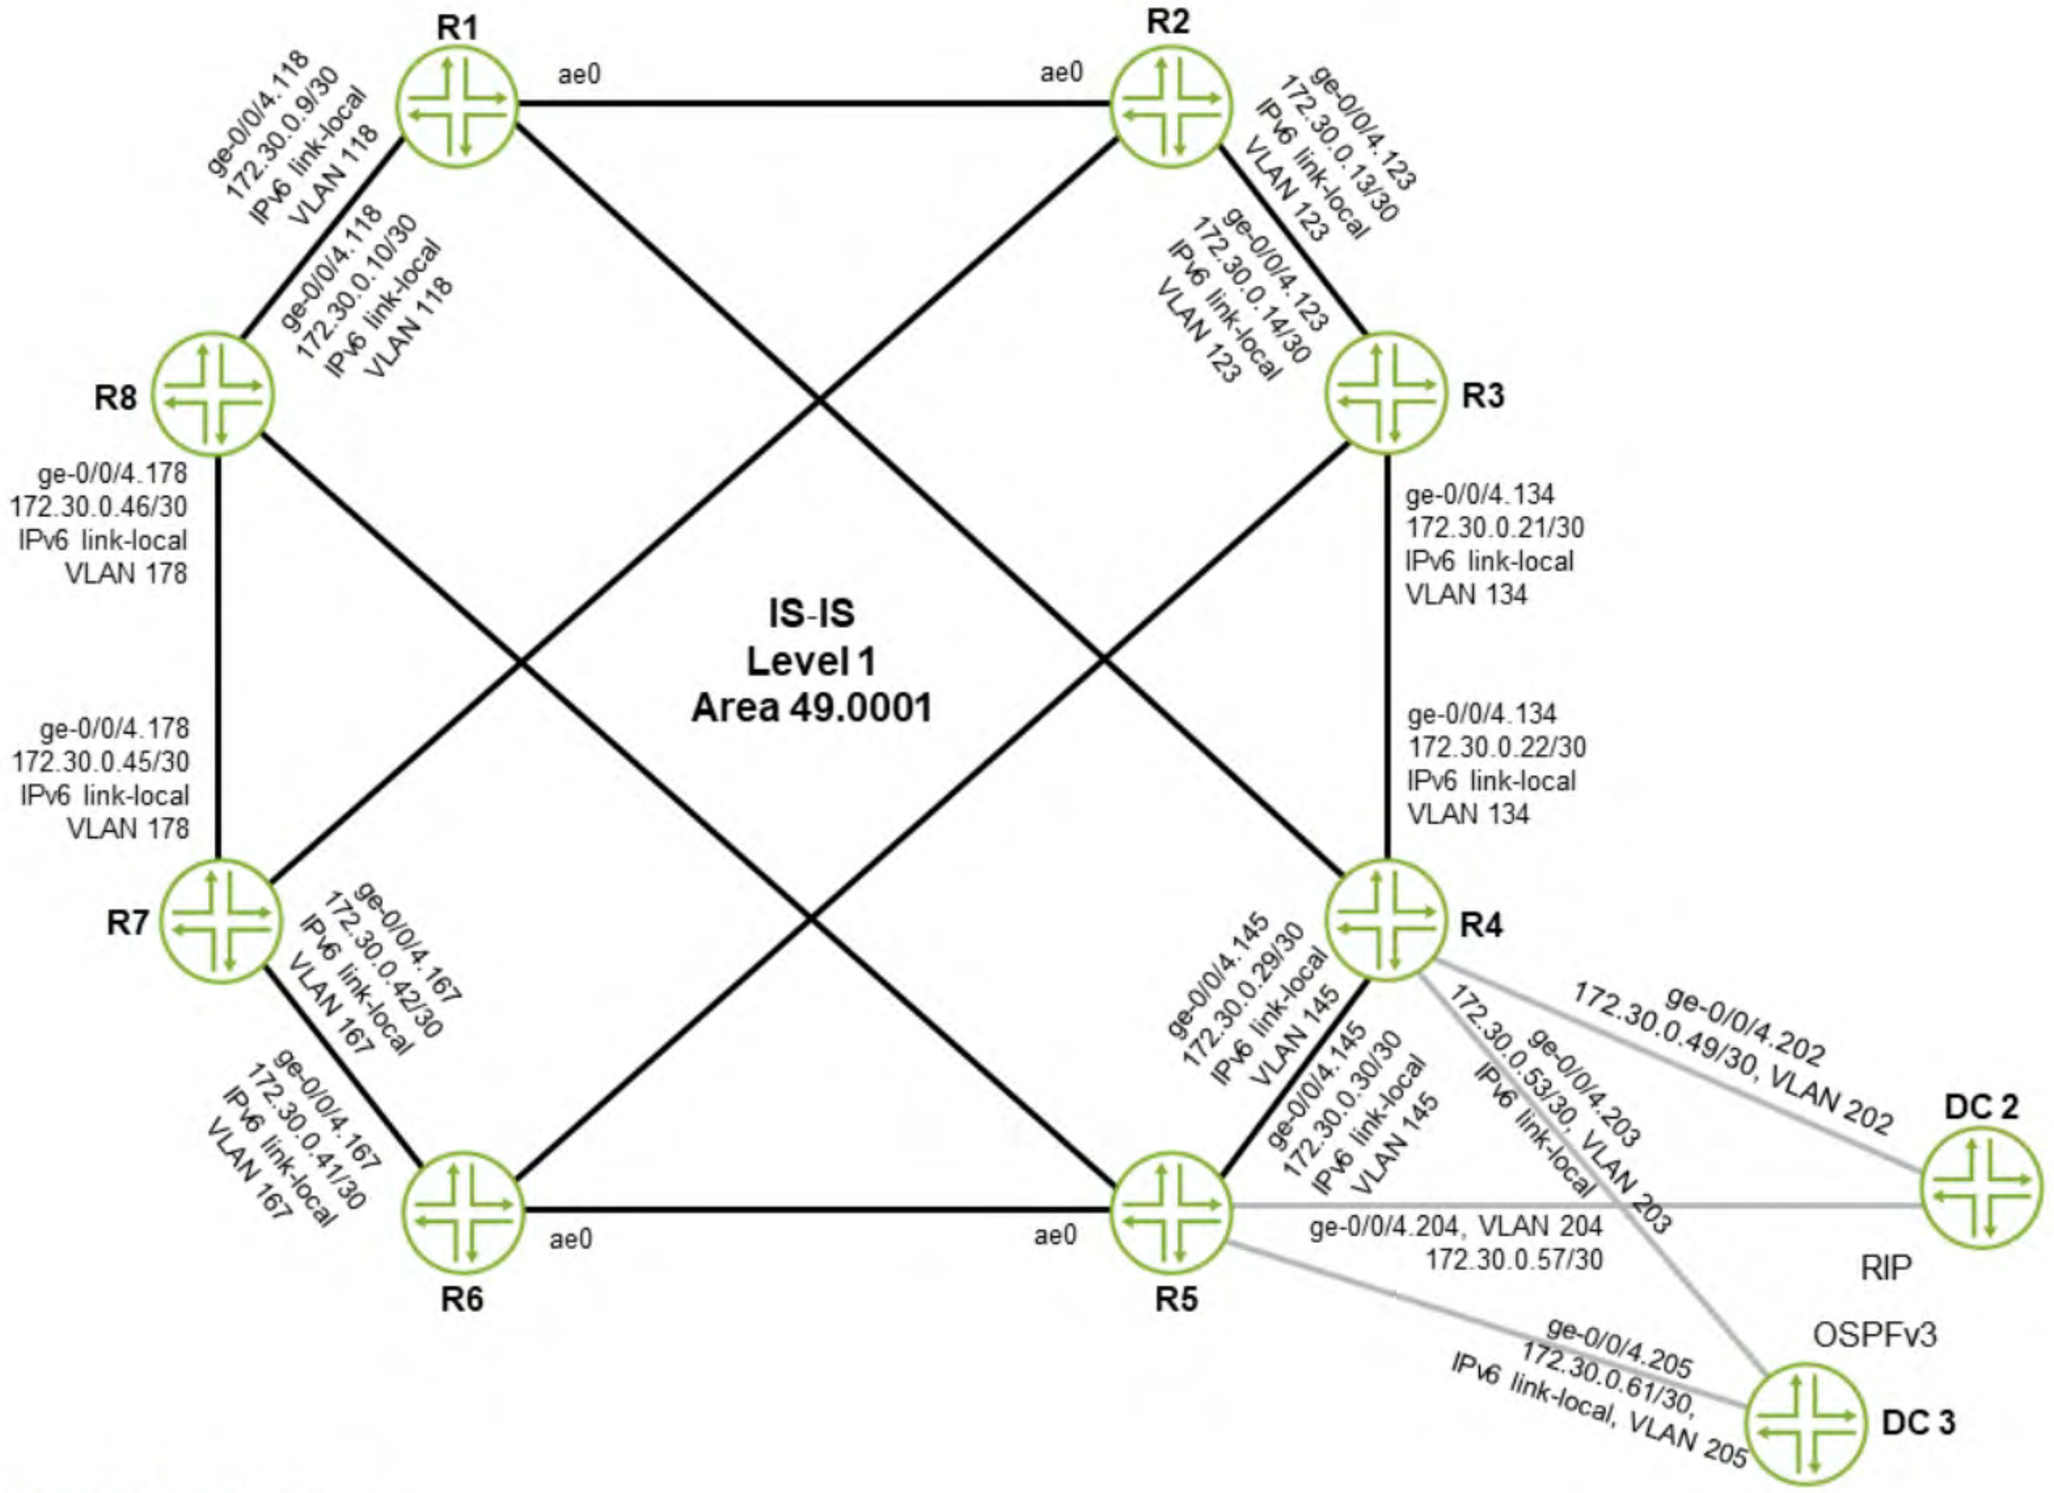
© 2020 Juniper Networks, Inc. All rights reserved.



© 2020 Juniper Networks, Inc. All rights reserved.

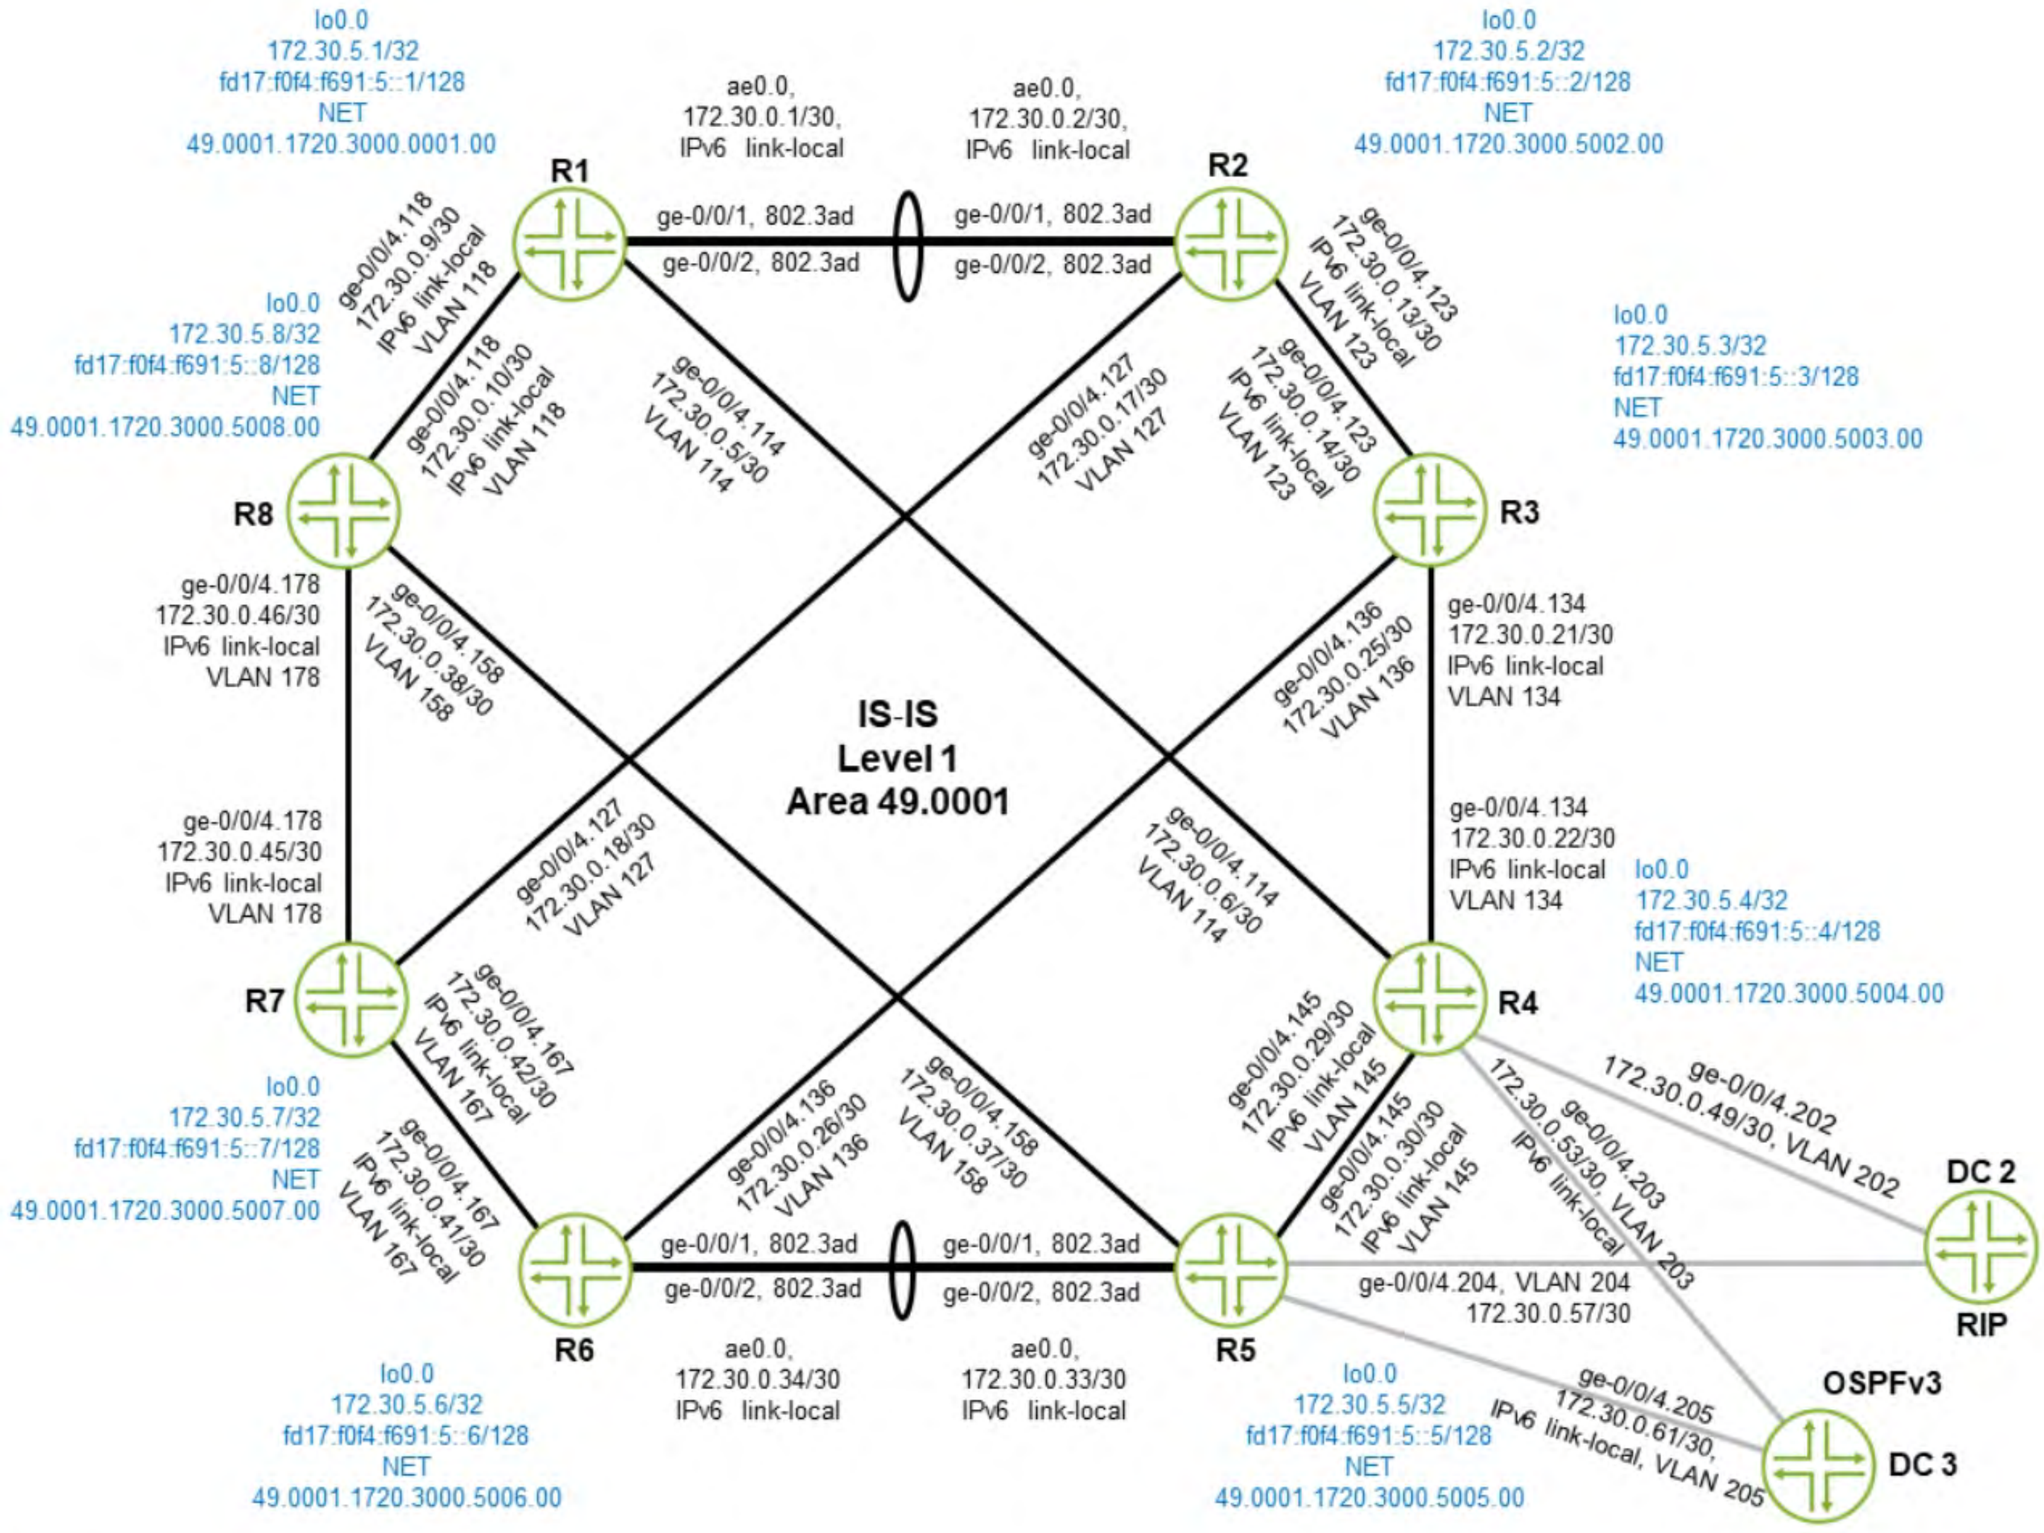


© 2020 Juniper Networks, Inc. All rights reserved.



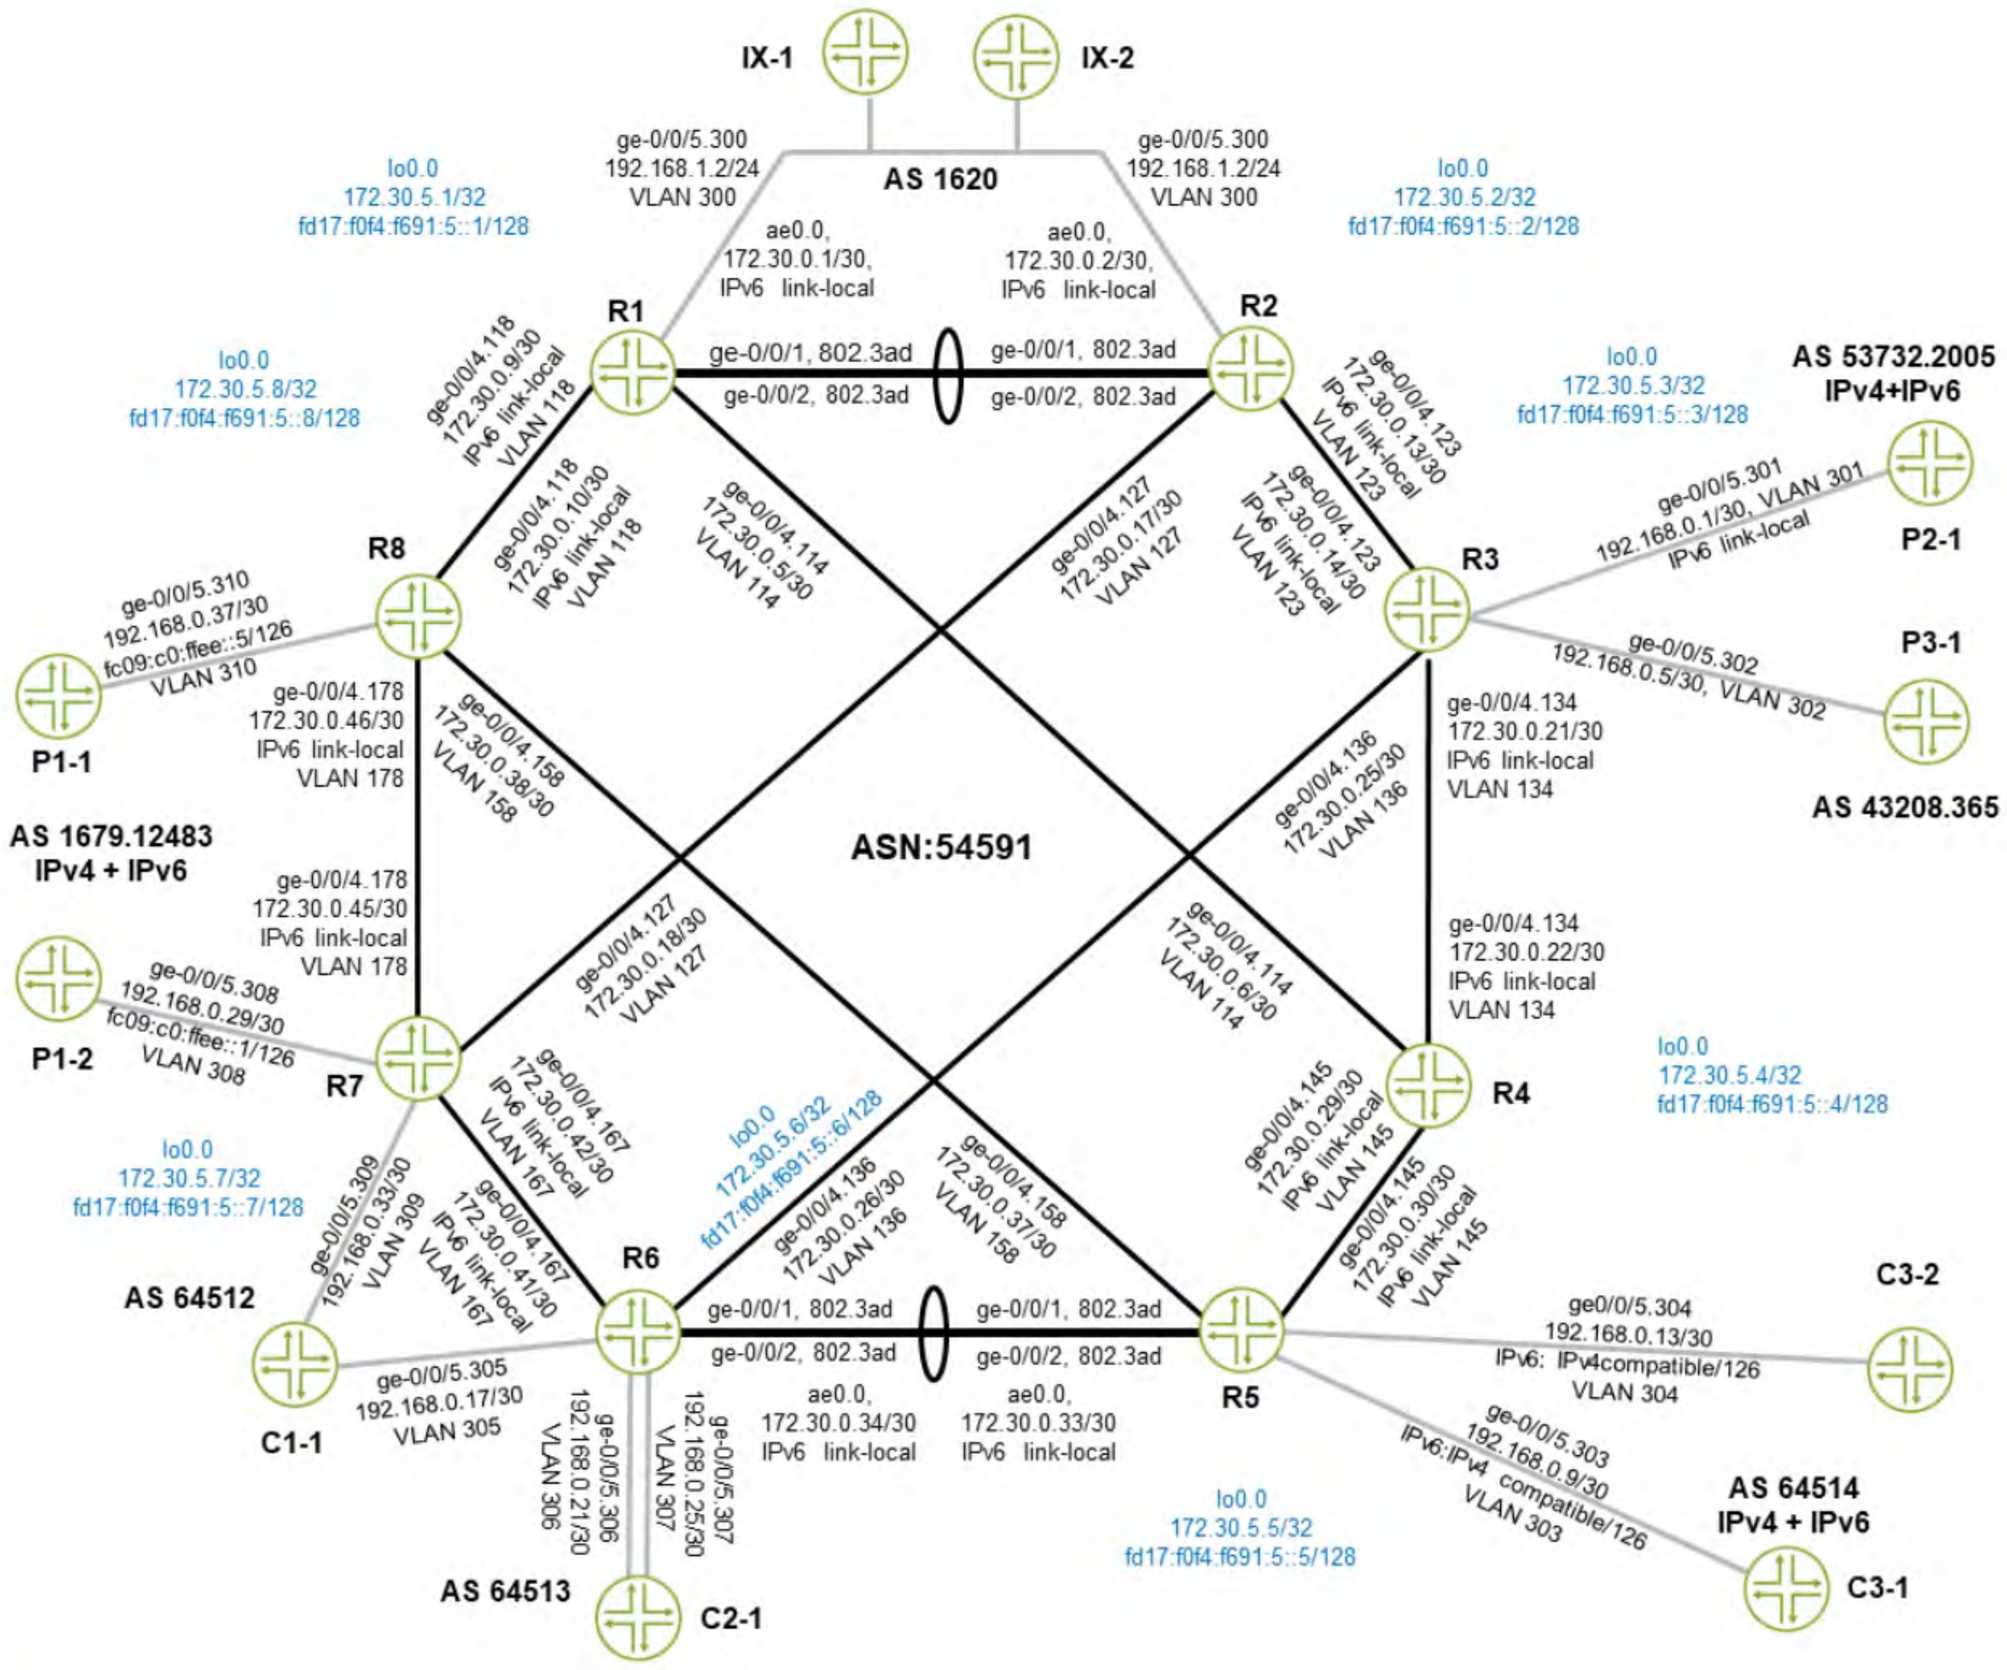
Chapter 2 - IGP rollout

© 2020 Juniper Networks, Inc. All rights reserved.

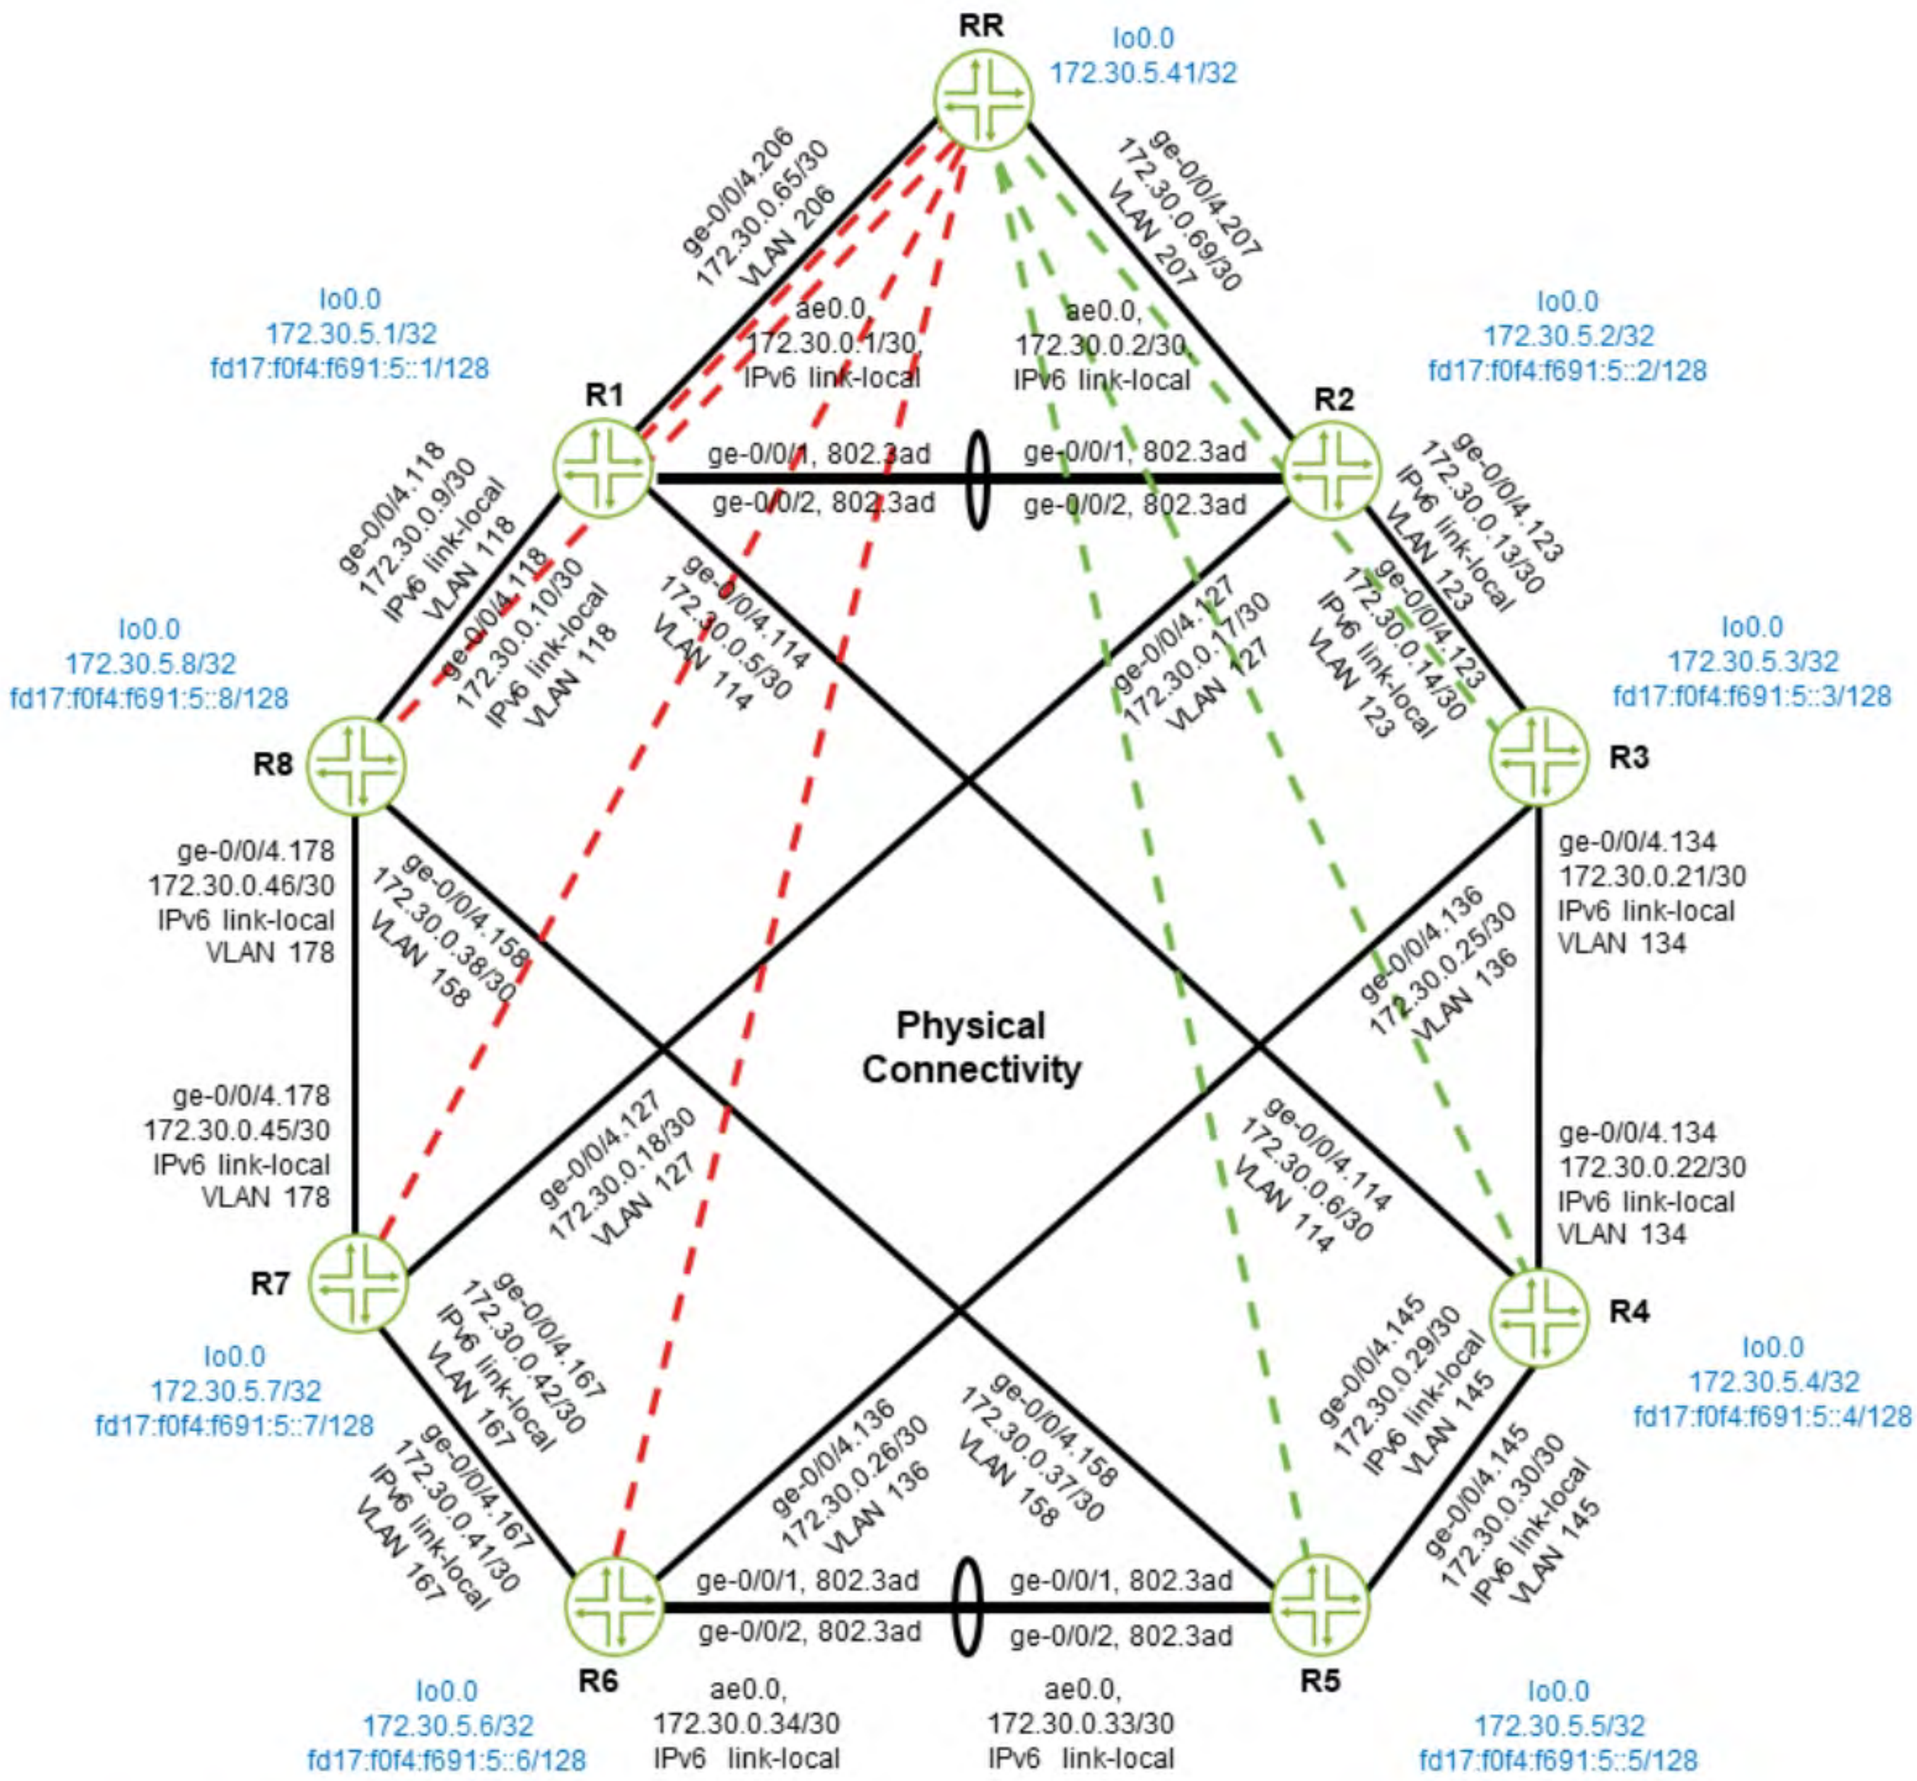


Chapter 2 - IGP rollout IS-IS

© 2020 Juniper Networks, Inc. All rights reserved.

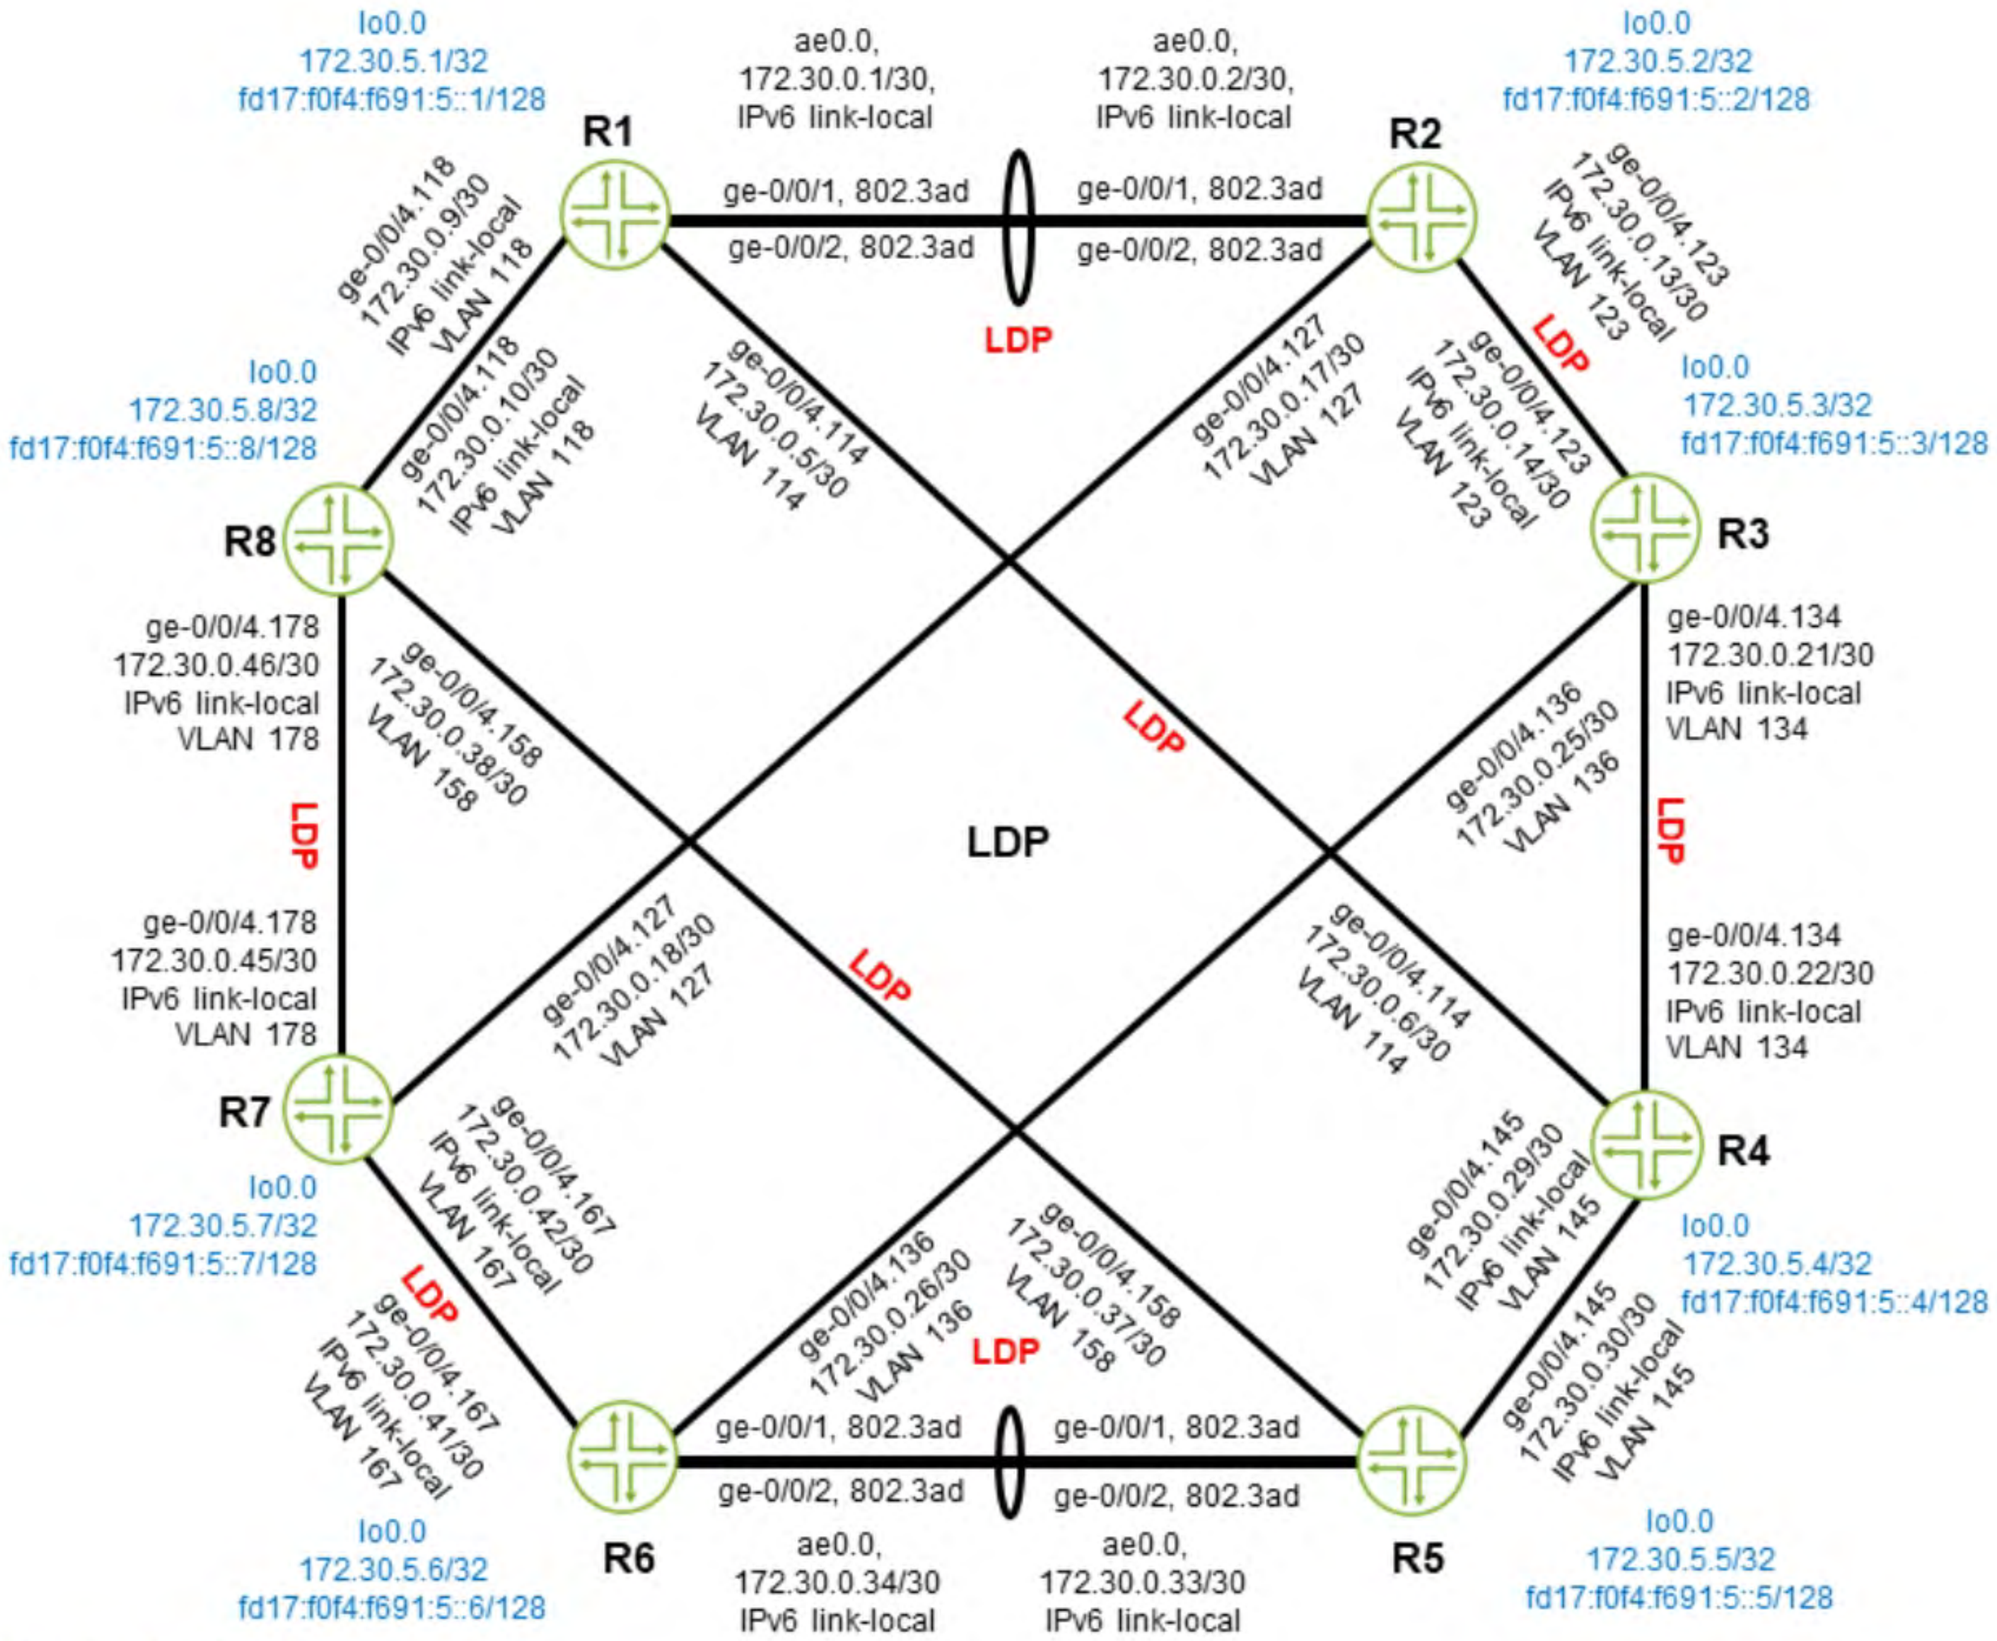


© 2020 Juniper Networks, Inc. All rights reserved.



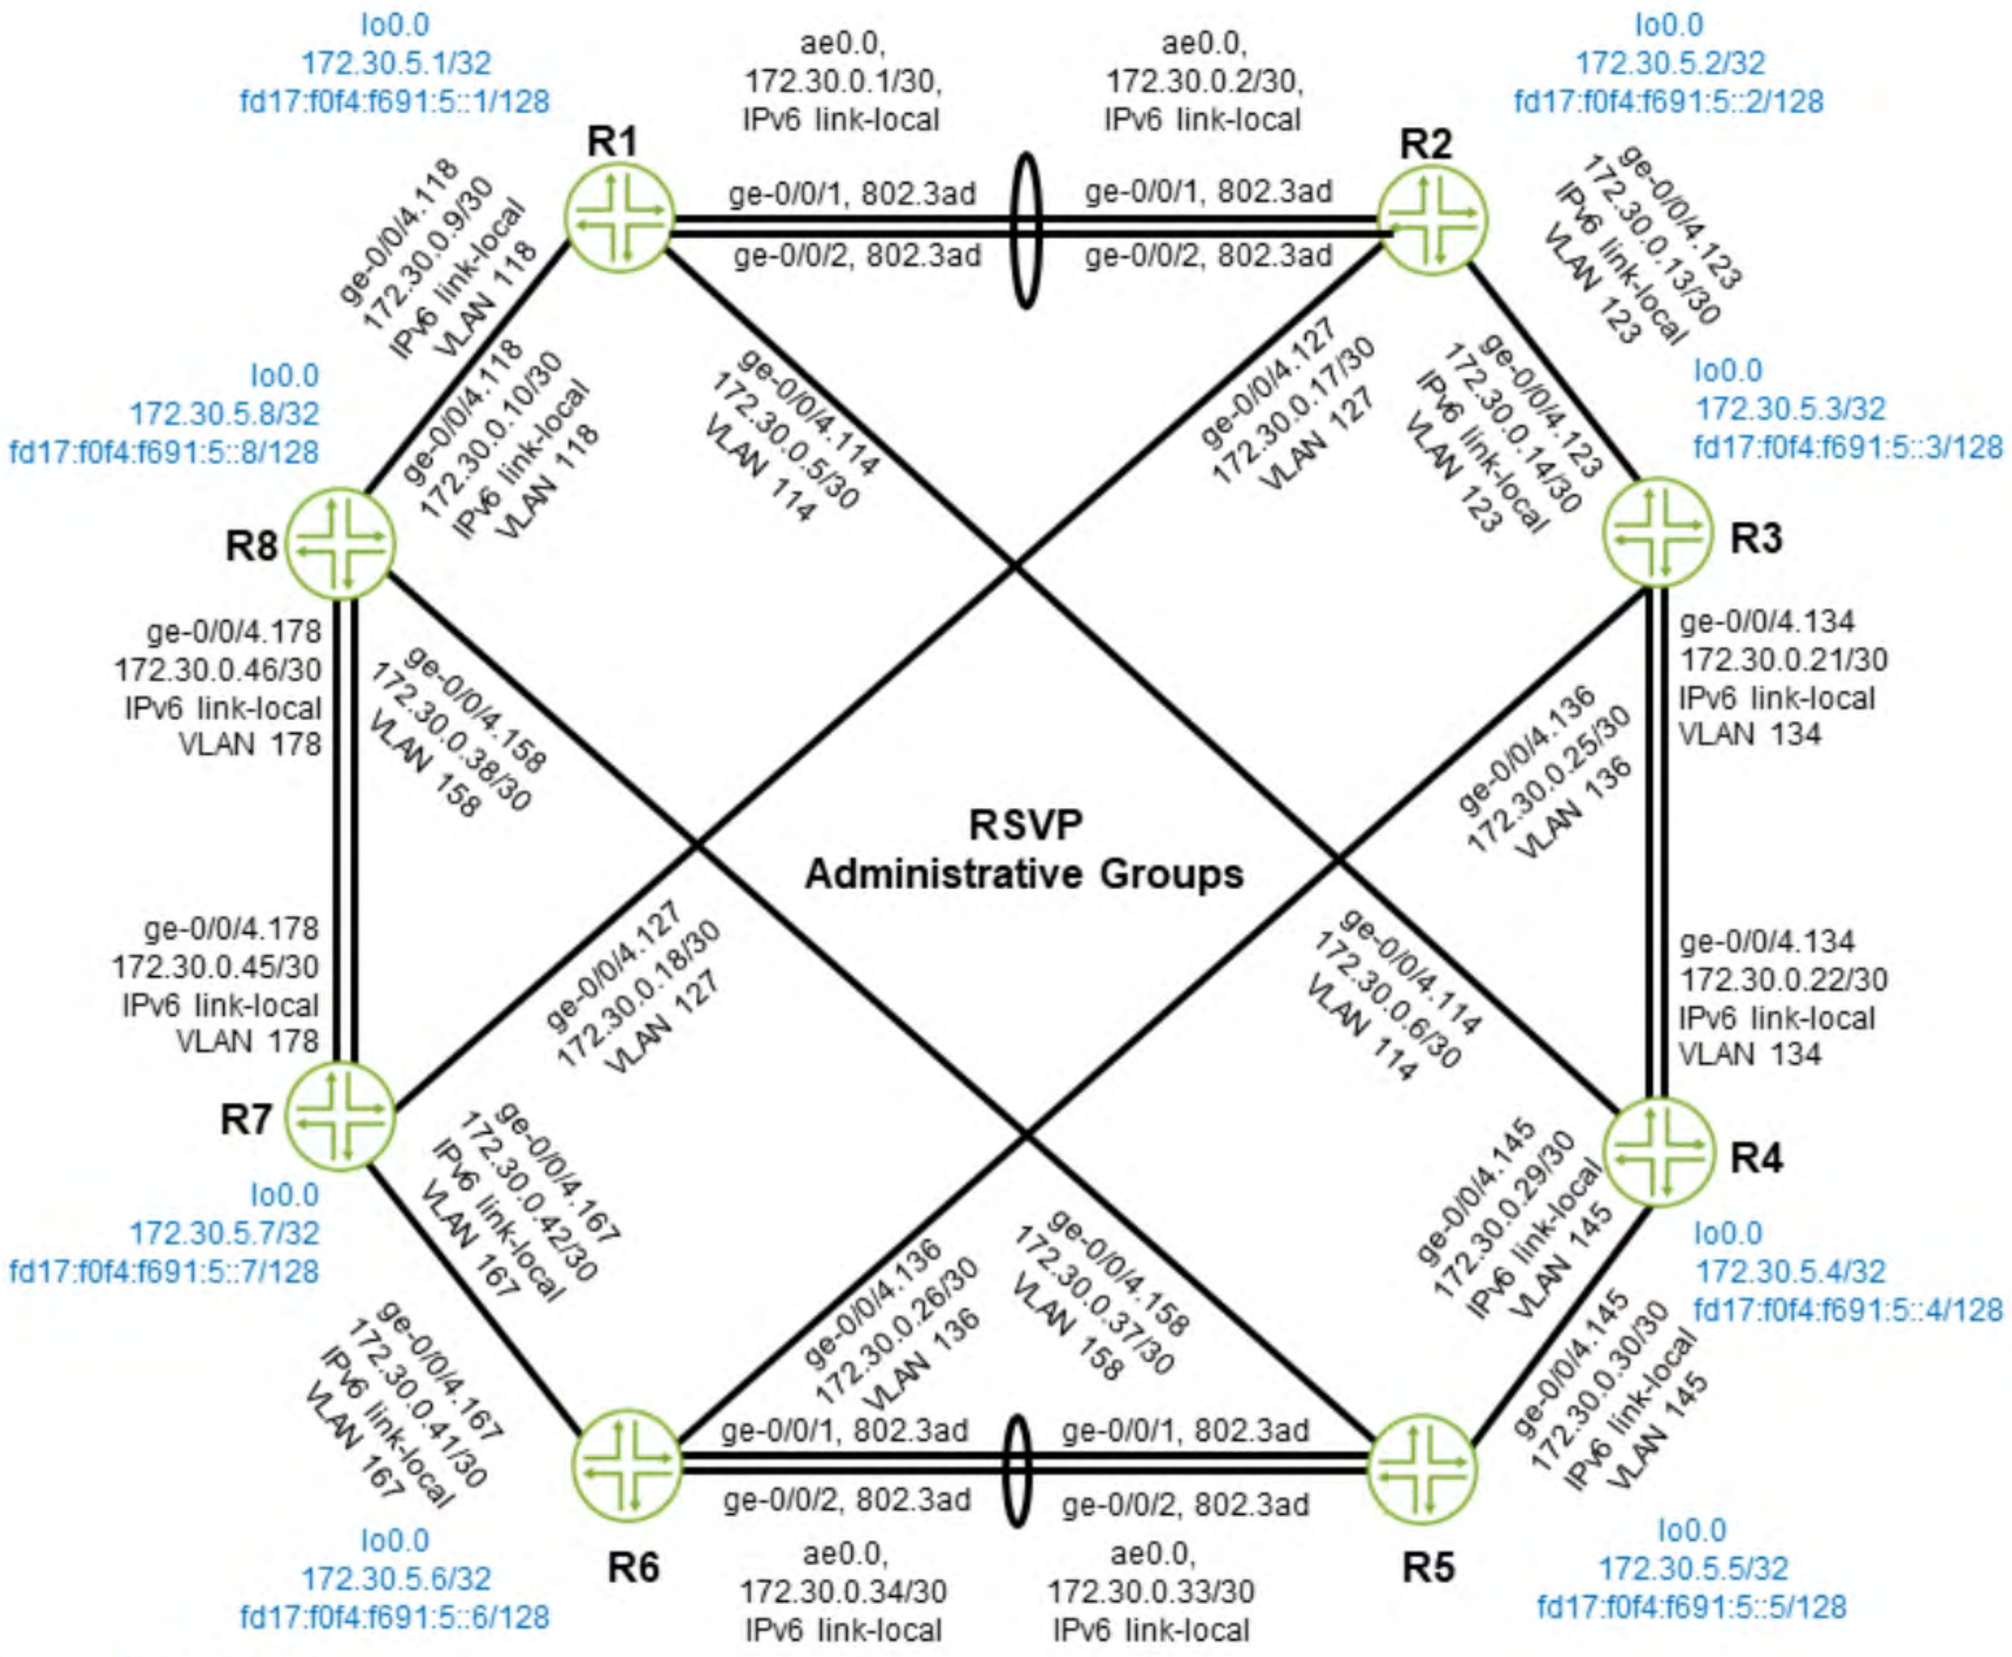
Chapter 3 - BGP 2

© 2020 Juniper Networks, Inc. All rights reserved.



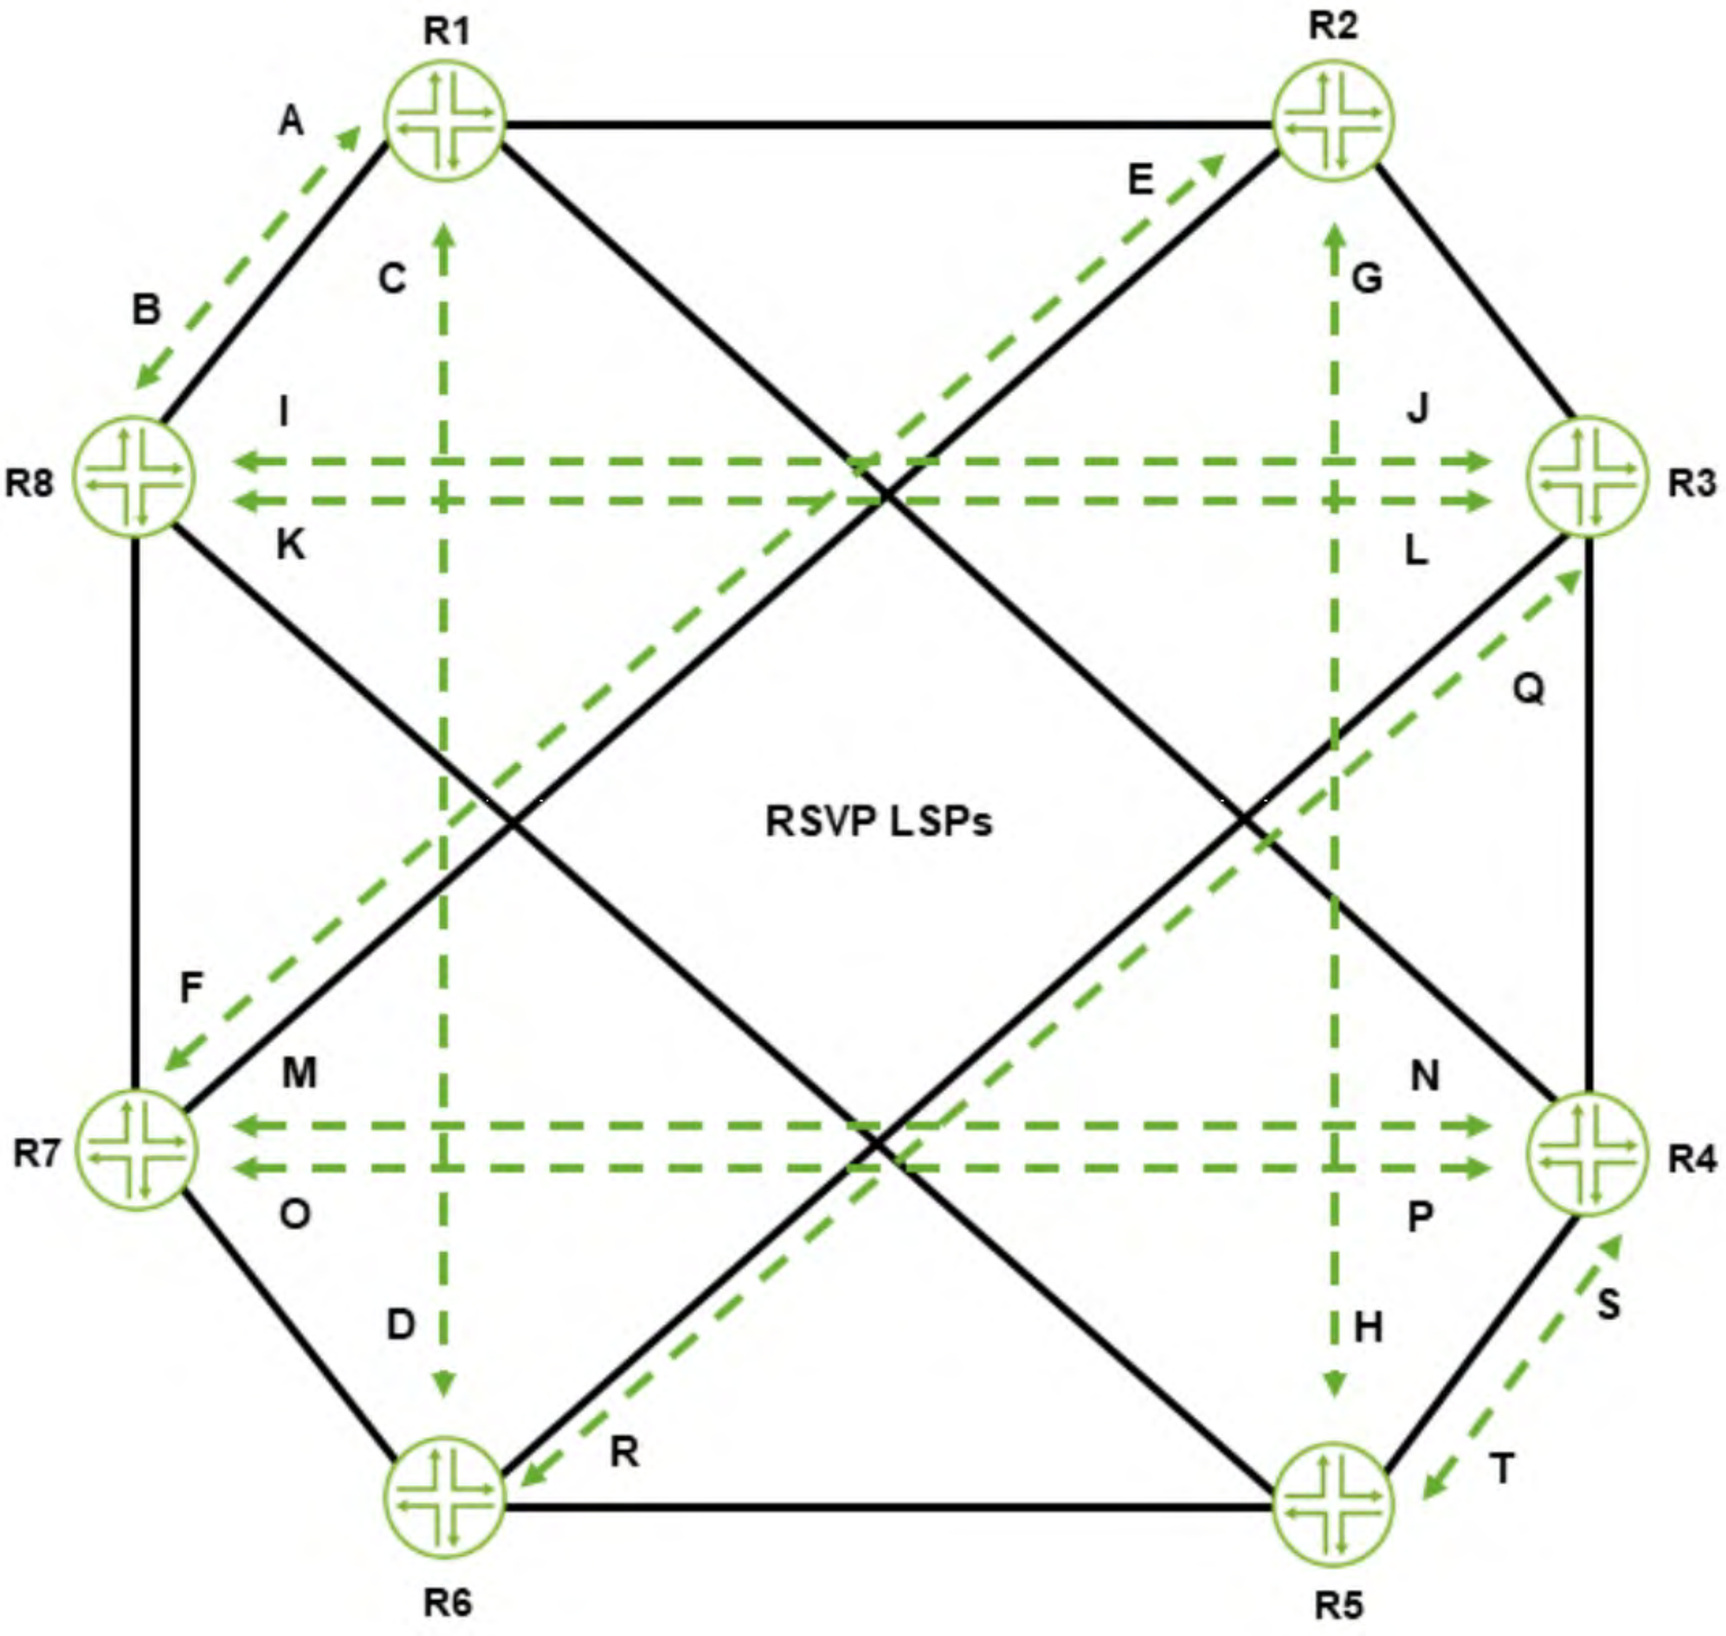
Chapter 4 - MPLS 1

© 2020 Juniper Networks, Inc. All rights reserved.

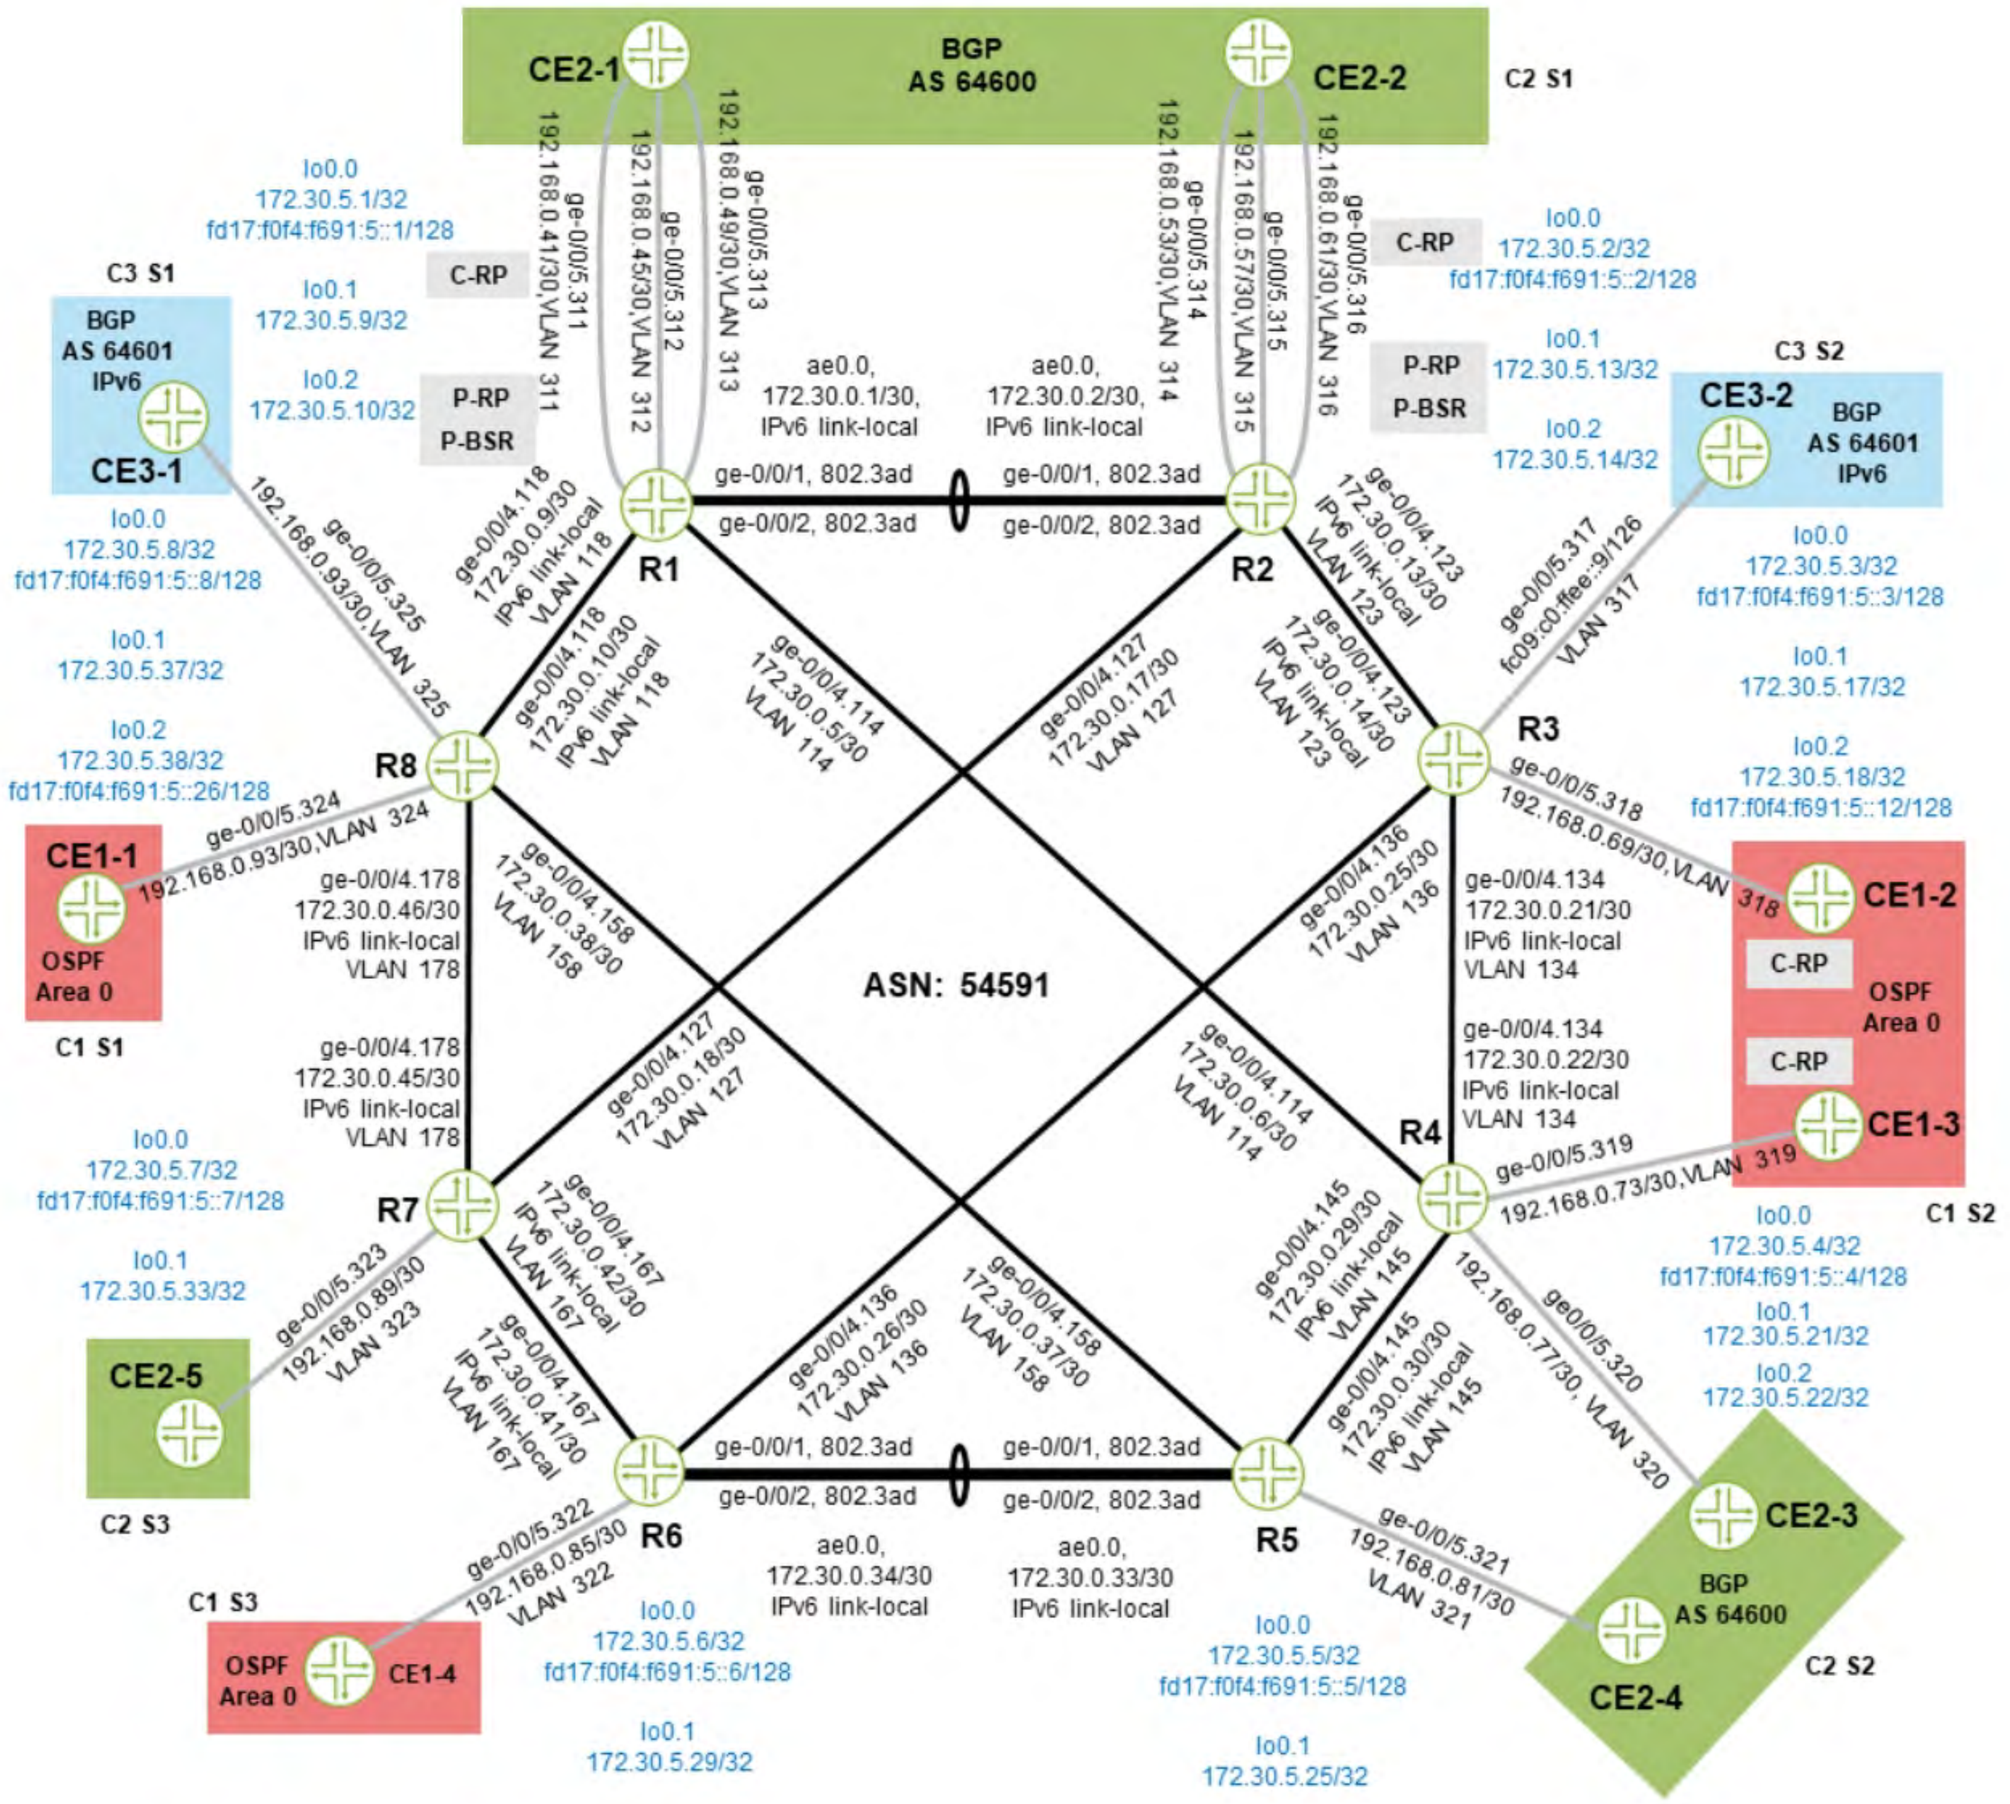


Chapter 4 - MPLS 2

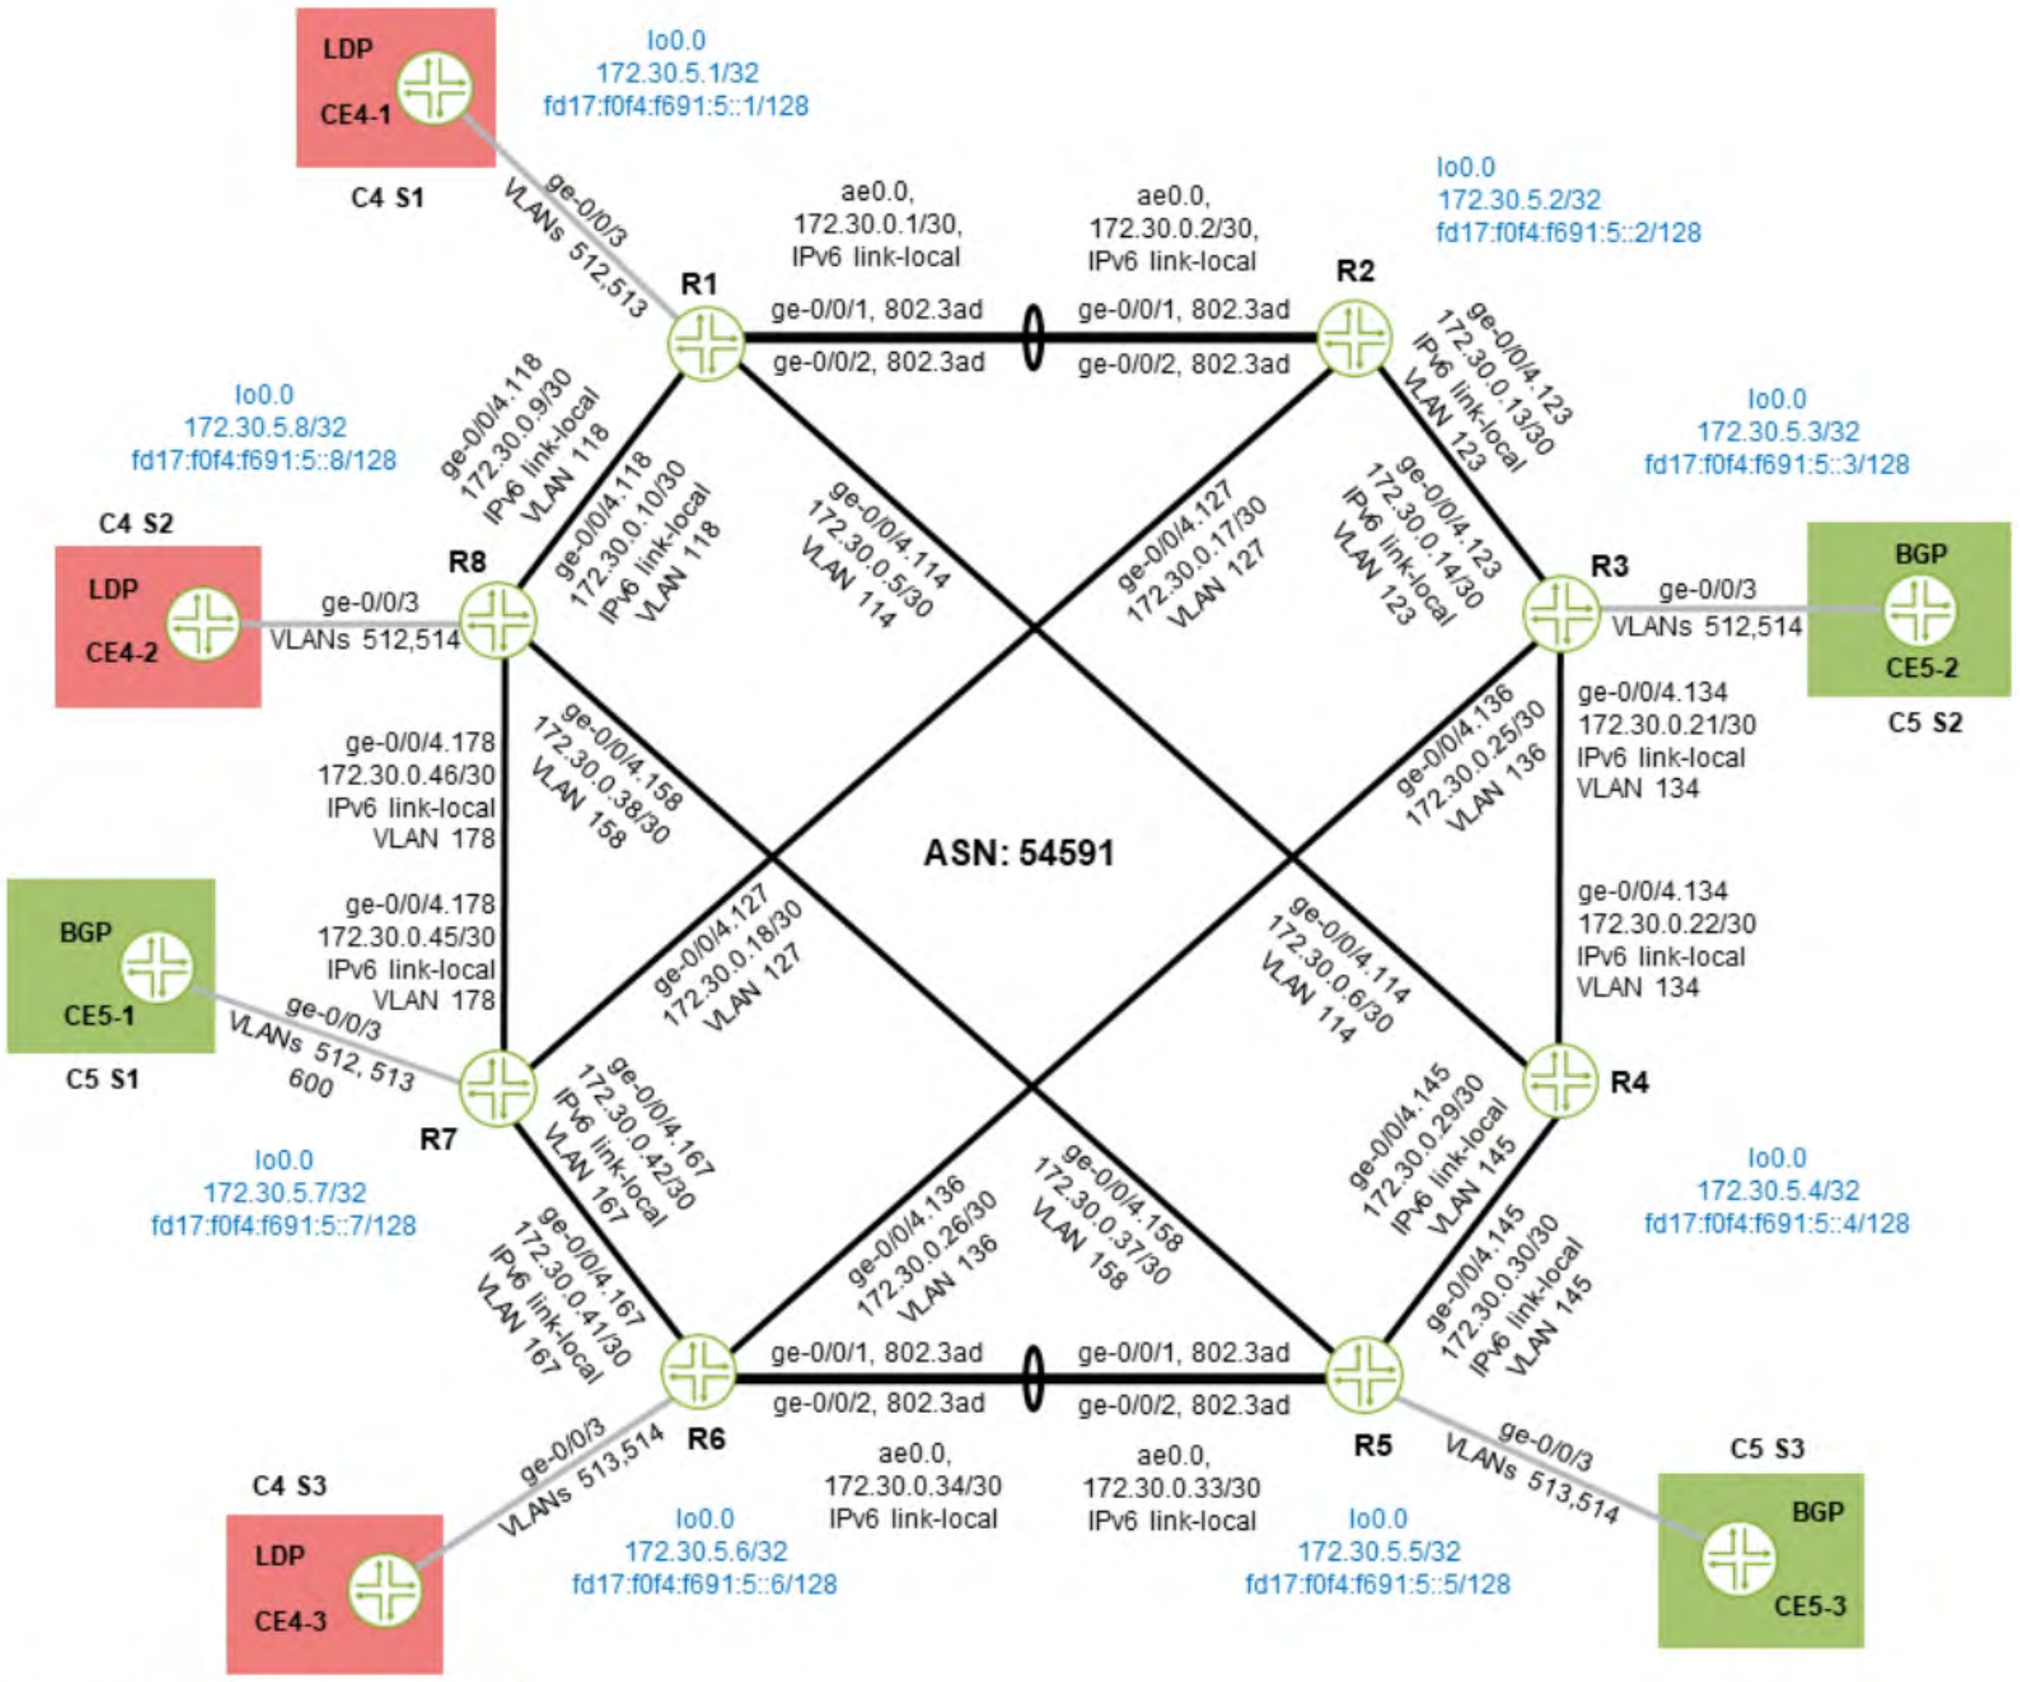
© 2020 Juniper Networks, Inc. All rights reserved.



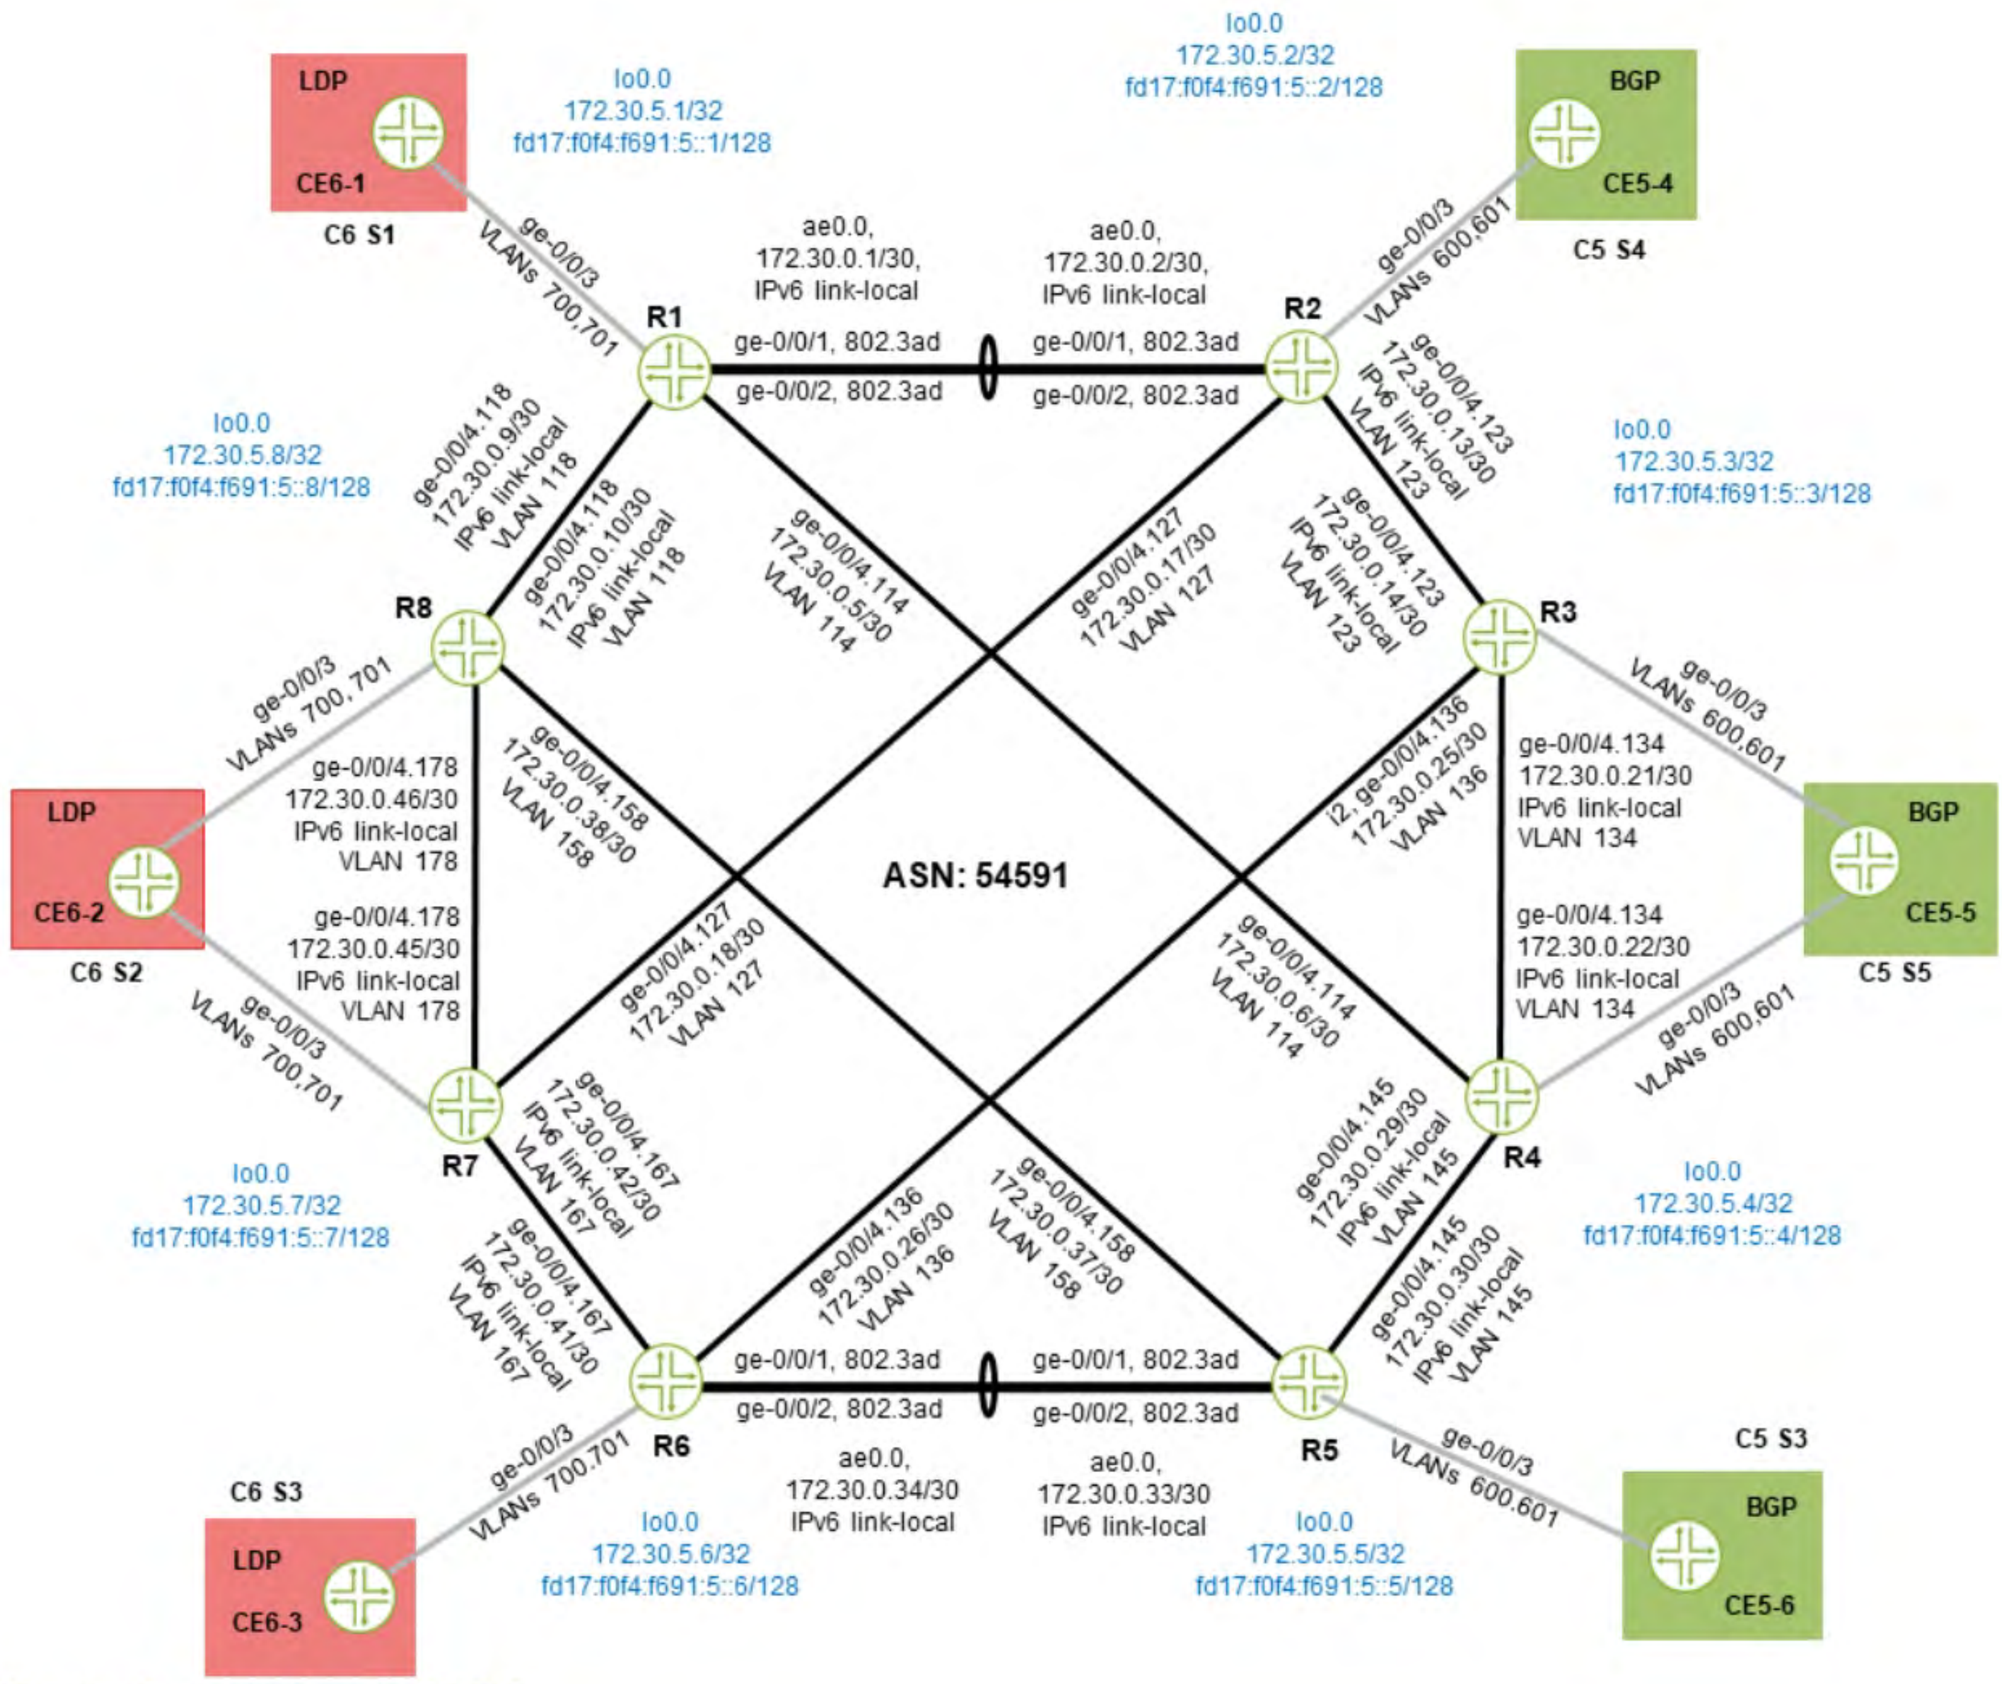
© 2020 Juniper Networks, Inc. All rights reserved.



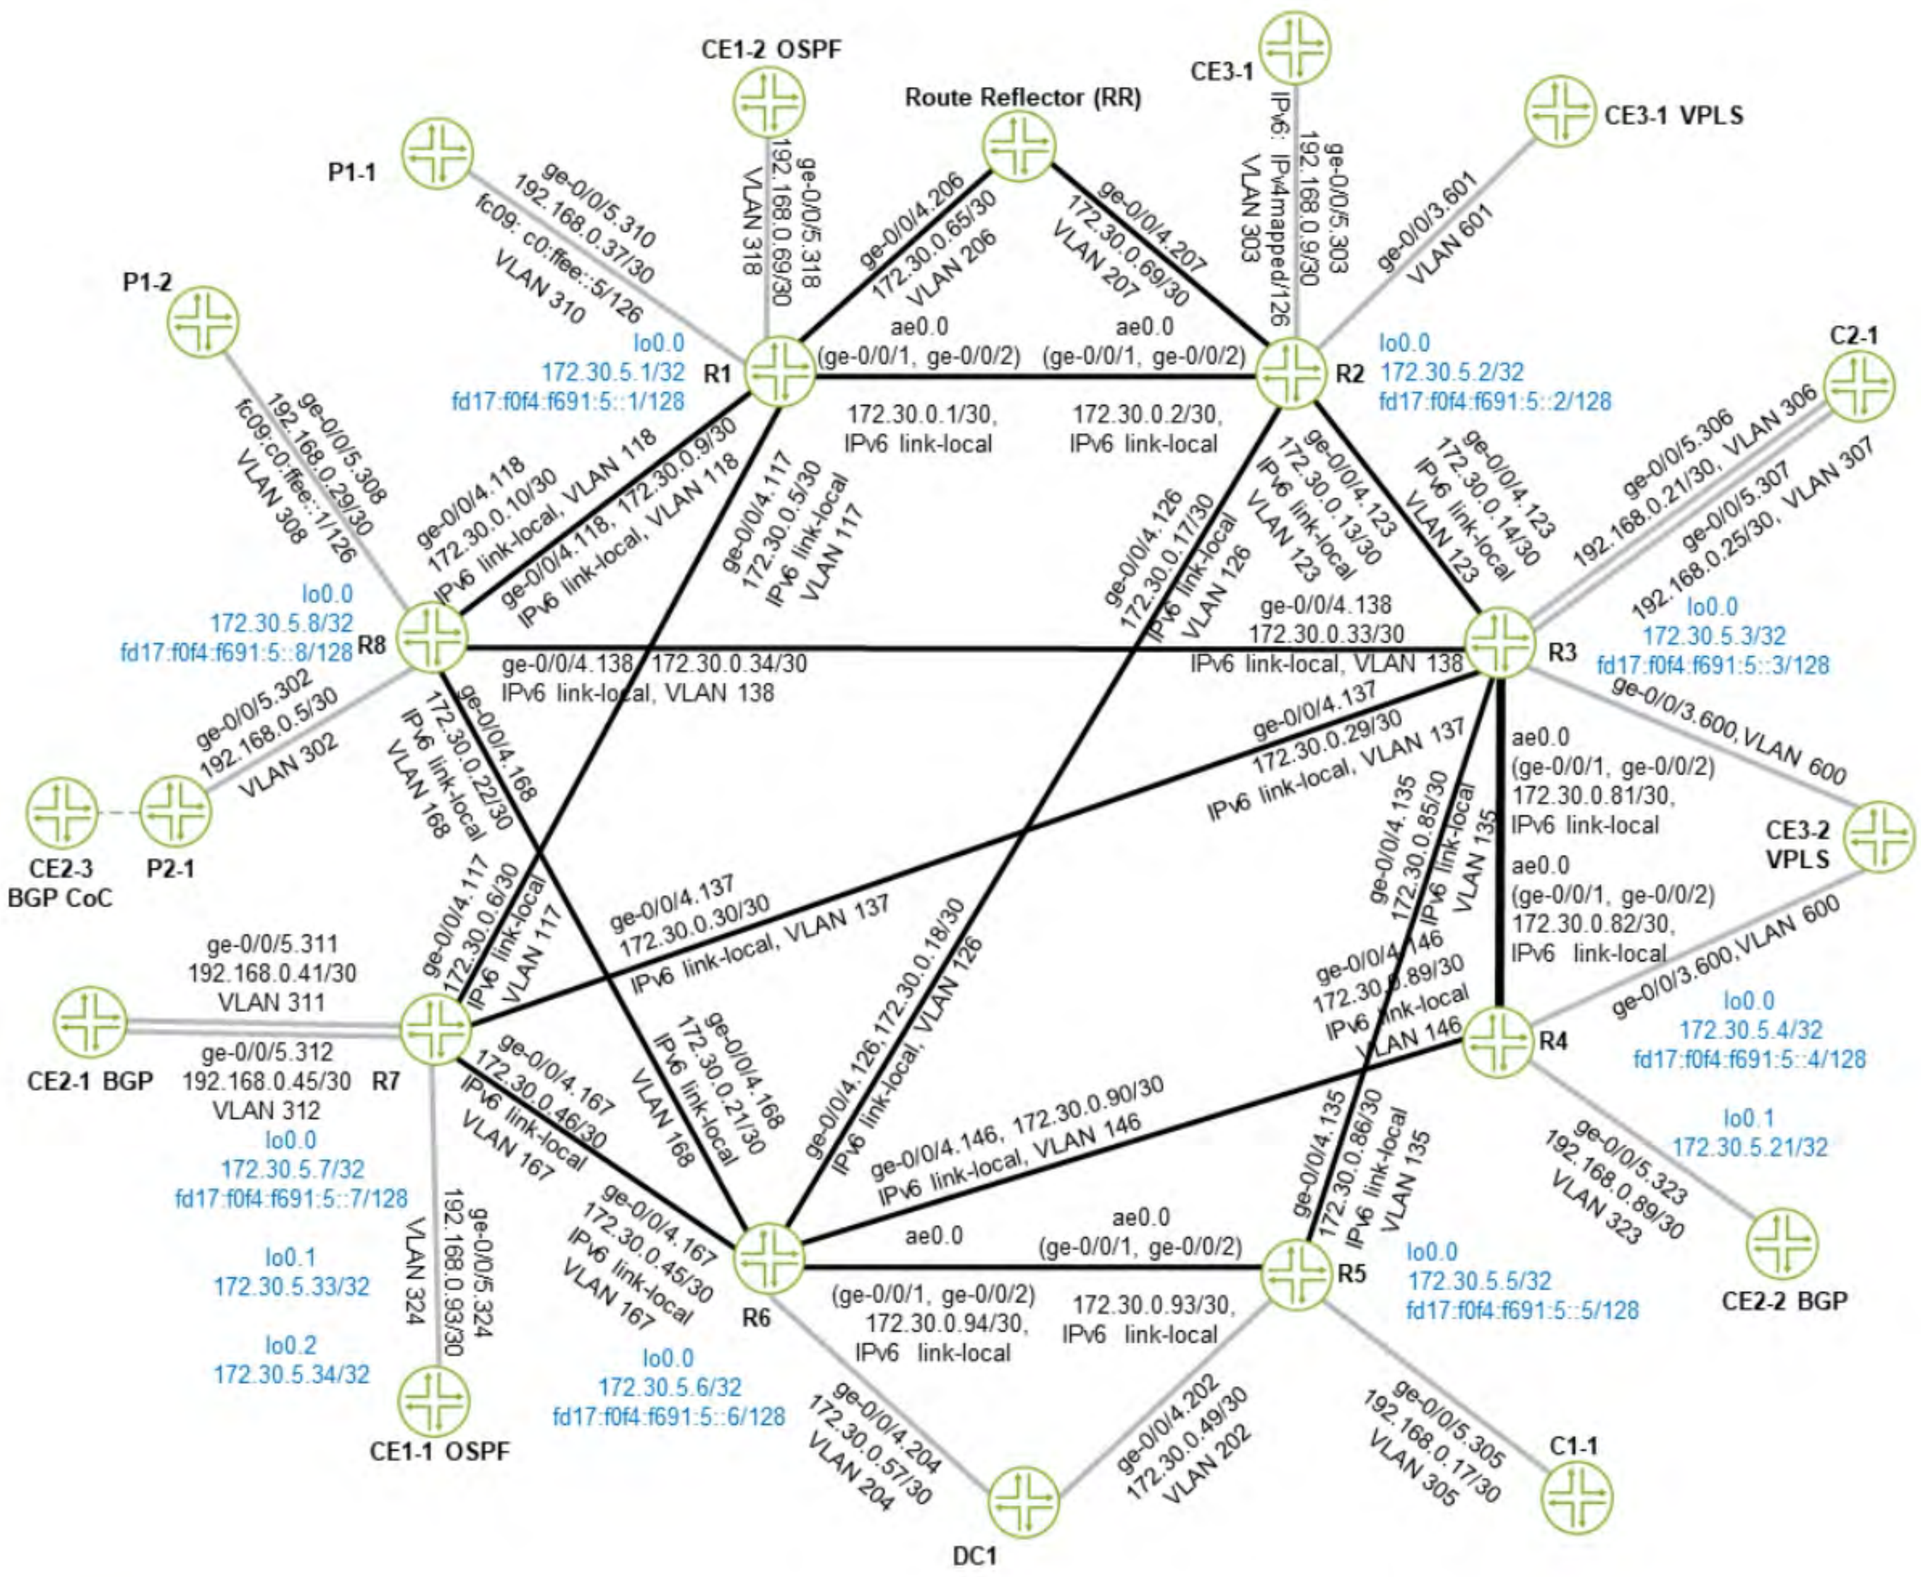
© 2020 Juniper Networks, Inc. All rights reserved.



© 2020 Juniper Networks, Inc. All rights reserved.

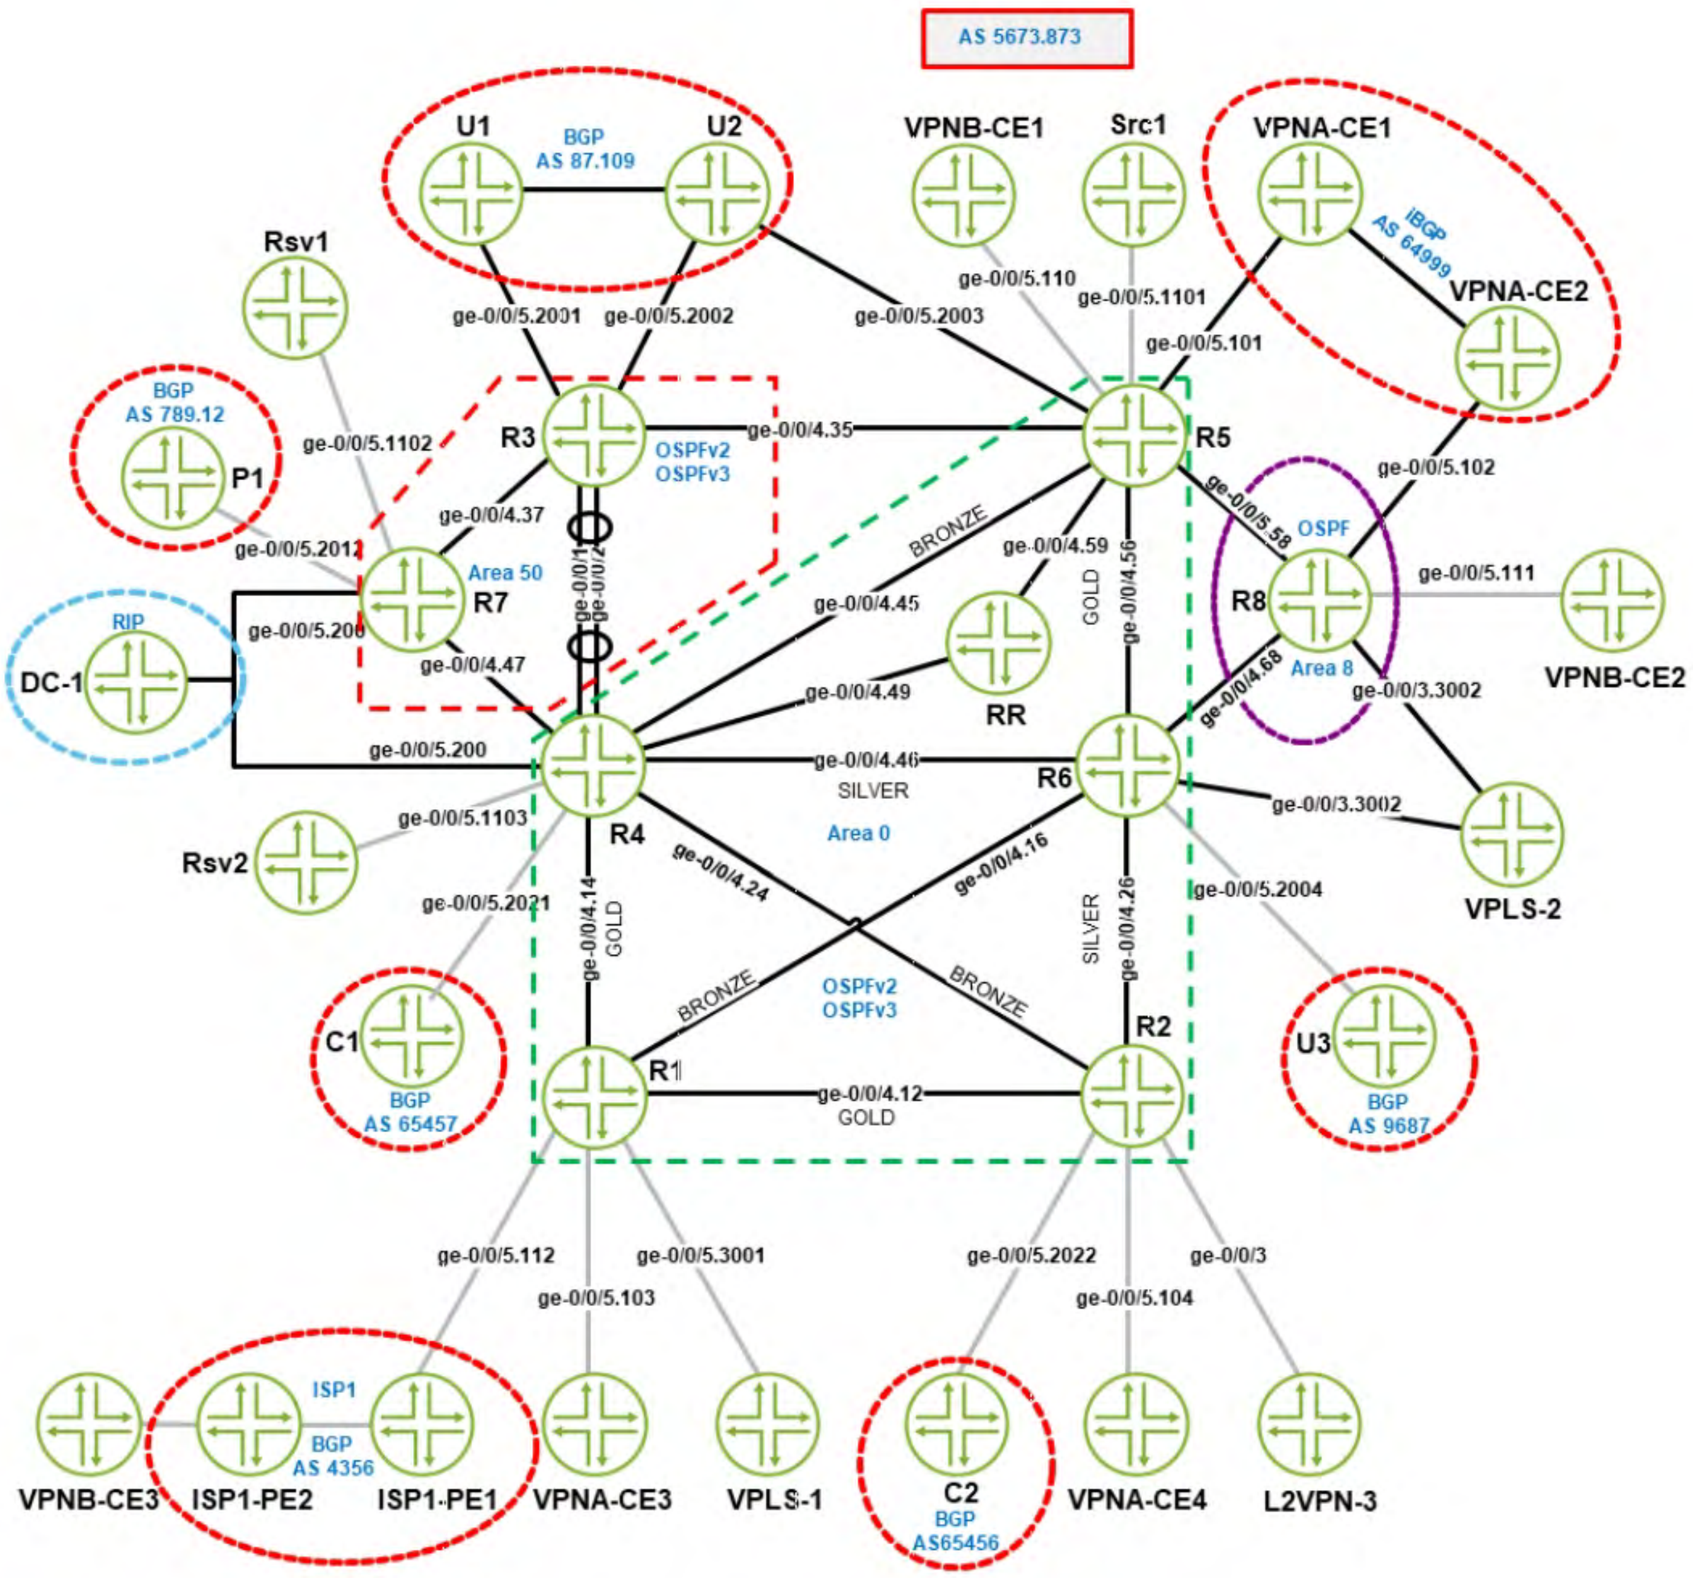


© 2020 Juniper Networks, Inc. All rights reserved.



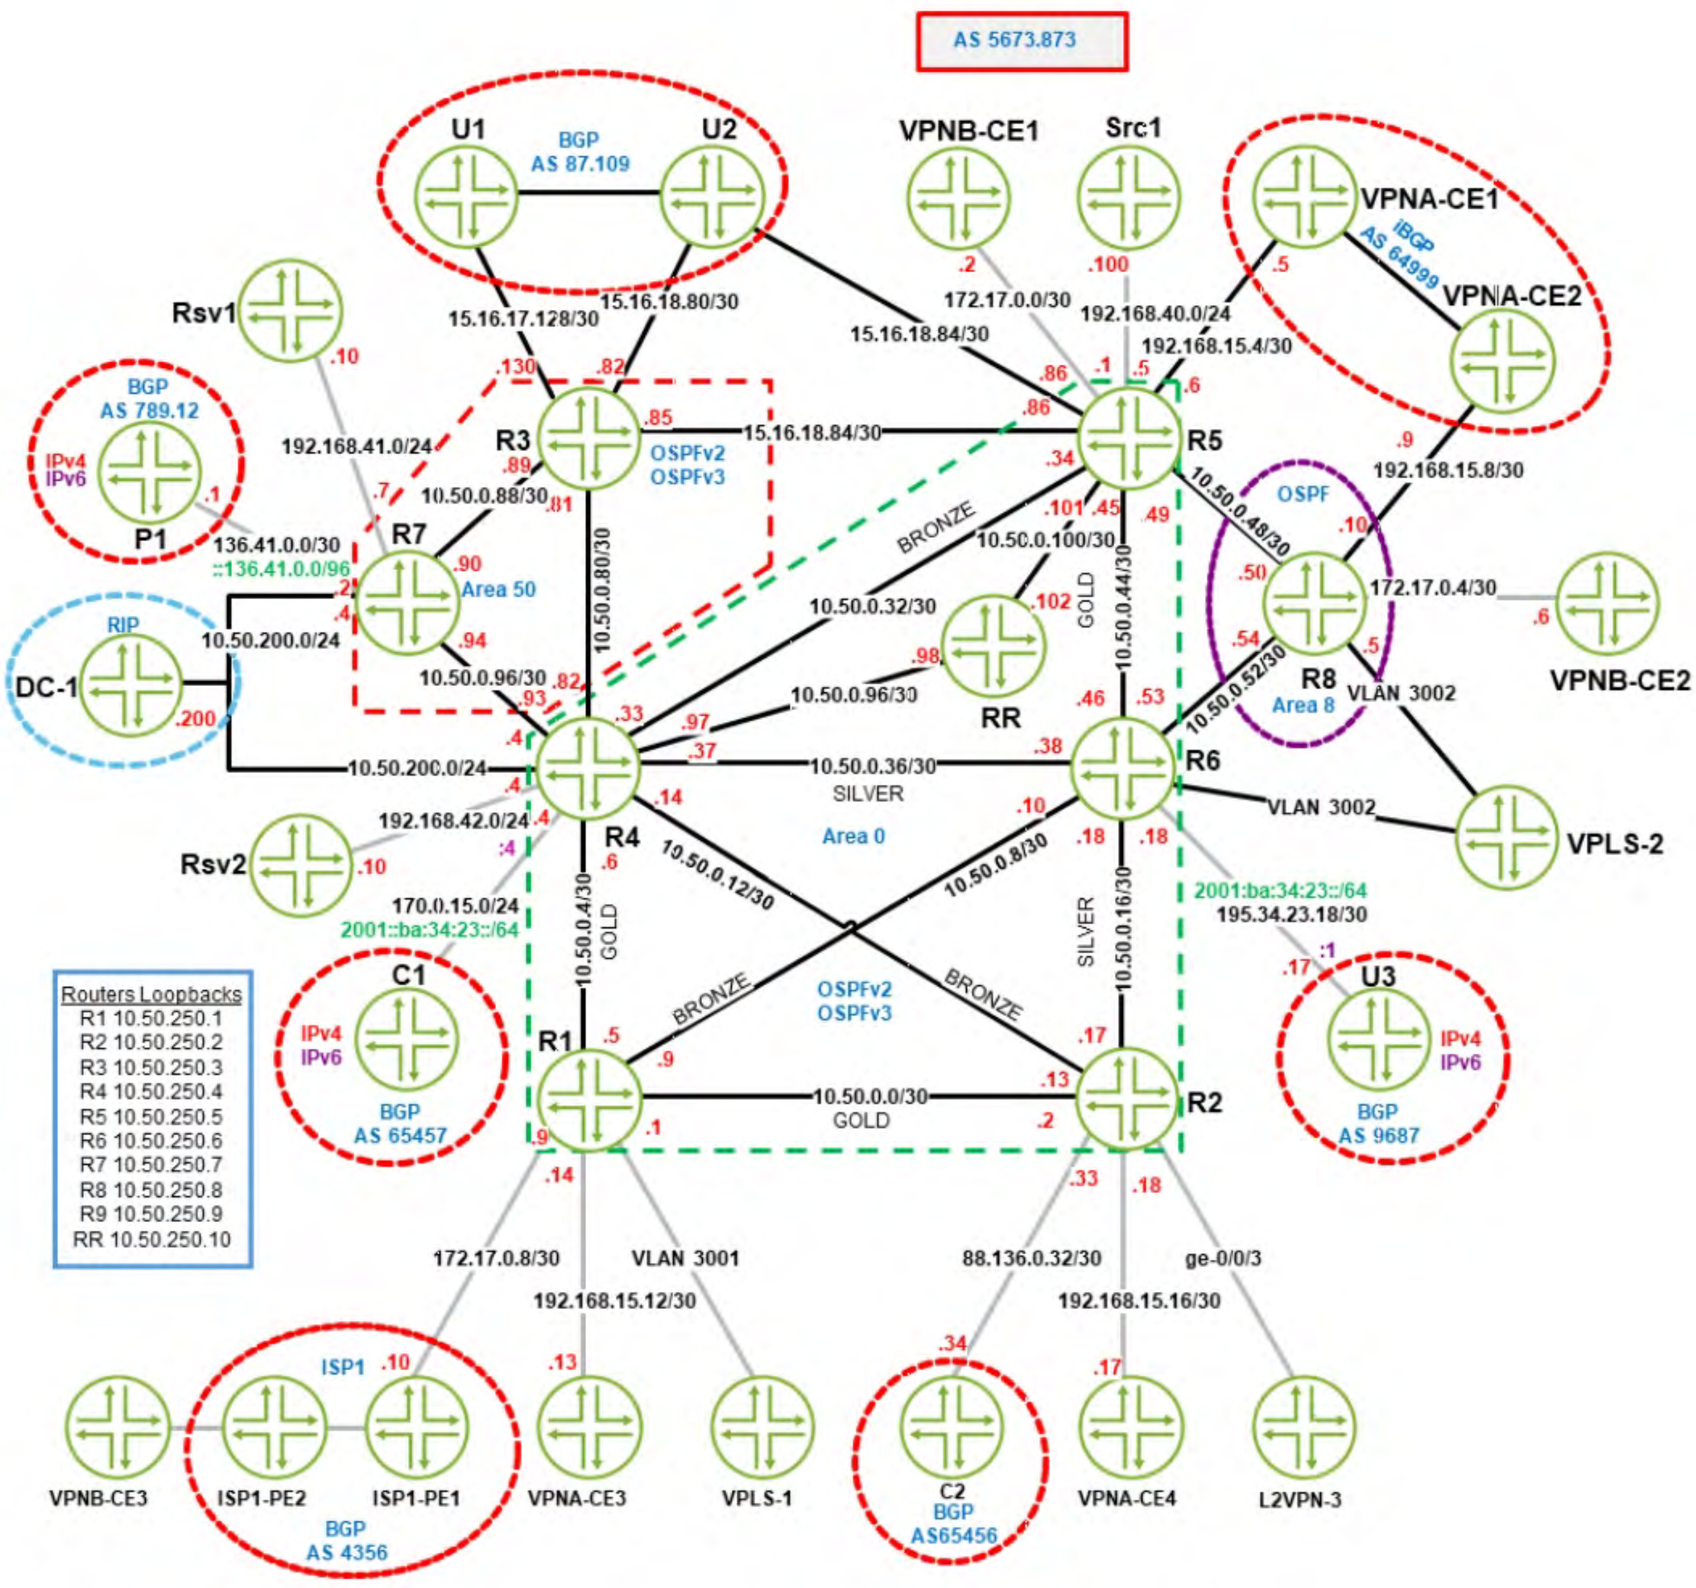
Full day lab 1

© 2020 Juniper Networks, Inc. All rights reserved.

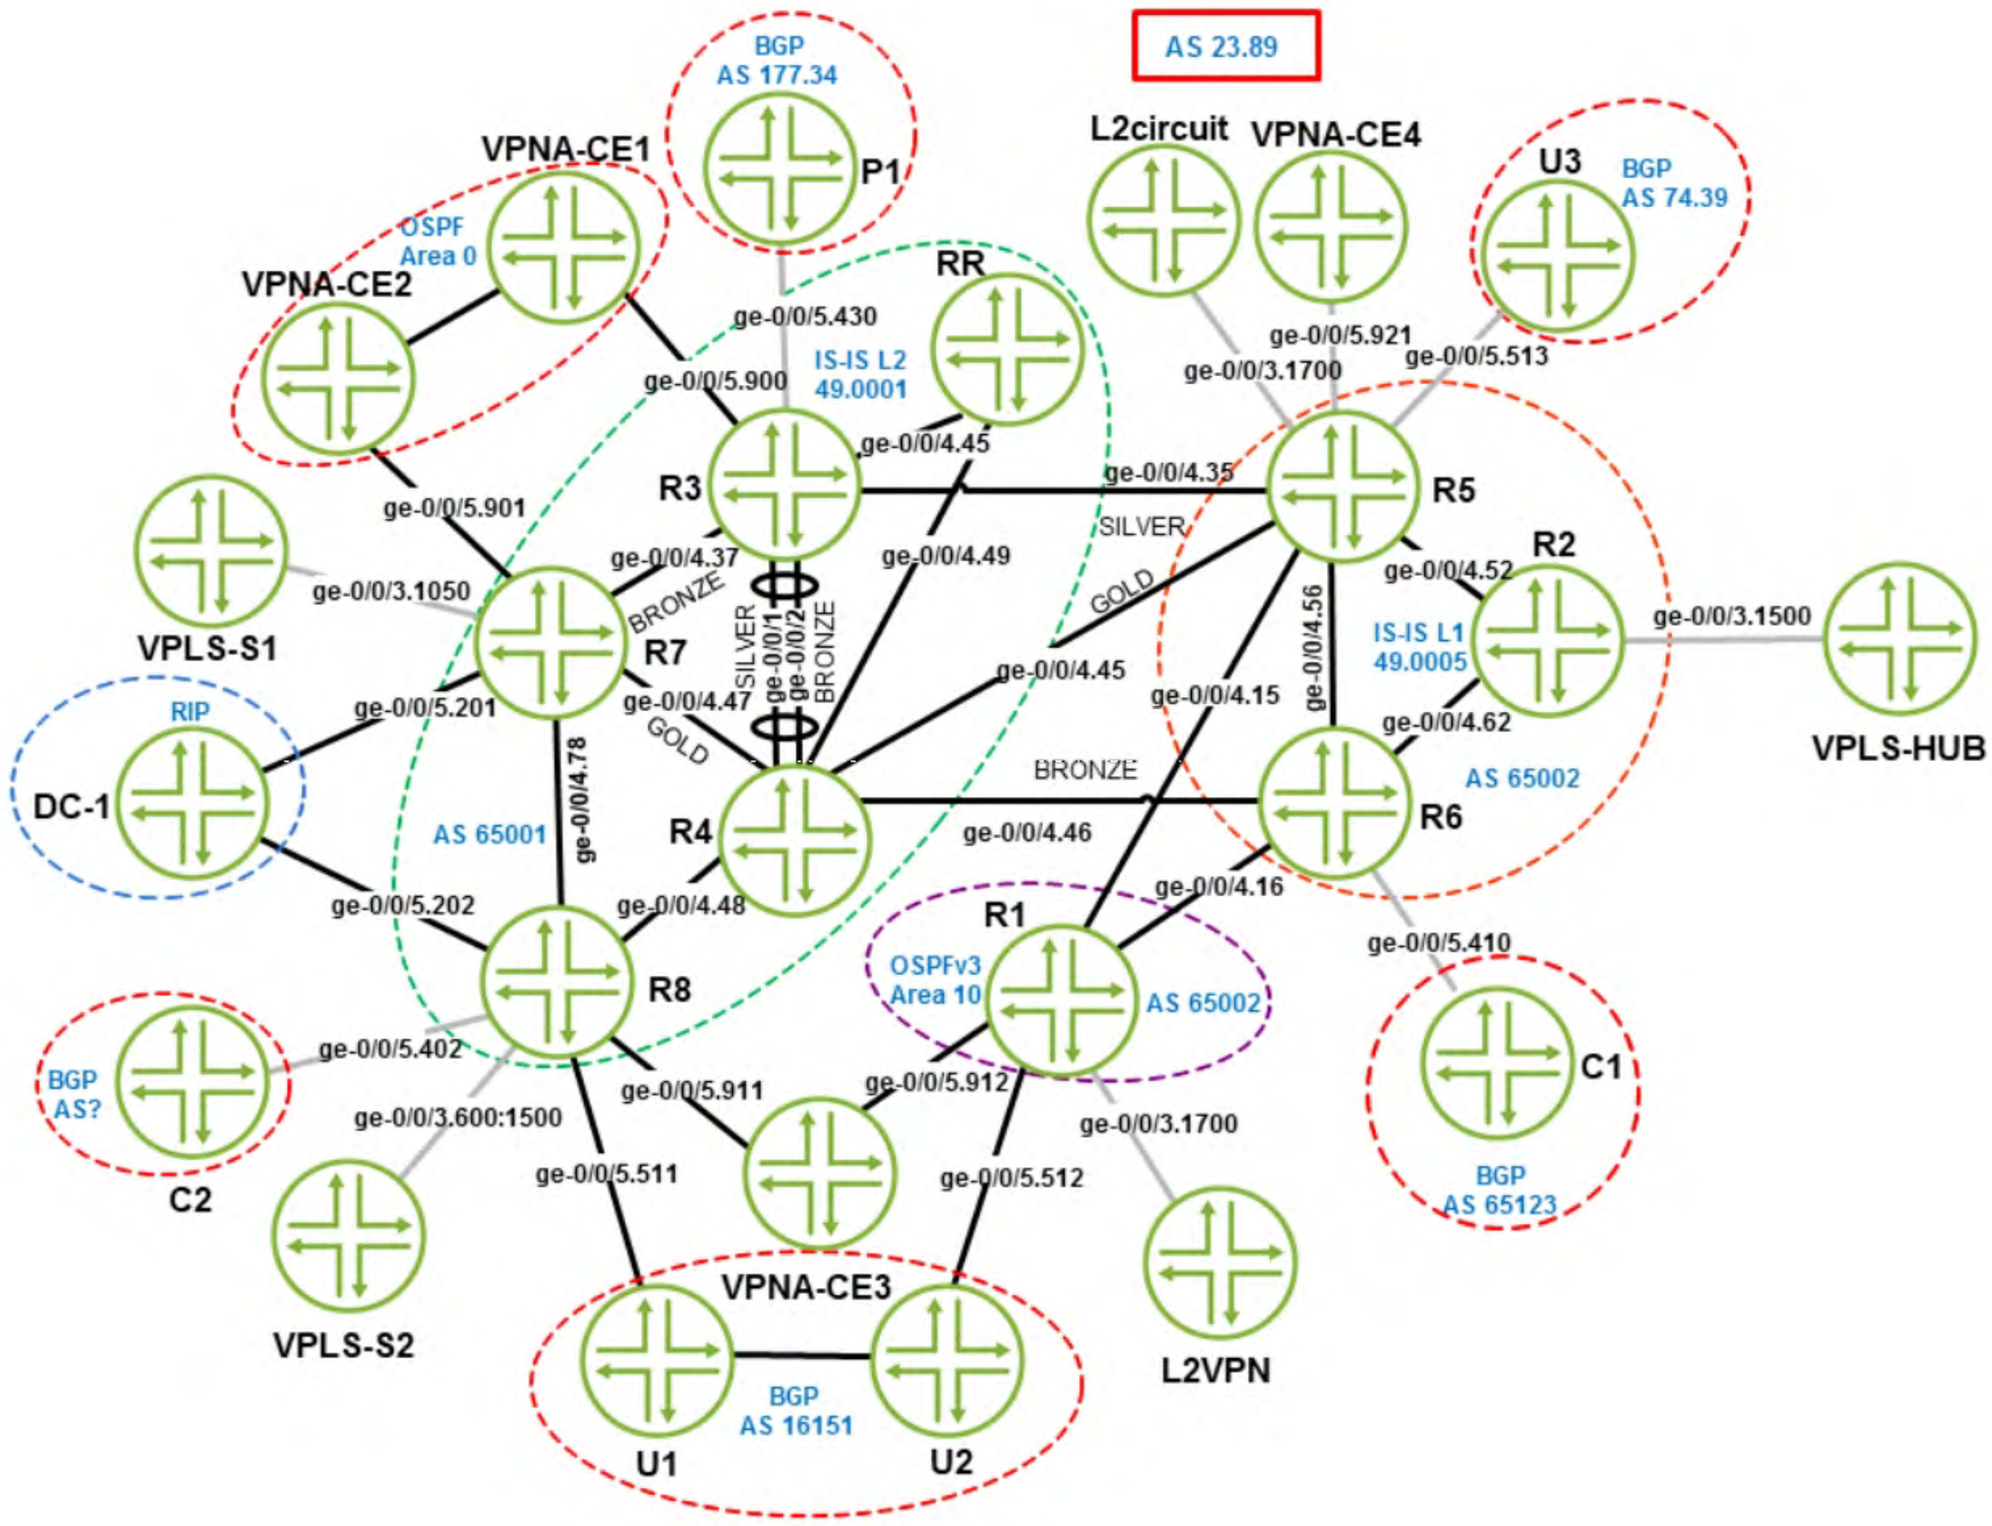


Full day lab 1

© 2020 Juniper Networks, Inc. All rights reserved.



© 2020 Juniper Networks, Inc. All rights reserved.



© 2020 Juniper Networks, Inc. All rights reserved.

Appendix - Chapter 1: General System Features

This appendix provides a detailed walkthrough focusing on the configuration of initial system services, management access to the devices and initial interface configuration. All 8 devices must be configured based on the requirements in the tasks.

Junos Network Operation System Overview

The Junos network operation system runs on most of Juniper Networks hardware. It is powerful, robust, modular, secured and based on a FreeBSD UNIX kernel. All Juniper Networks devices running Junos software have a separate control and data plain architecture. On software-based devices this separation is implemented as different software processes. Whereas high-end devices include different dedicated hardware modules for the control plain (RE) running Junos and the data plain (PFE) running Junos microcode downloaded from the RE. The Routing Engine (RE) is based on x86 architectures on most of the devices and is responsible for handling the protocol processes, all other software processes that control the device's interfaces, other chassis components, system management and user access to the device.

Juniper Networks devices can be configured using a Web GUI and a highly powerful and easy to use CLI. The CLI is used in all Juniper Networks lab exams and will be used throughout this book.

Junos CLI Overview

The Junos CLI can be used for configuration of all features and protocols on the Junos device. One way accessing the CLI is through the dedicated *out of band (OoB) console* interface.

Another option for accessing the CLI is through *inband* via SSH or Telnet terminal protocols. Unlike the console connection inband requires network configuration on the device.

Logging in to the CLI requires a username and password. All Junos devices have only the *root* user enabled in the factory default configuration. Initially, the root user can log in without a password. Log in with user *root* places you into the UNIX shell.

```
Amnesiac (ttyu0)
Login: root
--- Junos 11.4R8.4 built 2013-05-23 05:53:07 UTC
root@R1%
```


© 2020 Juniper Networks, Inc. All rights reserved.

The shell prompt is recognized by the sign **%.** From there you can start the Junos CLI by typing **cli**.

Note: When logged in with user *root,* exiting the CLI with exit returns you to UNIX shell again. A second `exit` command is needed to exit the shell and the device completely.

Starting the Junos CLI with `cli`, the prompt changes to **>** showing the *operational* mode.

```
root@R1% cli
root@R1>
```

Operational mode is used mainly for displaying different information and monitoring all the protocols and processes of the device. Requesting software updates and restarting individual processes are performed also from that mode.

In order to enter *configuration* mode, execute the `configure` or `edit` commands at the operational mode prompt

```
changes to #.
root@R1> configure
Entering configuration mode

[edit]
root@R1#
```

In *configuration* mode the configuration statements are hierarchically organized, so they are easily readable. Throughout this book we will configure many configuration statements and parameters.

Configuration statements entered in *configuration* mode do not take effect immediately. This allows for Configure multiple statements but applying them at single point. Also, this design features provides tools for better and faster interaction with Junos CLI.

The Junos CLI has two types of configuration files – active and candidate. The active configuration is the configuration, which is currently actively running on the device. This is also the configuration, which is loaded when the device boots. The candidate configuration is the configuration that is a candidate to become the active configuration. When entering *configuration* mode, a copy of the active configuration is populated into the candidate configuration. All the changes made in *configuration* mode will also change that candidate configuration file. At any point when you decide to commit the candidate configuration and it will become the active configuration. The commit process checks the candidate configuration for proper syntax and error message are printed on if found.

© 2020 Juniper Networks, Inc. All rights reserved.

A maximum of 50 previous versions (ranging from 0 to 49) of the configuration are kept on the device. 0 represents copy of the actual configuration file. Using `rollback` command followed by number from 0 to 49 loads the previous configuration file specified by that number. The `commit` command is needed to make the roll back effective.

Note: The command rollback `0` wipes out all changes entered to the candidate configuration file and returns to a clear copy of the actual configuration. The `rollback` command can be executed only from the top of the hierarchical CLI tree.

The `set` command inserts changes to the candidate configuration whereas `delete` removes statements. This command deletes changes from both active and candidate configuration files.

Throughout this book many tools available in Junos CLI will be shown. Those tools help to speed up Configure the tasks. Configuration speed is important for JNCIE lab exam exams. For repeated configurations, which you have to apply on several devices it is recommended to use a text editor so it's easy to copy and paste several CLI commands to these devices. During the real JNCIE lab exam you are allowed to use the Windows notepad application for this purpose.

Task 1: Initial System Configuration

The first part from Chapter 1 requires the configuration of the device host names, OoB management, static routing for management, management services, user access, DNS and NTP servers.

1) The initial JNCIE-SP workbook start configurations need to be loaded on all devices (except for the vVR device) before you can start to configure the tasks. Ensure you also load the chapters configuration (vrX-chapters.config) on both VR devices. Typically, you would need to start from scratch using the root user account. For the purposes of this lab you will be using the lab user account (lab/lab123). In this section you will configure this user again as well as other options that have already been configured to ensure they meet the task requirements. Some items may have been intentionally been misconfigured as part of the starting configuration files.

```
lab@R1> configure
Entering configuration mode

[edit]
lab@R1# load override ch1-start.config
load complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit and-quit
commit complete
Exiting configuration mode
```

The configuration command `load` is used to load complete or partial configuration blocks from a file or directly from the terminal. In the screen capture above the option `overwrite` ch1-start.config is used to replace the complete configuration with the initial chapter 1 configuration file.

Note: You can save the complete candidate configuration or parts of it with the `save` command.

2) This step requires that you configure hostnames according the Table 1.

Router	**Router Type**	**Host Name**
vmx1	vMX	R1
vmx2	vMX	R2
vmx3	vMX	R3
vmx4	vMX	R4
vmx5	vMX	R5
vmx6	vMX	R6
vmx7	vMX	R7
vmx8	vMX	R8

Table 1

Note: All the tasks in this chapter should be performed on all 8 routers. We will show the configuration on just one router. You can use the `load merge terminal` command to load the configuration on the other devices after changing the router specific parameters.

Configure host name with following command:

```
[edit]
lab@R1# set system host-name R1
```

Note: We will not commit the configuration after every single change when configured. A commit will be performed after several commands have been entered.

3) Then dedicated management interfaces should be configured for OoB management. IP addresses are listed in Table 2.

© 2020 Juniper Networks, Inc. All rights reserved.

To access the CLI using ssh and Telnet you must first configure an OoB network interface. Most Juniper Networks devices have a dedicated interface for OoB management. The interfaces that should be used for this purpose are listed below.

Host Name	**OoB Interface Name**	**OoB Interface IP Address**
R1	fxp0	10.10.1.1/24
R2	fxp0	10.10.1.2/24
R3	fxp0	10.10.1.3/24
R4	fxp0	10.10.1.4/24
R5	fxp0	10.10.1.5/24
R6	fxp0	10.10.1.6/24
R7	fxp0	10.10.1.7/24
R8	fxp0	10.10.1.8/24

Table 2

First, we will configure the interface descriptions.

```
[edit]
lab@R1# set interfaces fxp0 unit 0 description "OoB Management"
```

The logical properties of the interfaces like the IP address of the interface are configured under *unit* <#>. You also have to specify the protocol family, which must be enabled on the specific interface. In this particular case IP version 4 should be configured which requires the `family inet` configuration under the unit.

```
[edit]
lab@R1# set interfaces fxp0.0 family inet address 10.10.1.1/24
```

4) Now you have to enable SSH, Telnet, and FTP access to the router.

All the management protocols are configured under the `system services` hierarchy.

```
[edit]
lab@R1# set system services ssh

[edit]
lab@R1# set system services telnet

[edit]
lab@R1# set system services ftp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete

[edit]
lab@R1#
```

Note: At this point we perform the `commit` command.

Configure the management protocols under the `system services` hierarchy. You can check that the device is listening on the respective service ports with following command. It is always a good practice to double-check all configuration changes to be sure that you performed the task as requested.

```
lab@R1# run show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address  Foreign Address          (state)
tcp4       0      0  *.23                *.*                LISTEN
tcp4       0      0  *.22                *.*                LISTEN
tcp4       0      0  *.21                *.*                LISTEN
tcp4       0      0  128.0.0.1.33040     128.0.0.1.49832    ESTABLISHED
tcp4       0      0  128.0.0.1.49832     128.0.0.1.33040    ESTABLISHED
tcp4       0      0  *.33039             *.*                LISTEN
tcp4       0      0  *.38                *.*                LISTEN
...
```

5) Next step you have to configure a static route to the remote management network and to make sure that route is not redistributed to any dynamic routing protocols . Also, you must ensure that this network is accessible when the RPD process is not running.

The static route is configured under the `routing-options` hierarchy. Static routes have many optional settings, one of them is the `no-readvertise` setting. This setting ensures that the route is ineligible to be advertised in dynamic routing protocols.

```
[edit]
lab@R1# set routing-options static route 10.10.10.0/24 next-hop 10.10.1.254 no-
readvertise

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

The following output shows that the static route is marked with NoReadvrt flag.

```
[edit]
lab@R1# run show route 10.10.10.0/24 extensive

inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.10.10.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.10.0/24 -> {10.10.1.254}
        *Static Preference: 5
                Next hop type: Router, Next hop index: 550
                Address: 0x156c4d0
                Next-hop reference count: 3
                Next hop: 10.10.1.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 58
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

The second part of this step requires that the route must be reachable when the RPD process is not running. This can happen if the RPD process crashes or when the device is booting. To accomplish this, the backup-router command under the system hierarchy has to be configured. When the RPD process has started this route is removed from the local routing and forwarding tables.

```
[edit]
lab@R1# set system backup-router 10.10.1.254 destination 10.10.10.0/24

[edit]
lab@R1# commit
commit complete
```

6) Configure the routers to use DNS server.

Configure a DNS server provides a capability for resolving names to IP address. The DNS server IP address is provided.

```
[edit]
lab@R1# set system name-server 10.10.1.100
```

Note: If the requirement is to map a static hostname to an IP address, then the command to accomplish this is system static-host-mapping

7) Set the time zone location of the device.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set system time-zone America/Los_Angeles
```

8) Next, all routers should be synchronized with the provided NTP server. Configure the routers to synchronize with that NTP server at boot time as well. Ensure that all communication between the router and the NTP server are authenticated with MD5 hash.

NTP time synchronization is very important for consistent time stamps of the log messages between all devices in the network. This helps to correlate the sequence of events in case of troubleshooting activities.

Junos operating system supports four mode of NTP operation: client, server, symmetrical (peer) and broadcast. Authentication of NTP communication is also supported. NTP boot server is used to synchronize the time during the boot of the router.

```
[edit]
lab@R1# set system ntp boot-server 10.10.1.100

[edit]
lab@R1# set system ntp authentication-key 1 type md5 value workbook

[edit]
lab@R1# set system ntp server 10.10.1.100 key 1

[edit]
lab@R1# set system ntp trusted-key 1

[edit]
lab@R1# commit
commit complete
```

The configured server should send NTP packets that contain the trusted key number specified. The value of the key should match the configured value for that key.

The following show command output checks the successful NTP server synchronization:

```
[edit]
lab@R1# run show ntp status
status=0644 leap_none, sync_ntp, 4 events, event_peer/strat_chg,
version="ntpd 4.2.0-a Thu Sep 20 22:19:19  2018 (1)", processor="amd64",
system="FreeBSDJNPR-11.0-20180816.8630ec5_buil", leap=00, stratum=4,
precision=-23, rootdelay=13.721, rootdispersion=78.033, peer=12620,
refid=10.10.1.100,
reftime=e0460d9a.679acef4  Wed, Mar 27 2019  7:56:26.404, poll=6,
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
clock=e0460dec.7ccde5d1  Wed, Mar 27 2019  7:57:48.487, state=4,
offset=7.281, frequency=5.307, jitter=0.039, stability=0.024

root@R1# run show ntp associations
   remote          refid           st t when poll reach   delay   offset  jitter
==============================================================================
*10.10.1.100       10.210.8.72      3 -   36   64    3    0.385    7.281   0.187
```

The * shows that server is selected for synchronization.

Note: The outputs above are *operational* display commands but are executed from the *configuration* mode by using the `run` command. Every command from *operational* mode can be executed in *configuration* mode like shown in the example. The *configuration* command `show` without `run` prepended displays the candidate configuration. It can be executed from the top of the hierarchical tree to display the whole configuration file. Whereas executed deeper from into the hierarchy displays only that specific part from the candidate configuration.

9) This step requires to configure FTP transfers of the configuration file each time the candidate configuration is committed. The FTP server IP address and credentials are provided.

Archiving files to remote location is helpful in cases where the internal storage on the device fails. Junos provides the capability of automatically uploading the configuration to an FTP or SCP server. To meet the task requirements the configuration must be archived each time the configuration is committed.

```
[edit]
lab@R1# set system archival configuration transfer-on-commit

[edit]
lab@R1# set system archival configuration archive-sites "ftp://lab@10.10.1.100"
password lab123

[edit]
lab@R1# commit
commit complete
```

Note: The remote location where the files will be transferred to will appear in the configuration with quotes since this is an URL. You can configure it without quotes but Junos will add them. Quotes are mandatory when configuring a text string with space symbols or other special characters like the “?” symbol because they have another meaning in the CLI if the quotes are not specified. This is the also the case with the interface description of the management interface configured in step 3.

© 2020 Juniper Networks, Inc. All rights reserved.

10) Next you have to configure the authentication methods that are used by the device to authenticate the users accessing the device by CLI. The RADIUS server IP address and parameters are provided.

The Junos CLI provides a local database of users for authentication and authorization. The local database contains users and passwords stored locally on the device. For authorization the user should be configured locally.

RADIUS and TACACS+ are centralized systems, which can be used for user authentication. Junos devices can send authentication requests to RADIUS or TACACS+ servers. The communication between the Junos devices and the RADIUS or TACACS+ server is authenticated with a shared password. Since for authorization the user should be configured locally on the Junos device, you can create a local template user that can be mapped to different users on the RADIUS or TACACS+ server. This avoids the need of duplicating all RADIUS or TACACS+ users locally. You can override authorization parameters of the template user by passing regular expressions from the server.

You can configure Junos to be RADIUS and TACACS+ client at the same time. The order in which the software checks one or more of the three authentication methods can be prioritized. For each login attempt Junos tries all methods in order and stops when the password is accepted. If the request is rejected or the previous method fails to reply, the next authentication methods in the order is consulted. Local authentication is always consulted as last resort if no reply is received from the previous listed methods. This is valid even if the `password` is not configured in the authentication order.

Under the `radius-server` hierarchy you can configure the RADIUS server port for authentication, the IP address used for sourcing the packets, retry attempts, timeout of the requests. The default port is 1812, 3 retry attempts and the timeout is 3 seconds if not changed.

First, you configure the authentication order as requested for this step.

```
[edit]
lab@R1# set system authentication-order [ radius password ]
```

Then the RADIUS server configuration with provided settings.

```
[edit]
lab@R1# set system radius-server 10.10.1.100 secret workbook

[edit]
lab@R1# set system radius-server 10.10.1.100  timeout 2

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set system radius-server 10.10.1.100 retry 1

[edit]
lab@R1# show | compare
[edit system]
+  authentication-order [ radius password ];
+  radius-server {
+      10.10.1.100 {
+          secret "$9$fz9pKvL7dscyeWxdg4fTz6Cu"; ## SECRET-DATA
+          timeout 2;
+          retry 1;
+      }
+  }

[edit]
lab@R1# commit
commit complete
```

Note: Another powerful tool of the Junos CLI is demonstrated. The `show | compare` command displays the changes between the active configuration and the candidate configuration. This shows what will become active after issuing a `commit`. The "+" sign marks the lines which will be added, whereas the "-" sign marks the lines that are removed. The `show | compare rollback #` command compares the candidate configuration with previous configuration with number "#".

Note: On the output shown the RADIUS shared secret is displayed as a MD5 hash with tag ## *SECRET-DATA*. The clear text secret is not visible.

11) New user lab has to be created with super user privileges.

When configuring a local user, a `class` must also be specified.

The `class` parameter is a group of privileges to which one or more users can belong. Classes contain explicit permissions and other parameters. For example, the only way you can configure the `idle-timeout` for a user is setting it in a class. Users can only be member of single class.

Four predefined classes are available in Junos CLI. The following permissions are configured for those classes:

- *Super-user*: all permissions
- *Operator*: clear, network, trace, reset and view
- *Read*-only: view permissions only
- *Unauthorized*: no permissions

© 2020 Juniper Networks, Inc. All rights reserved.

These permissions cannot be overwritten. If you want to create a user that has all permissions, but need an idle-timeout configured also, you have to configure a new class with super-user privilege and configure the idle-timeout for it. Then apply the users to that class.

Permissions are a predefined set of commands and configuration statements that require the same access privilege level. Many permission-sets are preconfigured. You can list and read a description of all predefined permissions with the command `help topic system permissions`.

The requirement of this step is to configure a `lab` user and assign it to `super-user` class.

```
[edit]
lab@R1# set system login user lab class super-user authentication plain-text-password
New password:
Retype new password:

[edit]
lab@R1# commit
commit complete
```

Note: All passwords must be minimum of 6 characters and require a change in case, digits or punctuation. As seen below.

```
[edit]
lab@R1# set system login user lab class super-user authentication plain-text-password
New password:
error: minimum password length is 6
error: require change of case, digits or punctuation
```

12) The next step requires configuration of additional users according Table 3.

Username	**Password**	**Privileges**
any	-	Permissions “view” and “view-configuration”. Authenticated only by the RADIUS
ops	ops123	Permissions “clear”, “network”, “reset”, “trace” and “view”
noc	noc123	Permissions “all”. Additionally, cannot execute any of the “clear”, “configure”, “edit” or “start shell” commands

Table 3

© 2020 Juniper Networks, Inc. All rights reserved.

The first user can have literally any username and be given `view` and `view-configuration` permissions. The restriction is that this user should be authenticated only by RADIUS. However, the local database is always consulted as last resort.

As mentioned earlier, users can be mapped to a local or remote template account. Local user template accounts can have any name. Whereas the remote template must use the predefined name: `remote`. The difference between the normal user account and the *template* account is that the *template* user account does not have a password configured locally on the device. The `remote` user template account is used when the authenticated user does not exist locally on the device.

To comply to the restrictions as mentioned in the task you can use the predefined user template `remote`. This way you can use any configured user account on the RADIUS server to login.

You should restrict that user *template* to have only `view` and `view-configuration` permission. Because none of the predefined classes have exactly these permissions you have to create a new one.

```
[edit]
lab@R1# set system login class limited permissions [ view view-configuration ]
```

Then you can assign user `remote` to that class.

```
[edit]
lab@R1# set system login user remote class limited
```

The second user in the table should have the name `ops` with the predefined `operator` class permissions.

```
[edit]
lab@R1# set system login user ops class operator authentication plain-text-password
New password:
Retype new password:
```

The last user in the table with name `noc` should have all permissions, but with the restrictions of executing some commands.

You can use `deny-commands` and `allow-commands` to restrict specific *operational* commands. The `deny-configuration` and `allow-configuration` options can be used for configuration commands. Regular expressions for matching the respecting commands

© 2020 Juniper Networks, Inc. All rights reserved.

can be used. Junos applies the deny statements before the `allow` statements, whereas both take precedence over the permissions sets configured for the users.

For user noc you have to configure again a new login class, as none of the predefined ones can be used.

```
[edit]
lab@R1# set system login class privileged permissions all
```

First you attach all permissions set to new class.

```
[edit]
lab@R1# set system login class privileged deny-commands
"(clear)|(configure)|(edit)|(start shell)"

[edit]
lab@R1# set system login user noc class privileged

[edit]
lab@R1# set system login user noc authentication plain-text-password
New password:
Retype new password:

[edit]
lab@R1# commit
commit complete
```

Note: Each expression is separated by the pipe | symbol and must be a complete standalone expression enclosed with parentheses ( ). Do not use spaces between the regular expressions separated by parentheses and connected with the pipe | symbol.

Note: The command `edit` like the command `show` can be used in *operational* and *configuration* modes. In *configuration* mode `edit` is used to move deeper into the configuration tree. Whereas used in *operational* mode this command does the same as `configure`, it enters *configuration* mode. In *operational* mode this command is hidden. This means that it should be typed as a whole command to be executed. Auto completion for hidden commands cannot be used. Also "?" cannot be used to display the hidden commands into the list of possible commands.

The following command displays the privileges of the current logged in user. You must login as user noc to display the authorization of user noc.

```
noc@R1> show cli authorization
Current user: 'noc          ' class 'privileged'
Permissions:
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    admin        -- Can view user accounts
    admin-control-- Can modify user accounts
    clear        -- Can clear learned network info
    configure    -- Can enter configuration mode
    control      -- Can modify any config
    edit         -- Can edit full files
    field        -- Can use field debug commands
    floppy       -- Can read and write the floppy
    interface    -- Can view interface configuration
    interface-control-- Can modify interface configuration
    network      -- Can access the network
    reset        -- Can reset/restart interfaces and daemons
    routing      -- Can view routing configuration
    routing-control-- Can modify routing configuration
    shell        -- Can start a local shell
    snmp         -- Can view SNMP configuration
    snmp-control-- Can modify SNMP configuration
    system       -- Can view system configuration
    system-control-- Can modify system configuration
    trace        -- Can view trace file settings
    trace-control-- Can modify trace file settings
    view         -- Can view current values and statistics
    maintenance -- Can become the super-user
    firewall     -- Can view firewall configuration
    firewall-control-- Can modify firewall configuration
    secret       -- Can view secret statements
    secret-control-- Can modify secret statements
    rollback     -- Can rollback to previous configurations
    security     -- Can view security configuration
    security-control-- Can modify security configuration
    access       -- Can view access configuration
    access-control-- Can modify access configuration
    view-configuration-- Can view all configuration (not including secrets)
    flow-tap     -- Can view flow-tap configuration
    flow-tap-control-- Can modify flow-tap configuration
    idp-profiler-operation-- Can Profiler data
    pgcp-session-mirroring-- Can view pgcp session mirroring configuration
    pgcp-session-mirroring-control-- Can modify pgcp session mirroring
configuration
    all-control -- Can modify any configuration
Individual command authorization:
    Allow regular expression: none
    Deny regular expression: (clear)|(configure)|(edit)|(start shell)
    Allow configuration regular expression: none
    Deny configuration regular expression: none
```


© 2020 Juniper Networks, Inc. All rights reserved.

13) Configure syslog as shown in Table 4.

Receiver	**Message Type**
File “jncie-sp-messages”	All info level messages
Syslog server S1	Interactive commands
	Configuration changes
	All notice level messages
File “user-commands”	All user interactive commands
User “ops”	All warning level messages

Table 4

Syslog is UNIX mechanism for collecting system and process messages. The main syslog file for storing messages on Junos devices (default) is `messages`, all files generated by the syslog process are stored in the `/var/log/` directory. Syslog provides a client server architecture, where the Junos device is a client sending messages to a remote server.

Junos syslog supports several facilities and severity levels. The facilities define the class of logged messages. The severity levels define the level of detail of the messages.

Junos factory default configuration defines basic syslog configuration. Before you configure the task requirements, you should check the current configuration under the `system syslog` hierarchy and delete the whole `syslog` configuration.

```
[edit]
lab@R1# show system syslog
file user-commands {
    interactive-commands any;
}
file jncie-sp-messages {
    any notice;
}

[edit]
lab@R1# delete system syslog

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Now you can start configuring this step according to the table provided in the workbook.

The first requirement is to configure a local file named `jncie-sp-messages`, which collects all up to the `info` facility level.

```
[edit]
lab@R1# set system syslog file jncie-sp-messages any info
```

Next you have to configure a remote syslog server as required in the table.

```
[edit]
lab@R1# set system syslog host 10.10.1.100 interactive-commands any

[edit]
lab@R1# set system syslog host 10.10.1.100 change-log any

[edit]
lab@R1# set system syslog host 10.10.1.100 any notice
```

Then again you have to configure a local file capturing all commands issued on CLI.

```
[edit]
lab@R1# set system syslog file user-commands interactive-commands any
```

For the last two requirements you have to log messages, which must appear directly in the user CLI session. Usually it is a good idea to ensure that all warning and critical messages appear immediately to all users which are currently logged in.

```
[edit]
lab@R1# set system syslog user ops any warning
```

14) Last step of Task 1 is to configure the automatic archival of the syslog file.

Local storage has limited space. Because of this, by default Junos stops writing messages to a log file if the file reaches 128kb. When the size limit has been reached, a numbered suffix is added to the filename and a new file name using the original name is created. Junos starts to write the log messages to the new one until it again reaches 128kb. By default, the files are rotated like this until 10 files are created. After that the syslog utility starts to overwrite the files.

For the task requirements the default settings need to be adjusted. The log files have to be rotated until 3 files are created with a size of 100kb.

© 2020 Juniper Networks, Inc. All rights reserved.

Settings for log file rotation can be configured individually for each syslog file under `system syslog file <filename>` or globally for all configured syslog files under the `system syslog` hierarchy. For this task you can use the global configuration.

```
[edit]
lab@R1# set system syslog archive files 3 size 100k

[edit]
lab@R1# commit
commit complete
```

This was the last step from Task 1. All changes made so far have been done on a single router (R1) only. To speed up the configuration on rest of the devices you can use the active configuration file from R1. Paste it in text editor, change the appropriate parameters for the other routers. Then use `load override terminal` to load the configuration on the other routers.

Task 2: SNMP Configuration

Simple Network Management Protocol (SNMP) is a network management protocol standardized in three different versions 1, 2c, 3. SNMP is based on the concept of agents and managers. Junos acts as a SNMP agent, whereas the management function is provided by an external network management system. Management stations poll the agents for specific Management Information Base (MIB) values. MIBs are a collection of objects in a hierarchical fashion maintained by the agents. SNMP supports a mechanism to notify the SNMP manager in case an event occurred on the network device. Version 2 uses traps for notifications, which are not acknowledged. Whereas SNMP version 3 uses informs to notify the manager, which is required to acknowledge the receipt of packets. SNMP allows for read and write access to agents. SNMP relies on the UDP protocol.

SNMP v1 is the initial version of SNMP and does not have built in security features. Version 2c makes use of community strings and clear text password to determine the access to the agent MIBs. Version 3 is the latest enhanced protocol specification, which provides sophisticated security features. Defining a user-based security model (USM) and a view-based access control model (VACM), both representing the access control service of the protocol. Messages exchanged can be send in clear text (noAuthNoPriv), only authenticated (authNoPriv) or authenticated and encrypted (authPriv).

1) Configure the access from the network management station to the agent. Table 5 defines the restrictions.

    a) You have to configure a user with both authentication and encryption.

© 2020 Juniper Networks, Inc. All rights reserved.

Parameter	Value
USM username	lab
USM user authentication	SHA
USM user authentication password	workbook
USM user encryption	3DES
USM user encryption password	workbook
VACM security model	usm
VACM user	lab
VACM security level	privacy
VACM read view OID	.1

Table 5

```
[edit]
lab@R1# set snmp v3 usm local-engine user lab authentication-sha authentication-key
workbook

[edit]
lab@R1# set snmp v3 usm local-engine user lab privacy-3des privacy-key workbook
```

b) Then you have to define the MIB view. MIBs are structures in the objects identifiers (OIDs) hierarchy. Table 5 requires configuration view from the top of the tree which is OID .1

```
[edit]
lab@R1# set snmp view root-view oid .1 include
```

2) Next you have to configure the view-based access control. Users are assigned to a group, which defines the access policy.

```
[edit]
lab@R1# set snmp v3 vacm security-to-group security-model usm security-name lab group
primary-group

[edit]
lab@R1# set snmp v3 vacm access group primary-group default-context-prefix security-
model usm security-level privacy read-view root-view
```

Note: The requirement that SNMPv3 must provide read access only is accomplished with `read-view` configured under the `security-level` hierarchy.

© 2020 Juniper Networks, Inc. All rights reserved.

3) Next you have to configure SNMPv3 notification parameters. This will enable sending traps to the management station.

a) Configure traps to be used for notification and filter which type of traps must be sent.

Parameter	Value
Target address	lab
Target processing model	SHA
Target security model	workbook
Target security level	3DES
Target security name	workbook
Notification OID filter	usm
Notification type	lab

Table 6

```
[edit]
lab@R1# set snmp v3 notify-filter all-traps oid snmpTraps

[edit]
lab@R1# set snmp v3 notify-filter all-traps oid jnxTraps

[edit]
lab@R1# set snmp v3 notify traps type trap

[edit]
lab@R1# set snmp v3 notify traps tag all-nms
```

Note: OIDs under the `notify-filter` should be entered exactly as how they are specified in Table 6. You cannot see the list of all traps supported in the Junos CLI. If you make a typo Junos will commit but with a warning message that this OID is not recognized. You can use the `exclude` keyword to ensure that a trap will not be sent to the management station.

```
[edit]
lab@R1# set snmp v3 notify-filter all-traps oid linkUp exclude
```

b) Configure the destination parameters of the management station.

```
[edit]
lab@R1# set snmp v3 target-parameters S1-parameters parameters security-model usm
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set snmp v3 target-parameters S1-parameters parameters security-level privacy

[edit]
lab@R1# set snmp v3 target-parameters S1-parameters parameters message-processing-model
v3

[edit]
lab@R1# set snmp v3 target-parameters S1-parameters parameters security-name lab
```

Note: All parameters under `target-parameters` are mandatory. If you omit just one, Junos displays a warning when the candidate configuration is listed like the output below. If you commit the candidate configuration without the mandatory statements, the commit will fail with the following error *"missing mandatory statement"*

```
target-parameters S1-parameters {
        parameters {
            security-model usm;
            security-level privacy;
            ## Warning: missing mandatory statement(s): 'message-processing-
model', 'security-name'
        }
}
```

c) Next you have to configure the destination server collecting the traps. Using the target-parameters and notify stanzas configured in previous steps.

```
[edit]
lab@R1# set snmp v3 target-address S1 address 10.10.1.100

[edit]
lab@R1# set snmp v3 target-address S1 target-parameters S1-parameters

[edit]
lab@R1# set snmp v3 target-address S1 tag-list all-nms

[edit]
lab@R1# commit
commit complete
```

Display all SNMPv3 configurations.

```
[edit]
lab@R1# run show snmp v3 | no-more

Local engine ID: 80 00 0a 4c 01 ac 1e 05 01
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Engine boots:             7
Engine time:         177213 seconds
Max msg size:         65507 bytes

Engine ID: local
    User                            Auth/Priv   Storage      Status
    lab                              sha/3des   nonvolatile  active

Group name           Security                Storage      Status
                     model     name                  type
primary-group        usm       lab                   nonvolatile  active

Access control:
Group                Context Security      Read        Write      Notify
                     prefix  model/level   view        view       view
primary-group                 usm/privacy  root-view

SNMP Target:
Address     Address                    Port  Parameters  Storage     Status
name                                         name        type
S1          10.10.1.100                162   S1-paramete nonvolatile active

Parameters     Security        Security     Notify  Storage      Status
name           name            model/level  filter  type
S1-parameters  lab              usm/privacy         nonvolatile  active

SNMP Notify:
Notify               Tag                Type         Storage      Status
name                                                 type
traps                all-nms            trap         nonvolatile  active

Filter               Subtree            Filter       Storage      Status
name                                    type         type
all-traps            1.3.6.1.4.1.2636.  include      nonvolatile  active
all-traps            1.3.6.1.6.3.1.1.5  include      nonvolatile  active
all-traps            1.3.6.1.6.3.1.1.5  exclude      nonvolatile  active
```


© 2020 Juniper Networks, Inc. All rights reserved.

SNMP statistics reveal the number of the traps sent to the management station.

```
[edit]
lab@R1# run show snmp statistics
SNMP statistics:
  Input:
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 0, Total set varbinds: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0, Duplicate request drops: 0
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 0
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 0, Wrong digests: 0, Decryption errors: 0
  Output:
    Packets: 2, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 2
...
```

Now that you are at the end of this step you can copy the configuration and load it on the other devices to save time.

Task 3: Firewall filters

This task requires the configure of firewall filters to control traffic destined to the routing engine. Firewall filters are stateless filters, which can be applied to transit traffic and also the traffic entering or leaving the router itself. Firewall filters can perform several actions on the traffic matched. They consist of one or more (named) terms. Each term has group of match conditions and one or multiple actions or non-terminating actions. If no traffic is matched by any of the terms, there is an implicit rule that discards all traffic. This last rule is hidden and not displayed into the configuration.

Firewall filters are evaluated sequentially from top to down. If a match is found the actions are performed. Firewall filters match on different conditions for the different protocol families. If you configure the filter without a protocol family, Junos recognize it as a `family inet` filter matching IPv4 traffic by default.

© 2020 Juniper Networks, Inc. All rights reserved.

You can match different conditions in firewall filters, including source or destination IP addresses, IP subnets, Layer 4 protocols, source or destination ports, QoS settings and many others. Some of the actions applied to the matched traffic include:

- *Accept* – packet is accepted
- *Discard* – packet is silently dropped
- *Reject* – packet is rejected and an ICMP message is generated
- *Policer* – rate-limit is applied to the traffic
- *Log* – logs the packet header directly on the PFE
- *Count* – counts the matched packets

1) Configure firewall filter rules permitting required protocols.

Note: This task requires that you know how the protocols work, how they are transported through the network and how you can match them in your filter. Note that Junos has predefined names for ports used by the well-known protocols running in the network.

Configure a separate term for each protocol that must be accepted by the routing engine.

```
[edit]
lab@R1# set firewall family inet filter protect-re term ah from protocol ah

[edit]
lab@R1# set firewall family inet filter protect-re term ah then accept

[edit]
lab@R1# set firewall family inet filter protect-re term bfd from protocol udp

[edit]
lab@R1# set firewall family inet filter protect-re term bfd from port 3784

[edit]
lab@R1# set firewall family inet filter protect-re term bfd then accept
```

Note: UDP ports 3784 and 3785 are reserved for BFD. Port 3785 is used for BFD echo function. Here you can use only port 3784 used for normal BFD function for the sender and responder.

```
[edit]
lab@R1# set firewall family inet filter protect-re term vrrp from protocol vrrp

[edit]
lab@R1# set firewall family inet filter protect-re term vrrp then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set firewall family inet filter protect-re term rip from protocol udp

[edit]
lab@R1# set firewall family inet filter protect-re term rip from port rip

[edit]
lab@R1# set firewall family inet filter protect-re term rip then accept

[edit]
lab@R1# set firewall family inet filter protect-re term ospf from protocol ospf

[edit]
lab@R1# set firewall family inet filter protect-re term ospf then accept
```

Note: OSPF is not using TCP or UDP as transport protocols. It is encapsulated directly in IP packets with protocol number 89.

```
[edit]
lab@R1# set firewall family inet filter protect-re term ldp from protocol udp

[edit]
lab@R1# set firewall family inet filter protect-re term ldp from protocol tcp

[edit]
lab@R1# set firewall family inet filter protect-re term ldp from port ldp

[edit]
lab@R1# set firewall family inet filter protect-re term ldp then accept
```

Note: The Label Distribution Protocol LDP uses port 646 over UDP and TCP. The UDP protocol is used for transporting hello discovery messages, whereas TCP is used for label exchange.

```
[edit]
lab@R1# set firewall family inet filter protect-re term rsvp from protocol rsvp

[edit]
lab@R1# set firewall family inet filter protect-re term rsvp then accept

[edit]
lab@R1# set firewall family inet filter protect-re term pim from protocol pim

[edit]
lab@R1# set firewall family inet filter protect-re term pim then accept


[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set firewall family inet filter protect-re term igmp from protocol igmp

[edit]
lab@R1# set firewall family inet filter protect-re term igmp then accept
```

Note: IGMP is used for multicast group registration join messages. Routing protocols like RIPv2 and OSPF use multicast by sending join messages to well-known multicast IP address(es). For example, RIP uses 224.0.0.9 to send and receive protocol packets.

```
[edit]
lab@R1# set firewall family inet filter protect-re term msdp from protocol tcp

[edit]
lab@R1# set firewall family inet filter protect-re term msdp from port msdp

[edit]
lab@R1# set firewall family inet filter protect-re term msdp then accept

[edit]
lab@R1# commit
commit complete
```

2) Next you have to configure a term accepting BGP messages but only from configured neighbors.

   a) To accomplish this step, you need to configure filter term, which adds dynamically all BGP neighbors, even when a new neighbor is added subsequently; its IP address should be added automatically to the filter to accept traffic. For this purpose, you can use dynamic prefix-lists. Prefix-lists are a named list of IP addresses. The Junos apply-path feature transforms prefix-list in a dynamic prefix-list. This feature dynamically adds BGP peer information (or other information) from the configuration using wildcards into the prefix-list. Thus, if you add a BGP peer to the configuration, it is automatically added to the prefix list

```
[edit]
lab@R1# set policy-options prefix-list bgp-peers apply-path "protocols bgp group <*>
neighbor <*>"
```

Note: A `prefix-list` is configured under the `policy-options` hierarchy. They are also used for routing-policies. The wildcard “*” can be used to match "all" BGP groups and match "all" neighbor IP address.

© 2020 Juniper Networks, Inc. All rights reserved.

b) Next you can reference that prefix-list in the filter term which accepts BGP traffic only from the configured neighbors.

```
[edit]
lab@R1# set firewall family inet filter protect-re term bgp from source-prefix-list
bgp-peers

[edit]
lab@R1# set firewall family inet filter protect-re term bgp from protocol tcp

[edit]
lab@R1# set firewall family inet filter protect-re term bgp from port bgp

[edit]
lab@R1# set firewall family inet filter protect-re term bgp then accept

[edit]
lab@R1# commit
commit complete
```

Note: Adding a term to the existing filter inserts it at bottom of the filter. In other words, it will become the last entry to be evaluated. If you need to insert it higher into the list, you can use the command `insert` in the way shown below.

```
[edit]
lab@R1# insert firewall family inet filter protect-re term bgp before term igmp
```

3) Next you have to add terms for accepting traffic from different management and system protocols. The explicit source should be configured in the match condition.

```
[edit]
lab@R1# set firewall family inet filter protect-re term ntp from source-address
10.10.1.0/24

[edit]
lab@R1# set firewall family inet filter protect-re term ntp from protocol udp

[edit]
lab@R1# set firewall family inet filter protect-re term ntp from port ntp

[edit]
lab@R1# set firewall family inet filter protect-re term ntp then accept

[edit]
lab@R1# set firewall family inet filter protect-re term snmp from source-address
10.10.1.0/24
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set firewall family inet filter protect-re term snmp from protocol udp

[edit]
lab@R1# set firewall family inet filter protect-re term snmp from port snmp

[edit]
lab@R1# set firewall family inet filter protect-re term snmp then accept

[edit]
lab@R1# set firewall family inet filter protect-re term radius from source-address
10.10.1.0/24

[edit]
lab@R1# set firewall family inet filter protect-re term radius from protocol udp

[edit]
lab@R1# set firewall family inet filter protect-re term radius from port radius

[edit]
lab@R1# set firewall family inet filter protect-re term radius then accept

[edit]
lab@R1# set firewall family inet filter protect-re term dns from source-address
10.10.1.0/24

[edit]
lab@R1# set firewall family inet filter protect-re term dns from protocol udp

[edit]
lab@R1# set firewall family inet filter protect-re term dns from port domain

[edit]
lab@R1# set firewall family inet filter protect-re term dns then accept

[edit]
lab@R1# set firewall family inet filter protect-re term ssh from source-address
10.10.1.0/24

[edit]
lab@R1# set firewall family inet filter protect-re term ssh from protocol tcp

[edit]
lab@R1# set firewall family inet filter protect-re term ssh from port ssh

[edit]
lab@R1# set firewall family inet filter protect-re term ssh then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set firewall family inet filter protect-re term telnet from source-address
10.10.1.0/24

[edit]
lab@R1# set firewall family inet filter protect-re term telnet from protocol tcp

[edit]
lab@R1# set firewall family inet filter protect-re term telnet from port telnet

[edit]
lab@R1# set firewall family inet filter protect-re term telnet then accept

[edit]
lab@R1# set firewall family inet filter protect-re term ftp from source-address
10.10.1.0/24

[edit]
lab@R1# set firewall family inet filter protect-re term ftp from protocol tcp

[edit]
lab@R1# set firewall family inet filter protect-re term ftp from port [ ftp ftp-data ]

[edit]
lab@R1# set firewall family inet filter protect-re term ftp then accept

[edit]
lab@R1# commit
commit complete
```

4) Next you have to configure terms for ICMP messages and traceroute traffic. These terms should rate-limit the traffic.

   a) Traffic rate-limiting in Junos is accomplished with the use of policers. Policers are limiting the traffic to the configured bandwidth limit; exceeding traffic is dropped. The burst size limit enables the policer to accept traffic bursts for a limited time. The bandwidth limit is configured in bytes per second or percentages, whereas the burst size limit is configured in bytes. The actions applied to the policers include discard and reclassified. Below is the format of the policer.

```
firewall {
       policer $policer_name$ {
              if-exceeding {
                     bandwidth-limit $PIR$;
                     burst-size-limit $burst_size$;
                     }
              then $action$;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            }
        }
```

Now you can configure the policer based on the given requirements.

```
[edit]
lab@R1# set firewall policer re-policer if-exceeding bandwidth-limit 100k

[edit]
lab@R1# set firewall policer re-policer if-exceeding burst-size-limit 25k

[edit]
lab@R1# set firewall policer re-policer then discard
```

b) After the policer is configured it can be added to the terms matching ICMP and traceroute traffic.

```
[edit]
lab@R1# set firewall family inet filter protect-re term icmp from protocol icmp

[edit]
lab@R1# set firewall family inet filter protect-re term icmp then policer re-policer

[edit]
lab@R1# set firewall family inet filter protect-re term icmp then accept

[edit]
lab@R1# set firewall family inet filter protect-re term traceroute from protocol udp

[edit]
lab@R1# set firewall family inet filter protect-re term traceroute from port 33434-
33534

[edit]
lab@R1# set firewall family inet filter protect-re term traceroute then policer re-
policer

[edit]
lab@R1# set firewall family inet filter protect-re term traceroute then accept

[edit]
lab@R1# commit
commit complete
```

Note: The *traceroute* term is matching UDP based traceroute traffic. ICMP traceroute is matched by the `term icmp`.

© 2020 Juniper Networks, Inc. All rights reserved.

5) Next you have to configure the last term, which is an explicit deny. Although there is implicit deny term at the end of each filter, it is recommended to configure explicit deny rule.

This step requires that last explicit deny term should log all packets, count the packets and bytes discarded.

```
[edit]
lab@R1# set firewall family inet filter protect-re term last then count dropped-packets

[edit]
lab@R1# set firewall family inet filter protect-re term last then log

[edit]
lab@R1# set firewall family inet filter protect-re term last then discard

[edit]
lab@R1# commit
commit complete
```

Note: The last term does not have any match conditions. This way all traffic which is not evaluated before will match that term

6) Applying the firewall filter.

Firewall filters can be applied in the input or output direction. Applying a filter in the input direction matches ingress traffic through that interface, output direction matches egress traffic. The task requires filtering traffic to the routing engine, to accomplish this you can apply the filter to the loopback interface lo0 in the input direction. That interface has no configured IP address yet, but the filter applied to it will match all traffic destined to the RE coming from any of the interfaces.

```
[edit]
lab@R1# set interfaces lo0 unit 0 family inet filter input protect-re

[edit]
lab@R1# commit
commit complete
```

Note: When applying a filter to the RE it is best to use `commit confirm <min.>`. This way if you make mistake, the router will roll back the configuration after the specified time if a second `commit` is not issued.

© 2020 Juniper Networks, Inc. All rights reserved.

To verify policer operation, you can use ping with big packets to simulate increased icmp traffic. Ping from R2's management IP address to R1's management IP. Policer should rate-limit ping and traceroute traffic above 100kbit/s.

```
lab@R2> ping 10.10.1.1 size 1400 rapid count 100
PING 10.10.1.1 (10.10.1.1): 1400 data bytes
!!!!!!!!!!!!!!!!.!!!!.!!!!!.!!!!.!!!!!.!!!!.!!!!!.!!!!.!!!!.!!!!!.!!!!.!!!!!.!!!
!.!!!!!.!!!!.!!!!.!
--- 10.10.1.1 ping statistics ---
100 packets transmitted, 84 packets received, 16% packet loss
round-trip min/avg/max/stddev = 0.304/0.493/1.123/0.138 ms

lab@R2> ping 10.10.1.1 size 1400 rapid count 100
PING 10.10.1.1 (10.10.1.1): 1400 data bytes
!!!!!!!!!!!!!!!!.!!!!!.!!!!.!!!!.!!!!!.!!!!.!!!!!.!!!!.!!!!!.!!!!.!!!!!.!!!!.!!!
!.!!!!!.!!!!.!!!!!.
--- 10.10.1.1 ping statistics ---
100 packets transmitted, 84 packets received, 16% packet loss
round-trip min/avg/max/stddev = 0.295/0.499/0.882/0.116 msYou can see that R1
drops some of the ICMP packets.
```

On R1 you can see the increased counters of dropped packets for the policer.

```
[edit]
lab@R1# run show firewall filter protect-re

Filter: protect-re
Counters:
Name                                                Bytes              Packets
dropped-packets                                     18368                   56
Policers:
Name                                                Bytes              Packets
re-policer-icmp                                     22848                   16
re-policer-traceroute                                   0                    0
```

Verify the operation of the other terms in the firewall filter:

You can test the accept terms by initiating a telnet session from one of the other routers. It should be successful.

```
lab@R3> telnet 10.10.1.1
Trying 10.10.1.1...
Connected to 10.10.1.1.
Escape character is '^]'.

R1 (ttyp1)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
login:
```

To verify the last explicit discard term, you can deactivate the `icmp` term and issue ping from other routers to the management IP address. Ping should be unsuccessful. A log entry has to be generated and the firewall counter must increment.

```
[edit]
lab@R1# run show firewall filter counter dropped-packets protect-re

Filter: protect-re
Counters:
Name                                Bytes              Packets
dropped-packets                      9576                  114
```

```
[edit]
lab@R1# run show firewall log | match ICMP
11:37:23  protect-re D     fxp0.0         ICMP           10.10.1.100
10.10.1.1
11:37:23  protect-re D     fxp0.0         ICMP           10.10.1.100
10.10.1.1
11:37:12  protect-re D     fxp0.0         ICMP           10.10.1.100
10.10.1.1
11:37:12  protect-re D     fxp0.0         ICMP           10.10.1.100
10.10.1.1
11:37:12  protect-re D     fxp0.0         ICMP           10.10.1.100
10.10.1.1
11:37:01  protect-re D     fxp0.0         ICMP           10.10.1.2
10.10.1.1
11:37:00  protect-re D     fxp0.0         ICMP           10.10.1.2
10.10.1.1
11:37:00  protect-re D     fxp0.0         ICMP           10.10.1.2
10.10.1.1
11:36:59  protect-re D     fxp0.0         ICMP           10.10.1.2
10.10.1.1
11:36:59  protect-re D     fxp0.0         ICMP           10.10.1.2
10.10.1.1
11:36:58  protect-re D     fxp0.0         ICMP           10.10.1.2
10.10.1.1
11:36:58  protect-re D     fxp0.0         ICMP           10.10.1.2
10.10.1.1
11:36:57  protect-re D     fxp0.0         ICMP           10.10.1.2
10.10.1.1
11:36:55  protect-re D     fxp0.0         ICMP           10.10.1.100
10.10.1.1
```

Note: Do not forget to activate the `icmp` term after you verified the filter.

© 2020 Juniper Networks, Inc. All rights reserved.

Now that you are at the end of this step you can copy the configuration and load it on the other devices to save time.

Task 4: Interface Configuration

This task involves the configuration of the interfaces used in the workbook chapters. Figure 3 represents the logical connections after all device interfaces are configured according to Table 7.

Router	Interface Name	IP Address	IPv6 Address
R1	ge-0/0/1	802.3ad	
	ge-0/0/4.114	172.30.0.5/30	
	ge-0/0/4.118	172.30.0.9/30	link-local
	ge-0/0/2	802.3ad	
	ae0.0	172.30.0.1/30	link-local
	lo0.0	172.30.5.1/32	fd17:f0f4:f691:5::1/128
R2	ge-0/0/1	802.3ad	
	ge-0/0/4.127	172.30.0.17/30	
	ge-0/0/4.123	172.30.0.13/30	link-local
	ge-0/0/2	802.3ad	
	ae0.0	172.30.0.2/30	link-local
	lo0.0	172.30.5.2/32	fd17:f0f4:f691:5::2/128
R3	ge-0/0/4.134	172.30.0.21/30	link-local
	ge-0/0/4.136	172.30.0.25/30	
	ge-0/0/4.123	172.30.0.14/30	link-local
	ge-0/0/4.200	172.30.1.1/24	
	ge-0/0/4.201	172.30.2.1/24	
	lo0.0	172.30.5.3/32	fd17:f0f4:f691:5::3/128
R4	ge-0/0/4.134	172.30.0.22/30	link-local
	ge-0/0/4.114	172.30.0.6/30	
	ge-0/0/4.145	172.30.0.29/30	link-local
	ge-0/0/4.200	172.30.1.2/24	
	ge-0/0/4.201	172.30.2.2/24	
	lo0.0	172.30.5.4/32	fd17:f0f4:f691:5::4/128
R5	ge-0/0/1	802.3ad	

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/4.158	172.30.0.37/30	
	ge-0/0/4.145	172.30.0.30/30	link-local
	ge-0/0/2	802.3ad	
	ae0.0	172.30.0.33/30	link-local
	lo0.0	172.30.5.5/32	fd17:f0f4:f691:5::5/128
R6	ge-0/0/1	802.3ad	
	ge-0/0/4.136	172.30.0.26/30	
	ge-0/0/4.167	172.30.0.41/30	link-local
	ge-0/0/2	802.3ad	
	ae0.0	172.30.0.34/30	link-local
	lo0.0	172.30.5.6/32	fd17:f0f4:f691:5::6/128
R7	ge-0/0/4.178	172.30.0.45/30	link-local
	ge-0/0/4.127	172.30.0.18/30	
	ge-0/0/4.167	172.30.0.42/30	link-local
	lo0.0	172.30.5.7/32	fd17:f0f4:f691:5::7/128
R8	ge-0/0/4.178	172.30.0.46/30	link-local
	ge-0/0/4.158	172.30.0.38/30	
	ge-0/0/4.118	172.30.0.10/30	link-local
	lo0.0	172.30.5.8/32	fd17:f0f4:f691:5::8/128

Table 7

Junos Interfaces

The main function of Junos based devices is to forward traffic from one interface to another. Many different types of interfaces are supported on Juniper Network devices. In this lab only Gigabit Ethernet (ge-) interfaces are used where ge- represents the media type of the interface. All media types are shortened to two letters except the fxp interface. As stated before the fxp interface is dedicated for the management of the device and is not used for transit traffic. The Junos interface naming convention uses the following format:

```
media_type-fpc/pic/port.unit
```

For example, **`ge-0/0/0.0`** – Gigabit Ethernet interface on *FPC 0/PIC 0/Port 0.logical interface 0*

Note: The interfaces on the device are not listed in the factory default configuration if they don't have any configuration. You can list all available interfaces with the operational command `show interface terse`.

© 2020 Juniper Networks, Inc. All rights reserved.

Every interface has physical and logical properties. Physical properties define the physical parameters of the interface like *encapsulation, physical maximum transmit unit (MTU), VLAN-tagging* and others.

Every interface in Junos requires at least one logical *unit* where the protocol and address information are configured. Ethernet interface can have multiple logical units after the `vlan-tagging` capability is enabled for that interface. In that case each logical `unit` is assigned different VLAN number. Loopback interface can have just one logical unit in particular routing instance. Logical units can have any number between 0 and 16384.

Logical interfaces are capable to support one or more *protocol families.* Enabling those families allows the interface to process data packet of different protocols. Some of the protocol families are `inet`, `inet6`, `iso`, `mpls`. Protocol addresses are the Layer 3 addresses used to route user traffic in a network. Protocol addresses are configured under the respective protocol families.

Junos Aggregated Interfaces (LAG)

IEEE 802.3ad link aggregation is used to form a single link layer interface from a group of physical Ethernet interfaces, known as a Link Aggregation Group (LAG). Aggregated interfaces create a single logical point-to-point link between two devices. Traffic is balanced over the physical member links within the LAG and increases bandwidth. High availably is also increased; if one of the links fails the LAG continues to carry traffic over the remaining links. All interfaces participating in the same LAG must have the same speed.

By default, Ethernet links in a LAG do not exchange information about the state of the links. IEEE802.3ad defines the use of the Link Aggregation Control Protocol (LACP) for monitoring the links within a LAG group.

1) Configure of the interfaces of router R1

```
[edit]
lab@R1# set chassis aggregated-devices ethernet device-count 1
```

Note: By default, no aggregated interfaces exist, they have to be created. Under the `chassis` hierarchy you can define the number of LAG interfaces available on the device. The aggregate-devices devices are counted from 0. Thus, if you specify a device count of 1, only 1 LAG interface can be created and must be named *ae0* to become operational. Configure a *device-count of 6* enables LAG interfaces from *ae0* to *ae5.*

Some physical parameters like link speed and duplex mode can be configured on the LAG member interfaces. Some other parameters like `vlan-tagging`, are not supported on the

© 2020 Juniper Networks, Inc. All rights reserved.

LAG member interfaces. The `vlan-tagging` option is supported on aggregated interfaces, but it has to be configured under the `ae0` group interface.

```
[edit]
lab@R1# set interfaces ge-0/0/1 gigether-options 802.3ad ae0

[edit]
lab@R1# set interfaces ge-0/0/2 gigether-options 802.3ad ae0
```

LACP is enabled on the LAG to monitor the aggregated interface. R1 will actively send LACP packets. R2 can participate passively or actively in the LACP process. At least one of the two devices should be configured in the active state. The link will not come up if both sides are in passive LACP mode.

```
 [edit]
lab@R1# set interfaces ae0 aggregated-ether-options lacp active

[edit]
lab@R1# set interfaces ae0 unit 0 description "R2 connection"

[edit]
lab@R1# set interfaces ae0 unit 0 family inet address 172.30.0.1/30

[edit]
lab@R1# set interfaces ae0 unit 0 family inet6

[edit]
lab@R1# commit
commit complete
```

Ensure that you also configure aggregated interfaces on R2, R5, and R6.

Family `inet6` enables the interface to process IPv6 packets, but the IPv6 address is not explicitly configured. Enabling `family inet6` under the interface automatically configures the *link-local* IPv6 address on the interface. Link-local addresses are unique on each single link and generated from the fe80:0:0:0::/64 subnet prefix. Link-local addresses are never routed. The Interface ID is the second half of the address and is generated based on the MAC address of the physical interface. This way it is guaranteed that it is unique on the link. A MAC address is 48 bits long, whereas the Interface ID of the IPv6 address is 64 bits according to the IEEE EUI-64 format. 16 bits – FF:FE are added between the first 24 bits of the MAC address and the last 24 bits.

```
[edit]
lab@R1# run show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 501
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, BPDU Error:
None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled,
Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Current address: 00:26:88:1d:38:40, Hardware address: 00:26:88:1d:38:40
  Last flapped   : 2013-06-26 08:37:25 UTC (00:08:37 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 984 bps (0 pps)

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 528)
    Flags: Hardware-Down Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Security: Zone: Null
    Protocol inet, MTU: 1500
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 172.30.0.0/30, Local: 172.30.0.1, Broadcast:
172.30.0.3
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Dest-route-down Is-Preferred
        Destination: fe80::/64, Local: fe80::226:88ff:fe1d:3840
        INET6 Address Flags: Tentative
```

Note: The second lowest significant bit from the first octet is inverted to 1. This Universal/Local bit is of no importance to the network devices, a high value of 1 means that this is a Universal unique address.

© 2020 Juniper Networks, Inc. All rights reserved.

Display the operation of the LACP LAG interface.

```
[edit]
lab@R1# run show lacp interfaces
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast   Passive
      ge-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast   Passive
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-0/0/1                  Current   Fast periodic Collecting distributing
      ge-0/0/2                  Current   Fast periodic Collecting distributing
```

Next you must configure the interface terminating two VLANs by creating different logical interfaces for each VLAN.

Ensure you also configure all the other devices with their specific interface configuration.

Note: For interface configurations, using the `load merge terminal` command in not recommended for configuring the other routers. It is better to configure each device's interfaces separately because many parameters are different.

```
[edit]
lab@R1# set interfaces ge-0/0/4 vlan-tagging

[edit]
lab@R1# set interfaces ge-0/0/4 unit 114 description "R4 connection"

[edit]
lab@R1# set interfaces ge-0/0/4 unit 114 vlan-id 114

[edit]
lab@R1# set interfaces ge-0/0/4 unit 114 family inet address 172.30.0.5/30

[edit]
lab@R1# set interfaces ge-0/0/4 unit 118 description "R8 connection"

[edit]
lab@R1# set interfaces ge-0/0/4 unit 118 vlan-id 118

[edit]
lab@R1# set interfaces ge-0/0/4 unit 118 family inet address 172.30.0.9/30

[edit]
lab@R1# set interfaces ge-0/0/4 unit 118 family inet6
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: `vlan-tagging` enables the Ethernet interface to process VLAN tagged traffic. With the `vlan-id` parameter, you can specify the VLAN tag for each logical interface.

The last interface that you have to configure on R1 is the loopback interface.

```
[edit]
lab@R1# set interfaces lo0 unit 0 family inet address 172.30.5.1/32

[edit]
lab@R1# set interfaces lo0 unit 0 family inet6 address fd17:f0f4:f691:5::1/128

[edit]
lab@R1# commit
commit completed
```

Note: The IPv6 address is explicitly configured and is part from the Globally unique address space. Because every IPv6 enabled interface must have a link-local address, the lo0 interface has two IPv6 addresses as shown below.

```
[edit]
lab@R1# run show interfaces lo0.0
  Logical interface lo0.0 (Index 66) (SNMP ifIndex 16)
    Flags: SNMP-Traps Encapsulation: Unspecified
    Input packets : 0
    Output packets: 0
    Security: Zone: Null
    Protocol inet, MTU: Unlimited
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Default Is-Primary
        Local: 172.30.5.1
    Protocol inet6, MTU: Unlimited
      Flags: None
      Addresses, Flags: Is-Default Is-Primary
        Local: fd17:f0f4:f691:5::1
        Local: fe80::226:880f:fc12:ce00
```

2) Configure and testing of VRRP on the interfaces of R3 and R4

The Virtual Router Redundancy Protocol (VRRP) is an industry-standardized protocol that provides gateway redundancy in LAN environments. VRRP elects one router as the virtual gateway for the segment while the other VRRP routers are in the backup state. All communication between the VRRP routers is through a common Layer 2 switch.

The VRRP router elected as virtual gateway is the master router on the LAN segment. The master router is the router that is actually forwarding the traffic from the clients towards other destinations. The backup router becomes the master VRRP router if a failure occurs

© 2020 Juniper Networks, Inc. All rights reserved.

and it cannot communicate anymore with the VRRP master router. A Virtual IP (VIP) address on the master router serves as the default gateway IP address.

The VRRP protocol uses heartbeat messages to monitor the availability of the master router. The default value for the advertisements is 1 second. It can be configured between 1 and 255 seconds. VRRP supports MD5 authentication.

VRRP election is based on priorities. The priority values are in the range from 1 to 255, the higher priority is preferred. By default, the priority on Junos devices is set to 100. Two routers which have equal priorities use the highest IP address to determine who becomes the master VRRP router.

VRRP can dynamically adjust the priority value by monitoring (tracking) interface states. Tracking the state of the WAN (uplink) interfaces of the LAN segment can trigger a master VRRP router switch. You can configure the priority value to be reduced from the master router if one or several interfaces become unavailable. This way you can assure optimal path selection through the network in case of link failures.

VRRP must be configured on two interfaces on R3 and R4.

a) R3.

```
[edit]
lab@R3# set interfaces ge-0/0/4 vlan-tagging

[edit]
lab@R3# set interfaces ge-0/0/4 unit 200 description "DC1 LAN 1"

[edit]
lab@R3# set interfaces ge-0/0/4 unit 200 vlan-id 200

[edit]
lab@R3# set interfaces ge-0/0/4 unit 200 family inet address 172.30.1.1/24 vrrp-group 1
virtual-address 172.30.1.254

[edit]
lab@R3# set interfaces ge-0/0/4 unit 200 family inet address 172.30.1.1/24 vrrp-group 1
priority 150

[edit]
lab@R3# set interfaces ge-0/0/4 unit 200 family inet address 172.30.1.1/24 vrrp-group 1
authentication-type md5

[edit]
lab@R3# set interfaces ge-0/0/4 unit 200 family inet address 172.30.1.1/24 vrrp-group 1
authentication-key workbook
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set interfaces ge-0/0/4 unit 200 family inet address 172.30.1.1/24 vrrp-group 1
track interface ge-0/0/4.123 priority-cost 30

[edit]
lab@R3# set interfaces ge-0/0/4 unit 200 family inet address 172.30.1.1/24 vrrp-group 1
track interface ge-0/0/4.136 priority-cost 30
```

Note: The unit 200 has a priority of 150, which is higher than the default. The `vrrp-group` 1 has 172.30.1.254 as VIP assigned. VRRP in that group is tracking interfaces ge-0/0/4.127 and ge-0/0/4.123. Both interfaces will reduce the priority with 30 in case of a failure. This means that only both interfaces must be unavailable for the priority to be lower than the default value of 100 advertised by R4. In that case the VRRP master router will switch to R4.

```
[edit]
lab@R3# set interfaces ge-0/0/4 unit 201 description "DC1 LAN 2"

[edit]
lab@R3# set interfaces ge-0/0/4 unit 201 vlan-id 201

[edit]
lab@R3# set interfaces ge-0/0/4 unit 201 family inet address 172.30.2.1/24 vrrp-group 2
virtual-address 172.30.2.254

[edit]
lab@R3# set interfaces ge-0/0/4 unit 201 family inet address 172.30.2.1/24 vrrp-group 2
authentication-type md5

[edit]
lab@R3# set interfaces ge-0/0/4 unit 201 family inet address 172.30.2.1/24 vrrp-group 2
authentication-key workbook

[edit]
lab@R3# commit
commit complete
```

Note: unit 201 has default priority of 100. The VIP IP address is 172.30.2.254 for `vrrp-group 2`

Note: md5 authentication has to be used for advertisement exchange. Other authentication options are *none* and *simple password*.

b) R4.

```
[edit]
lab@R4# set interfaces ge-0/0/4 vlan-tagging
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set interfaces ge-0/0/4 unit 200 description "DC1 LAN 1"

[edit]
lab@R4# set interfaces ge-0/0/4 unit 200 vlan-id 200

[edit]
lab@R4# set interfaces ge-0/0/4 unit 200 family inet address 172.30.1.2/24 vrrp-group 1
virtual-address 172.30.1.254

[edit]
lab@R4# set interfaces ge-0/0/4 unit 200 family inet address 172.30.1.2/24 vrrp-group 1
authentication-type md5

[edit]
lab@R4# set interfaces ge-0/0/4 unit 200 family inet address 172.30.1.2/24 vrrp-group 1
authentication-key workbook
```

Note: Unit 200 has default priority of 100. IP address 172.30.1.254 is VIP for *vrrp-group 1*

```
[edit]
lab@R4# set interfaces ge-0/0/4 unit 201 description "DC1 LAN 2"

[edit]
lab@R4# set interfaces ge-0/0/4 unit 201 vlan-id 201

[edit]
lab@R4# set interfaces ge-0/0/4 unit 201 family inet address 172.30.2.2/24 vrrp-group 2
virtual-address 172.30.2.254

[edit]
lab@R4# set interfaces ge-0/0/4 unit 201 family inet address 172.30.2.2/24 vrrp-group 2
priority 150

[edit]
lab@R4# set interfaces ge-0/0/4 unit 201 family inet address 172.30.2.2/24 vrrp-group 2
authentication-type md5

[edit]
lab@R4# set interfaces ge-0/0/4 unit 201 family inet address 172.30.2.2/24 vrrp-group 2
authentication-key workbook


[edit]
lab@R4# set interfaces ge-0/0/4 unit 201 family inet address 172.30.2.2/24 vrrp-group 2
track interface ge-0/0/4.114 priority-cost 30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set interfaces ge-0/0/4 unit 201 family inet address 172.30.2.2/24 vrrp-group 2
track interface ge-0/0/4.145 priority-cost 30

[edit]
lab@R4# commit
commit complete
```

Note: Unit 201 has priority 150, which is higher than the default. *vrrp-group 2* has 172.30.2.254 as VIP assigned. For this group VRRP tracks the interfaces ge-0/0/4.114 and ge-0/0/4.145. Again, both interfaces will reduce the priority of that group with 30 in case of a failure. In that case the VRRP master router will switch to R3.

Test VRRP functionality:

a) Checking the VRRP status when all interfaces are up and available:

```
lab@R3> show vrrp
Interface     State       Group   VR state VR Mode   Timer    Type   Address
ge-0/0/4.200  up              1   master   Active      A  0.324 lcl    172.30.1.1
                                                                vip
172.30.1.254
ge-0/0/4.201  up              2   backup   Active      D  3.407 lcl    172.30.2.1
                                                                vip
172.30.2.254
                                                                mas    172.30.2.2

lab@R3> show vrrp track
Track Int   State         Speed   VRRP Int   Group   VR State      Current prio
ge-0/0/4.123 up             1g   ge-0/0/4.200     1 master                150
ge-0/0/4.136 up             1g   ge-0/0/4.200     1 master                150
```

```
lab@R4> show vrrp
Interface     State       Group   VR state VR Mode   Timer    Type   Address
ge-0/0/4.200  up              1   backup   Active      D  2.904 lcl    172.30.1.2
                                                                vip
172.30.1.254
                                                                mas    172.30.1.1
ge-0/0/4.201  up              2   master   Active      A  0.519 lcl    172.30.2.2
                                                                vip
172.30.2.254

lab@R4> show vrrp track
Track Int   State         Speed   VRRP Int   Group   VR State      Current prio
ge-0/0/4.114 up             1g   ge-0/0/4.201     2 master                150
ge-0/0/4.145 up             1g   ge-0/0/4.201     2 master                150
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) Checking the master switchover and tracking functionality:

To verify the master switchover and tracking functionality you can deactivate the interfaces configured for tracking in *vrrp-group 1* on router R3

```
[edit]
lab@R3# deactivate interfaces ge-0/0/4.123

[edit]
lab@R3# deactivate interfaces ge-0/0/4.136

[edit]
lab@R3# commit
commit complete
```

On R3 both VRRP groups are in backup state:

```
[edit]
lab@R3# run show vrrp
Interface     State       Group   VR state VR Mode   Timer    Type   Address
ge-0/0/4.200  up              1   backup   Active      D  3.289 lcl    172.30.1.1
                                                               vip
172.30.1.254
                                                               mas    172.30.1.2
ge-0/0/4.201  up              2   backup   Active      D  3.282 lcl    172.30.2.1
                                                               vip
172.30.2.254
                                                               mas    172.30.2.2
```

```
[edit]
lab@R3#
lab@R3# run show vrrp track
Track Int   State         Speed   VRRP Int   Group   VR State      Current prio
ge-0/0/4.123 unknown          0   ge-0/0/4.200     1 backup                 90
ge-0/0/4.136 unknown          0   ge-0/0/4.200     1 backup                 90
```

On R4 both VRRP groups are in master state:

```
lab@R4> show vrrp
Interface     State       Group   VR state VR Mode   Timer    Type   Address
ge-0/0/4.200  up              1   master   Active      A  0.236 lcl    172.30.1.2
                                                               vip
172.30.1.254
ge-0/0/4.201  up              2   master   Active      A  0.783 lcl    172.30.2.2
                                                               vip
172.30.2.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: Remember to activate both interfaces to roll back to the requested configuration.

Task 5: Scripting

The last task for the first part requires demonstrating basic knowledge of working with Junos scripts.

Junos Scripting.

Junos scripting is used for automating different Junos tasks and functions. The scripts are simplifying complex tasks, automate manual procedures and help with management and troubleshooting.

Junos supports three types of scripts:

- *Operation (OP)* scripts are used to customize Junos commands. They can execute multiple `show` commands in one shot, process the data and display only the desired information. Other functions of OP scripts are to automate configuration changes in a controlled way.
- *Commit* scripts are executed each time when a commit is performed. The script verifies the candidate configuration and fails the commit operation if errors are found. Usually they are used for mitigating human errors when configuring the device.
- *Event* scripts are trigged automatically by event(s) occurring on the device. Events can be syslog message, SNMP traps and others. Event scripts can collect valuable troubleshooting information or automatically change the configuration.

You can use existing scripts located on the management server S1 for this Task.

1) Download the OP script.

```
lab@R1> file copy ftp://lab:lab123@10.10.1.100/show-interfaces.slax
/var/db/scripts/op/show-interfaces.slax
```

Note: You must be logged in with *root* user to copy files to /var/db/scripts folder.

2) Download the commit script.

```
lab@R1> file copy ftp://lab:lab123@10.10.1.100/interface-mask-check.slax
/var/db/scripts/commit/interface-mask-check.slax
```

3) Download the event script.

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1> file copy ftp://lab:lab123@10.10.1.100/syslog-int-desc-on-link-change.slax
/var/db/scripts/event/syslog-int-desc-on-link-change.slax
```

4) Enable and verify the script operation.

a) Enable operation script.

```
[edit]
lab@R1# set system scripts op file show-interfaces.slax

[edit]
lab@R1# commit
commit complete
```

Verify the script:

This operation script modifies the output of the command *show interface terse*. The descriptions of the interfaces are added to the output.

```
lab@R1# run op show-interfaces
Interface               Admin Link Proto    Local                 Remote
lc-0/0/0.32769                     vpls
pfe-0/0/0.16383                    inet
                                   inet6
pfh-0/0/0.16383                    inet
pfh-0/0/0.16384                    inet
ge-0/0/1.0                         aenet    --> ae0.0
ge-0/0/2.0                         aenet    --> ae0.0
ge-0/0/4.114            R4 connection
                                   inet     172.30.0.5           --> 172.30.0.4/30
                                   multiservice
ge-0/0/4.118            R8 connection
                                   inet     172.30.0.9           --> 172.30.0.8/30
                                   inet6    fe80::250:5600:76a9:278c--> fe80::/64
                                   multiservice
ge-0/0/4.32767                     multiservice
ae0.0                   R2 connection
                                   inet     172.30.0.1           --> 172.30.0.0/30
                                   inet6    fe80::2e6b:f5ff:fe86:c6c0-->
fe80::/64
                                   multiservice
em1.0                              inet     10.0.0.4             --> 10/8
                                            128.0.0.1            --> 128/2
...
```

b) Enable the commit script.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set system scripts commit file interface-mask-check.slax

[edit]
lab@R1# commit
commit complete
```

Verify the script:

This commit script displays a warning during a commit if the subnet mask of any configured interface is smaller than 24.

```
[edit]
lab@R1# set interfaces ge-0/0/4.115 family inet address 172.30.4.1/20

[edit]
lab@R1# commit check
warning: The address of 172.30.4.1 has a mask of /20
on interface ge-0/0/4 unit 115
[edit interfaces ge-0/0/4]
  'unit 115'
     VLAN-ID must be specified on tagged ethernet interfaces
error: configuration check-out failed

[edit]
lab@R1# rollback
load complete
```

c) Enable the event script.

Event scripts are triggered by event policies. They are not self-triggered. Event policies instruct the event daemon to perform specific actions in response to system events. One of such actions is executing an event script. Event policies are very powerful tools capable of performing different tasks based on events. A minimum configuration of the event policy is shown below:

```
event-options {
       policy example {
              events ui_commit;
              then {
                     event-script example-script.slax;
              }
       }
}
```

Event scripts should be enabled in the configuration.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set event-options event-script file syslog-int-desc-on-link-change.slax
```

Then event policy has to be created. Please the following example of an event-policy.

```
[edit]
lab@R1# run file show /var/db/scripts/event/syslog-int-desc-on-link-change.slax
*
*
* This event script will generate a new syslog message based on the triggering
* syslog message of SNMP_TRAP_LINK_DOWN or SNMP_TRAP_LINK_UP. It collects the
* related interface information from the syslog message and also grabs the
* interface description to form a new syslog message.
*
*
* To invoke this event script, place the syslog-interface-description-on-
* link-change.slax file in /var/db/scripts/event/ and enter the following
* into the device config.
* The second policy is to also create a trap on the newly created syslog
* message.
*
* ----Begin config snippet----
*
* root@JUNIPER_DEVICE# show event-options
* policy sylog_if_description {
*     events [ snmp_trap_link_up snmp_trap_link_down ];
*     then {
*         event-script syslog-int-desc-on-link-change.slax;
*     }
* }
* policy snmptrap_if_description {
*     events SYSTEM;
*     attributes-match {
*         SYSTEM.message matches NEW_SNMP_TRAP_LINK;
*     }
*     then {
*         raise-trap;
*     }
* }
* event-script {
*     file syslog-int-desc-on-link-change.slax;
* }
*
* ----End config snippet----
*
*
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: The operational command `file show` <file name> displays the content of the file.

Configure the event-policy.

```
[edit]
lab@R1# set event-options policy sylog_if_description events snmp_trap_link_up

[edit]
lab@R1# set event-options policy sylog_if_description events snmp_trap_link_down

[edit]
lab@R1# set event-options policy sylog_if_description then event-script syslog-int-
desc-on-link-change.slax

[edit]
lab@R1# set event-options policy snmptrap_if_description events SYSTEM

[edit]
lab@R1# set event-options policy snmptrap_if_description attributes-match
SYSTEM.message matches NEW_SNMP_TRAP_LINK

[edit]
lab@R1# set event-options policy snmptrap_if_description then raise-trap

[edit]
lab@R1# commit
commit complete
```

Note: The first policy sylog_if_description is executing the event-script when an interface goes UP or DOWN. The result of the script generates a syslog message, which includes the description of the interface. The second policy snmptrap_if_description generates a new SNMP trap message based on that new syslog message, again the description of the interface is included.

Verify the script and policy:

You can deactivate some of the interfaces to generate syslog message matched by the policy.

Note: Remember to reactivate the interface after testing.

```
[edit]
lab@R1# set interfaces ae0 disable

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

You may see a slightly different output from the script. The important part is that you see an entry for “R1 cscript”.

```
lab@R1> show log jncie-sp-messages
Jun 20 20:31:22  R1 mgd[3470]: UI_CHILD_START: Starting child '/sbin/ifinfo'
Jun 20 20:31:22  R1 mgd[3470]: UI_CHILD_STATUS: Cleanup child '/sbin/ifinfo', PID
3479, status 0
Jun 20 20:31:22  R1 mgd[3471]: UI_CHILD_STATUS: Cleanup child '/sbin/ifinfo', PID
3478, status 0
Jun 20 20:31:22  R1 mgd[3471]: UI_LOGOUT_EVENT: User 'root' logout
Jun 20 20:31:22  R1 cscript: NEW_SNMP_TRAP_LINK_DOWN, R1, , , ,
Jun 20 20:31:22  R1 mgd[3470]: UI_LOGOUT_EVENT: User 'root' logout
Jun 20 20:31:22  R1 cscript: NEW_SNMP_TRAP_LINK_DOWN, R1, ae0.0, up, down, R2
connection
Jun 20 20:32:14  R1 mgd[1934]: UI_CMDLINE_READ_LINE: User 'lab', command 'exit '
Jun 20 20:32:14  R1 mgd[1934]: UI_DBASE_LOGOUT_EVENT: User 'lab' exiting
configuration mode
Jun 20 20:32:37  R1 mgd[1934]: UI_CMDLINE_READ_LINE: User 'lab', command 'show
log jncie- '
```

Scripts have to be downloaded to the other devices and configured.

5) Last step is to enable the ae0 interface and save the configuration. If you do not want to save your own configuration, there is a saved copy on each device named `ch1-end.config` that can be used for subsequent labs.

```
[edit]
lab@R1# delete interfaces ae0 disable

[edit]
lab@R1# commit
commit complete [edit]

lab@R1# save my_baseline.config
```


© 2020 Juniper Networks, Inc. All rights reserved.

Appendix - Chapter 2: IGP Configuration and Troubleshooting

This chapter has three different tasks, which involve the configuration and troubleshooting of IGP routing protocols. IGPs are *interior gateway protocols* used to provide dynamic routing within autonomous system (AS). Several IGP routing protocols are supported by Junos OS. Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS) and Routing Information Base (RIP) are three IGP protocols able of building dynamic routing paths through the network. The first two protocols are *link-state* routing protocols capable of building and maintaining neighbor adjacencies, exchanging link-state information with their adjacent routers, creating databases containing those advertisements, and use the Dijkstra algorithm to build optimal loop free paths through the network based on the link-state database. RIP is a *distance vector* routing protocol using a hop by hop routing information exchange and simple calculations (hop count) to find the best path through the network. Distance vector protocols converge slowly and are vulnerable to routing loops. The JNCIE-SP lab mainly focuses on the configuration and troubleshooting of link-state routing protocols, but basic RIP operation and configuration knowledge is also required. Static routes provide a manual way of changing the routing information on a single router. The use of static routes during JNCIE-SP lab tasks are usually not allowed. Never configure a static route unless you are told to do so, or else you will likely loose points on your exam (if in doubt, always check with your exam proctor).

Task 1: OSPF Troubleshooting

OSPF Overview

OSPF exchanges information using Link State Advertisements (LSA), which are flooded through the network. The data described by the LSAs is used to create and maintain the link state database (LSDB). Shortest path first (SPF), which is based on the Dijkstra algorithm and used to calculate the best path for each prefix based on the information into the LSDB.

All routers have the same path information in the LSDB. LSAs are flooded through the network from router to router. This way all routers in one area have the same LSAs stored in the LSDB. Split horizon is used to prevent sending an LSA received from a neighbor back to the same neighbor. If there is a need to limit the flooding of LSAs and SFP calculations, OSPF has the ability to create individual Areas. This allows splitting the routing domain to different smaller SPF domains. This way flooding can be constrained within an Area. Having many different areas with SPF algorithm calculations without visibility to other areas requires the exchange of inter-area routing information. This role is handled by Area Border Routers (ABR), which have interfaces participating in two or more areas. This is done in a distance vector fashion. Routers in one area do not have visibility

© 2020 Juniper Networks, Inc. All rights reserved.

into other areas links, so they trust the information send from the ABR. An ABR is like a gateway to the routes in other areas. To protect against inter area routing loops, all inter area traffic should pass the *backbone area*, which is Area 0.

OSPF routers can form and maintain adjacencies with directly connected neighbors. OSPF Hello messages are used to discover, negotiate adjacencies and serve as keepalives. During the discovery of OSPF adjacencies, some attributes in the hello message must be the same between neighbors. These are: Area ID, Authentication, on broadcast networks the network mask, Hello interval, Dead interval, and some option bits like External bit. Also, the MTU on both ends of the link should match. This last check is performed before the database synchronization phase starts to define the maximum size of database packets possible without the need for fragmentation. The OSPF adjacency states are:

- **Down** – this is the initial state.
- **Init** – The local router sent a hello packet, but the communication is still not established.
- **2way** – bidirectional communication is started. Local router received a Hello packet with its own router ID listed.
- **ExStart** – start of the database synchronization process. During this state the routers negotiate which router is the master of this process.
- **Exchange** – LSA exchange between the routers.
- **Loading** – The local router finished sending LSA information but is still receiving database information.
- **Full** – this state indicates that the databases on both neighbors are synchronized and adjacency is formed.

OSPF uses different types of LSAs to exchange routing information. The scope of flooding of each LSA depends on its type and the function of the router sending it.

The following LSAs types are used in typical OSPF network:

- **Router (Type 1)** – this LSA is flooded only within the area. Router LSA describes the links and neighbors of the local router.
- **Network (Type 2)** – Type 2 LSA is used only for broadcast media (Ethernet by default) and is originated by the Designated Router (DR) of the segment. It is used to represent the segment with all attached routers. The flooding scope is within the area only.
- **Summary (Type 3)** – summary LSA is used from the ABR to exchange information from LSA type 1 and LSA type 2 to other attached areas. Its scope is whole OSPF domain if not limited by the area design.

© 2020 Juniper Networks, Inc. All rights reserved.

- **ASBR Summary (Type 4)** – Autonomous System Border Routers are routers that inject external routing information to the OSPF domain. ASBR Summary LSA describe the router ID of ASBR routers. This LSA is generated by an ABR attached to the area in which the ASBR is located to describe the reachability to this ASBR router-id. Its scope is the whole OSPF domain.
- **External (Type 5)** – this LSA describes IP routes redistributed from other routing protocol including static routes. These LSAs are generated by ASBRs. They are flooded to all areas and OSPF routers except the areas, which by design do not need them. By default, the metric type used by Junos OS is type 2, which includes only the external cost of the route and ignoring internal cost to reach the ASBR router. Type 1 consists of external cost plus the internal cost to the ASBR, which can be configured.
- **NSSA External (Type 7)** – NSSA External LSA is similar to External LSA Type 5 as it is representing IP prefixes external to OSPF domain. The difference is that these LSAs are send by ASBRs inside a NSSA area. Type 7 LSAs are translated to an LSA Type 5 by the ABR attached to the NSSA area from where the LSA Type 7 is generated.
- **Opaque (Type 10)** – This LSA has an area wide scope and is used for exchanging additional information about the links used for traffic-engineering purposes in MPLS networks. The information from that LSA is used to populate the Traffic Engineering Database (TED), used for the Constrained Shortest Path First (CSPF) algorithm.

OSPF uses a cost value as the metric on each link. The cost is calculated based upon the interface bandwidth and the OSPF reference bandwidth. The cost is the result of the reference bandwidth divided by the interface bandwidth. The default OSPF reference bandwidth is 100Mbps which means that an interface with speed 100Mbsp or faster will have cost of 1 by default. The cost of each interface can also be configured manually.

As stated above, OSPF uses the concept of areas to divide the routing domain in smaller areas, which limit the flooding of LSAs. Several different types of areas are available to address different topology designs. The backbone area (area 0) has a special purpose, as it ensures a loop-free network when interconnecting other areas. Non-backbone areas are connected to the OSPF backbone to reach prefixes in other areas. Two special non-backbone areas are available for optimal LSAs flooding:

- **Stub area** – this area has no ASBR inside and all packets to and from external prefixes pass through the ABR(s). Suppressing Type 5 LSAs and injecting a default route from the ABR is enough for reaching all external prefixes. Type 5 LSAs are not allowed in stub areas. Having no specific information for external routes, internal routers inside stub area do not need to know about ASBRs, so

© 2020 Juniper Networks, Inc. All rights reserved.

LSAs Type 4 are not flooded as well. A default route is generated by the ABR as LSA Type 3. You can further limit the number of LSAs inside the stub area by making the area totally stub. Having only a default route outside the area through the ABR, which will stop flooding of other LSA Type 3 into the area except the one caring the default route.

- **NSSA** – because stub areas are not flexible, you cannot selectively allow some Type 5 LSAs to be filtered. When an ASBR (external routes) are needed inside stub areas, a new type of area must be used – Not So Stubby Area (NSSA). Like stub areas, Type 5 LSA's are not allowed. For this reason, type 7 LSA are generated by ASBRs in NSSA areas. Type 7 LSAs are re-advertised as LSA Type 5 by the ABR when send to other areas. As with stub areas Type 5 LSAs are not allowed, default route is injected from the ABR as Type 3 LSA or Type 7.

OSPF requires a unique identifier for every router within the routing domain. The OSPF Router ID is a 32 bit dotted number. It is represented the same way as an IP address. It is recommended to manually configure the OSPF Router ID, which is often the same as the routers IPV4 loopback IP address. **Remember that it must be unique within the network, otherwise unexpected behavior could be observed.**

OSPF classifies all interfaces into one of the following network types:

- Broadcast networks
- Point-to-point networks
- Nonbroadcast multi-access (NBMA) networks
- Point-to-multipoint networks
- Virtual links

Note: Virtual links will be covered later during the troubleshooting steps.

The first two types are mostly used in today's network. Broadcast multi access network is the default type for Ethernet interfaces, where more than two devices can communicate on one segment. Multicast packets are used for exchanging OSPF messages using the well-known multicast addresses 224.0.0.5 and 224.0.0.6. All OSPF routers listen to the 224.0.0.5 address, the 224.0.0.6 address is used for Designated Router (DR) and Backup Designated Router (BDR) purposes. The DR and BDR routers are elected in multi-access segments to reduce the OSPF traffic exchange and SFP calculations. DR routers have a priority value, which can be a number between 0 and 255. It can be manually configured, by default Junos assigns the value of 128. If two routers share the same priority value the highest Router ID is used as a tiebreaker. Every router in the same broadcast segment forms *Full* adjacency with DR and BDR only. Non DR and BDR routers will not form a full OSPF adjacency between each other (2way). DR and BDR are the only neighbors, which are able to send OSPF type 2 LSAs in the OSPF domain. They represent the multi access

© 2020 Juniper Networks, Inc. All rights reserved.

segment. By default, keep-alive/hello messages are sent every 10 seconds on Ethernet Links. When four hellos are lost the neighbor is declared dead.

Point-to-point networks always connect two routers, thus there is no need for DR and BDR election.

The first task requires the troubleshooting of a preconfigured OSPF network according Table 8 and Figure 4. Again, it is recommended that before you start to read all tasks. You will notice that the last step is to describe the problems you found during this troubleshooting task. Knowing that, you can write down your findings while performing every individual task. At the end you will have a complete list with all the issues found.

1) Let's start with loading the configuration on all eight devices for chapter two part 1 (`ch2-start.config`).

```
[edit]
lab@R1# load override ch2-start.config
```

Note: You will need to load the **vrX-chapters.config** on both the **vr-devices** as well.

2) Verify the OSPF configuration and adjacencies.

a) R1

```
[edit]
lab@R1# show protocols
protocols {
    ospf {
        area 0.0.0.3 {
            interface ae0.0 {
                authentication {
                    md5 1 key "$9$eGdWNbHkPTF/ZUi.5FAtevWx-w"; ## SECRET-DATA
                }
            }
        }
        area 0.0.0.2 {
            interface ge-0/0/4.118 {
                authentication {
                    md5 1 key "$9$ecsWNbHkPTF/ZUi.5FAtevWx-w"; ## SECRET-DATA
                }
            }
            interface lo0.0;
        }
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: Protocols in Junos are configured under the `protocols` hierarchy.

OSPF interfaces configuration is applied under the `area` hierarchy. The area ID must be present and must match in the Hello packets exchanged between neighbors. The router R1 has interfaces in two areas; area 2 has two interfaces and area 3 with just one interface. The loopback interface is also configured for OSPF, fulfilling the requirement for loopback reachability of all routers.

Note: None of these areas are the backbone area.

Note: The area ID is displayed in dotted decimal format. The area ID is a 32-bit number and you can configure it as a single decimal number as Junos will convert it to a dotted decimal format.

Let's display the active adjacencies, used as a first check when troubleshooting the OSPF protocol. You may see slightly different results with the election of the DR and BDR.

```
[edit]
lab@R1# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.10      ge-0/0/4.118           Full      172.30.5.8       128    39
```

Only one neighbor in *Full* state is displayed on interface `ge-0/0/4.118`. Also, the router ID of the neighbor is listed. The priority is left to the default value of 128. The dead timer decreases until a next Hello is received.

The next command displays all configured OSPF interfaces and the active neighbors on each interface:

```
[edit]
lab@R1# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.118       BDR     0.0.0.2         172.30.5.8      172.30.5.1         1
lo0.0              DR      0.0.0.2         172.30.5.1      0.0.0.0            0
ae0.0              DR      0.0.0.3         172.30.5.1      0.0.0.0            0
```

Three interfaces are listed. As with the previous command only one interface has an active neighbor. The loopback interface has no neighbors, as a looped software interface it is not expected to have one.

Note: Router ID could be referencing an IP address, which is not advertised in OSPF.

© 2020 Juniper Networks, Inc. All rights reserved.

The state column is showing the role of the router on the segment. The local router on interface `ge-0/0/4.118` is a BDR. The other neighbor on the segment is DR with Router ID `172.30.5.8`.

The area in which an interface participates is displayed next to the State column. You can compare if each interface is configured in the correct Area according to Table 8.

b) R2

```
[edit]
lab@R2# run show protocols ospf
area 0.0.0.33 {
    interface ae0.0 {
        authentication {
            md5 1 key "$9$Sy9eLNUDkm5F4aGi.56/SreWX-"; ## SECRET-DATA
        }
    }
}
area 0.0.0.0 {
    interface lo0.0;
    interface ge-0/0/4.127 {
        authentication {
            md5 1 key "$9$38y990IXxdw2aKMLNb2GU369pOR"; ## SECRET-DATA
        }
    }
    interface ge-0/0/4.123 {
        authentication {
            md5 1 key "$9$QtsSntOW87dwgreMX-waJQFnCpB"; ## SECRET-DATA
        }
    }
}
```

Listing the active neighbors on R2.

```
[edit]
lab@R2# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.14      ge-0/0/4.123           ExStart   172.30.5.3       128    34
```

Here again only one neighbor is displayed but the state of the neighbor is *ExStart*.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.123       BDR     0.0.0.0         172.30.5.3      172.30.5.2         1
ge-0/0/4.127       DR      0.0.0.0         172.30.5.2      0.0.0.0            0
lo0.0              DR      0.0.0.0         172.30.5.2      0.0.0.0            0
ae0.0              DR      0.0.0.33        172.30.5.2      0.0.0.0            0
```

R2 has four interfaces configured in OSPF including the loopback0.0 interface.

c) R3

```
 [edit]
lab@R3# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/4.123 {
        authentication {
            md5 1 key "$9$T3A0MWx-b2ylvLNboaTz39tO"; ## SECRET-DATA
        }
    }
    interface lo0.0;
    interface ge-0/0/4.136 {
        authentication {
            md5 1 key "$9$L3KNs4f5F6CuHqPQnCB1LxNbYo"; ## SECRET-DATA
        }
    }
}
area 0.0.0.4 {
    nssa {
        area-range 172.30.32.0/20;
    }
    interface ge-0/0/4.134 {
        authentication {
            md5 1 key "$9$B6BElMg4ZDHmVw2aUH5TBIEyeW"; ## SECRET-DATA
        }
    }
}
```

List the active neighbors on R3.

```
[edit]
lab@R3# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.13      ge-0/0/4.123           ExStart   172.30.5.2       128    33
```

R3 has only one neighbor in the *ExStart* state.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.123       DR      0.0.0.0         172.30.5.3      172.30.5.2         1
ge-0/0/4.136       DR      0.0.0.0         172.30.5.3      0.0.0.0            0
lo0.0              DR      0.0.0.0         172.30.5.3      0.0.0.0            0
ge-0/0/4.134       DR      0.0.0.4         172.30.5.3      0.0.0.0            0
```

Only one of the configured interfaces for OSPF has a neighbor.

d) R4

```
[edit]
lab@R4# show protocols ospf
export rip-to-ospf;
area 0.0.0.4 {
    stub;
    interface ge-0/0/4.134 {
        authentication {
            md5 1 key "$9$eGdWNbHkPTF/ZUi.5FAtevWx-w"; ## SECRET-DATA
        }
    }
    interface ge-0/0/4.145 {
        authentication {
            md5 1 key "$9$GGiPQ1IcrK8tuBEyKx7GDi.fz"; ## SECRET-DATA
        }
    }
    interface lo0.0;
}
```

List the active neighbors on R4.

```
[edit]
lab@R4# run show ospf neighbor
```

List all interfaces configured for OSPF.

```
[edit]
lab@R4# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.134       DR      0.0.0.4         172.30.5.4      0.0.0.0            0
ge-0/0/4.145       DR      0.0.0.4         172.30.5.4      0.0.0.0            0
lo0.0              DR      0.0.0.4         172.30.5.4      0.0.0.0            0
```

No active neighbors are displayed in the output of both commands on R4.

© 2020 Juniper Networks, Inc. All rights reserved.

e) R5

```
[edit]
lab@R5# show protocols ospf
export rip-to-ospf;
area 0.0.0.4 {
    nssa;
    interface ge-0/0/4.145 {
        authentication {
            md5 1 key "$9$-PwoGF39t0IP5z6A0hc-VwgaU"; ## SECRET-DATA
        }
    }
    interface ae0.0 {
        authentication {
            md5 1 key "$9$gAJi.tpBRcln/A0IcvMgoJDHm"; ## SECRET-DATA
        }
    }
    interface lo0.0;
}
```

Display the active OSPF neighbors.

```
[edit]
lab@R5# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.34      ae0.0                  Full      172.30.5.2       128    34

[edit]
lab@R5# run show ospf interface
Interface        State   Area            DR ID           BDR ID          Nbrs
ae0.0            DR      0.0.0.4         172.30.5.5      172.30.5.2         1
ge-0/0/4.145     DR      0.0.0.4         172.30.5.5      0.0.0.0            0
lo0.0            DR      0.0.0.4         172.30.5.5      0.0.0.0            0
```

One active neighbor in the *Full* state in Area 4 on interface *ae0.0*.

f) R6

```
 [edit]
lab@R6# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/4.136 {
        authentication {
            md5 1 key "$9$z6dnn9peK87NbIElM"; ## SECRET-DATA
        }
    }
    interface ge-0/0/4.167 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        authentication {
            md5 1 key "$9$gyJi.tpBRcln/A0IcvMgoJDHm"; ## SECRET-DATA
        }
    }
    interface lo0.0;
}
area 0.0.0.4 {
    nssa {
        default-lsa default-metric 10;
        no-summaries;
    }
    area-range 172.30.32.0/20;
    interface ae0.0 {
        authentication {
            md5 1 key "$9$PQ/teK8x-whSlMX-2gP5Qn9p"; ## SECRET-DATA
        }
    }
}
```

Display the active OSPF neighbors.

```
[edit]
lab@R6# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.42      ge-0/0/4.167           Full      172.30.5.7       128    33
172.30.0.33      ae0.0                  Full      172.30.5.5       128    30

[edit]
lab@R6# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.136       DR      0.0.0.0         172.30.5.6      0.0.0.0            0
ge-0/0/4.167       BDR     0.0.0.0         172.30.5.7      172.30.5.6         1
lo0.0              DR      0.0.0.0         172.30.5.6      0.0.0.0            0
ae0.0              BDR     0.0.0.4         172.30.5.5      172.30.5.6         1
```

Two neighbors in *Full* state, one in Area 0 and one in Area 4.

g) R7

```
 [edit]
lab@R7# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/4.127 {
        authentication {
            md5 1 key "$9$3HFM90IXxdw2aKMLNb2GU369pOR"; ## SECRET-DATA
        }
    }
    interface ge-0/0/4.167 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        authentication {
            md5 1 key "$9$OeHXRrv2gJUi.dbYoGif5O1RSlM"; ## SECRET-DATA
        }
    }
    interface lo0.0;
}
area 0.0.0.1 {
    interface ge-0/0/4.178 {
        authentication {
            md5 1 key "$9$Nab4ZzF/Au1mfQnCuEhNdb2oG"; ## SECRET-DATA
        }
    }
}
```

Display the active OSPF neighbors.

```
[edit]
lab@R7# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.41      ge-0/0/4.167           Full      172.30.5.2       128    38
172.30.0.46      ge-0/0/4.178           Full      172.30.5.8       128    32

[edit]
lab@R7# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.127       DR      0.0.0.0         172.30.5.7      0.0.0.0            0
ge-0/0/4.167       DR      0.0.0.0         172.30.5.7      172.30.5.6         1
lo0.0              DR      0.0.0.0         172.30.5.7      0.0.0.0            0
ge-0/0/4.178       BDR     0.0.0.1         172.30.5.8      172.30.5.7         1
```

Two neighbors are in the *Full* state, one in Area 0 and one in Area 1.

h) R8

```
[edit]
lab@R8# show protocols ospf
area 0.0.0.1 {
    interface ge-0/0/4.178 {
        authentication {
            md5 1 key "$9$RzBSvLaJDH.5s2oGi.zFRhSeMX"; ## SECRET-DATA
        }
    }
}
area 0.0.0.2 {
    interface ge-0/0/4.118 {
        authentication {
            md5 1 key "$9$IDmcK8oaUiqfwY4ZjqQzIEclvL"; ## SECRET-DATA
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
    interface lo0.0;
}
```

As with R1, R8 has interfaces in two non-backbone areas.

Display the active OSPF neighbors.

```
[edit]
lab@R8# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.45      ge-0/0/4.178           Full      172.30.5.7       128    33
172.30.0.9       ge-0/0/4.118           Full      172.30.5.1       128    33

[edit]
lab@R8# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.178       DR      0.0.0.1         172.30.5.8      172.30.5.7         1
ge-0/0/4.118       DR      0.0.0.2         172.30.5.8      172.30.5.1         1
lo0.0              DR      0.0.0.2         172.30.5.8      0.0.0.0            0
```

R8 has active neighbors in *Full* state on all interfaces configured for OSPF.

3) Fix the broken OSPF adjacencies.

Now you have a brief overview of the OSPF configurations on all 8 routers. Most of the routers have interfaces configured for OSPF, but without neighbors. You have to fix these adjacencies before you continue with OSPF route advertisement and loopback reachability troubleshooting.

Note: Although all interfaces are multi access broadcast Ethernet type, the topology is point-to-point. All connections are back-to-back Ethernet, only two peers are expected on each segment.

From the list with active OSPF interfaces you saw that interface *ae0.0* has zero neighbors. To check why the adjacency is not established you can use the Junos traceoption feature. Traceoptions can be used for troubleshooting all protocols on a Junos device. To use traceoptions you must add additional configurations to the `protocol` hierarchy and commit the changes to start tracking and record debug information in specified file.

© 2020 Juniper Networks, Inc. All rights reserved.

a) R1 – R2 adjacency.

Let's configure traceoptions for OSPF on R1 for all errors with option *detail.*

```
[edit]
lab@R1# set protocols ospf traceoptions file ospf.log

[edit]
lab@R1# set protocols ospf traceoptions flag error detail

[edit]
lab@R1# commit
commit complete
```

Note: You can enable one or multiple flags for the same trace file. You can use the detail option to get more detailed information.

Now you can see the content of the file:

```
[edit]
lab@R1# run show log ospf.log
Jun 25 19:36:22 trace_on: Tracing to "/var/log/ospf.log" started
Jun 25 19:36:24.350792 OSPF packet ignored: area mismatch (0.0.0.33) from
172.30.0.2 on intf ae0.0 area 0.0.0.3
Jun 25 19:36:24.351273 OSPF rcvd Hello 172.30.0.2 -> 224.0.0.5 (ae0.0 IFL 69 area
0.0.0.3)
Jun 25 19:36:24.351331   Version 2, length 44, ID 172.30.5.2, area 0.0.0.33
Jun 25 19:36:24.351416   checksum 0x0, authtype 2
Jun 25 19:36:24.351472   mask 255.255.255.252, hello_ivl 10, opts 0x12, prio 128
Jun 25 19:36:24.351540   dead_ivl 40, DR 172.30.0.2, BDR 0.0.0.0
```

You can see that on interface ae0.0 there is an area mismatch. The neighbor is sending hello packets with Area ID 33.

Listing the interfaces configured for OSPF on both routers clearly shows that as well.

```
[edit]
lab@R1# run show ospf interface ae0.0
Interface          State   Area            DR ID           BDR ID          Nbrs
ae0.0              DR      0.0.0.3         172.30.5.1      0.0.0.0            0
```

```
[edit]
lab@R2# run show ospf interface ae0.0
Interface          State   Area            DR ID           BDR ID          Nbrs
ae0.0              DR      0.0.0.33        172.30.5.2      0.0.0.0            0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Fix the configuration on R2.

```
[edit]
lab@R2# rename protocols ospf area 33 to area 3

[edit]
lab@R2# commit
commit complete
```

Note: Using the **rename** command will rename the configuration parameters. This way you can fix the error with just one command.

Let's list again the interfaces enabled for OSPF.

```
[edit]
lab@R1# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.118       BDR     0.0.0.2         172.30.5.8      172.30.5.1         1
lo0.0              DR      0.0.0.2         172.30.5.1      0.0.0.0            0
ae0.0              DR      0.0.0.3         172.30.5.1      172.30.5.2         1
```

```
[edit]
lab@R2# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.123       BDR     0.0.0.0         172.30.5.3      172.30.5.2         1
ge-0/0/4.127       DR      0.0.0.0         172.30.5.2      0.0.0.0            0
lo0.0              DR      0.0.0.0         172.30.5.2      0.0.0.0            0
ae0.0              BDR     0.0.0.3         172.30.5.1      172.30.5.2         1
```

All interfaces on R1 have active neighbors. It is recommended to deactivate or delete the configuration for *traceoptions* on R1 especially in a production environment. Note that by default the trace files are archived when the file size reaches 128KB and overwritten when the number of files reach 3.

```
[edit]
lab@R1# delete protocols ospf traceoptions

[edit]
lab@R1# commit
commit complete
```

b) R2 – R3 adjacency.

Interface **ge-0/0/4.123** on R2 is connected to R3. From the previous Task you know that it has only reached the OSPF `ExStart` state. This is the state where the database exchange starts. For this state the MTU must match on each side of the link.

© 2020 Juniper Networks, Inc. All rights reserved.

Again, you can use traceoptions on R2 to troubleshoot neighbor establishment.

```
[edit]
lab@R2# set protocols ospf traceoptions file ospf.log

[edit]
lab@R2# set protocols ospf traceoptions flag error detail

[edit]
lab@R2# commit
commit complete
```

Display the content of the trace file.

```
[edit]
lab@R2# run show log ospf.log
Jun 25 21:29:36 trace_on: Tracing to "/var/log/ospf.log" started
Jun 25 21:29:37.576450 OSPF packet ignored: MTU mismatch from 172.30.0.14 on intf
ge-0/0/4.123 area 0.0.0.0
Jun 25 21:29:41.805027 OSPF packet ignored: MTU mismatch from 172.30.0.14 on intf
ge-0/0/4.123 area 0.0.0.0
Jun 25 21:29:45.586188 OSPF packet ignored: MTU mismatch from 172.30.0.14 on intf
ge-0/0/4.123 area 0.0.0.0
```

We see that OSPF ignores the packets because of an MTU mismatch on interface `ge-0/0/4.123`.

Display the list of OSPF configured interfaces with the `extensive` options shows the MTU for each interface.

```
[edit]
lab@R2# run show ospf interface ge-0/0/4.123 extensive
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.123       BDR     0.0.0.0         172.30.5.3      172.30.5.2         1
  Type: LAN, Address: 172.30.0.13, Mask: 255.255.255.252, MTU: 1400, Cost: 1
  DR addr: 172.30.0.14, BDR addr: 172.30.0.13, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan  1 01:00:00 CET
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show ospf interface ge-0/0/4.123 extensive
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.123       DR      0.0.0.0         172.30.5.3      172.30.5.2         1
  Type: LAN, Address: 172.30.0.14, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  DR addr: 172.30.0.14, BDR addr: 172.30.0.13, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan  1 01:00:00 CET
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

By default, the MTU on all Ethernet interfaces is 1500 bytes. The MTU of interface ge-0/0/4.123 is 1400. This means that the interface configuration on R2 has an explicit MTU configured for family inet.

```
[edit]
lab@R2# show interfaces ge-0/0/4.123
description "R3 connection";
vlan-id 123;
family inet {
    mtu 1400;
    address 172.30.0.13/30;
}
family inet6;
```

Fix the configuration by deleting the MTU configuration on R2. This way the interface MTU reverts back to the default value of 1500.

```
[edit]
lab@R2# delete interfaces ge-0/0/4.123 family inet mtu

[edit]
lab@R2# commit
commit complete
```

After equalizing the MTU on both ends, the OSPF neighbors reach the *Full* state.

```
[edit]
lab@R2# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.14      ge-0/0/4.123           Full      172.30.5.3       128    34
172.30.0.1       ae0.0                  Full      172.30.5.1       128    39
```

c) R2 – R7 adjacency.

© 2020 Juniper Networks, Inc. All rights reserved.

R7 should have neighbors with routers R2, R6 and R8. You can see that only the neighbor with R2 is not established.

R2 OSPF traceoptions are still enabled. Looking at the log file you can see that there is subnet mismatch on interface ge-0/0/4.127.

```
[edit]
lab@R2# run show log ospf.log
Jun 25 20:01:01 trace_on: Tracing to "/var/log/ospf.log" started
Jun 25 20:01:04.326414 OSPF packet ignored: subnet mismatch from 172.30.1.18 on
intf ge-0/0/4.127 area 0.0.0.0
Jun 25 20:01:04.326882 OSPF rcvd Hello 172.30.1.18 -> 224.0.0.5 (ge-0/0/4.127 IFL
74 area 0.0.0.0)
Jun 25 20:01:04.326952   Version 2, length 44, ID 172.30.5.7, area 0.0.0.0
Jun 25 20:01:04.327002   checksum 0x0, authtype 2
Jun 25 20:01:04.327091   mask 255.255.255.252, hello_ivl 10, opts 0x12, prio 128
Jun 25 20:01:04.327157   dead_ivl 40, DR 172.30.1.18, BDR 0.0.0.0
```

Display the extensive OSPF interface information and you can see the error in subnet configuration.

```
[edit]
lab@R2# run show ospf interface ge-0/0/4.127 extensive
Interface          State   Area        DR ID           BDR ID          Nbrs
ge-0/0/4.127       DR      0.0.0.0     172.30.5.2      0.0.0.0            0
  Type: LAN, Address: 172.30.0.17, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  DR addr: 172.30.0.17, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 01:00:00 CET
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

```
[edit]
lab@R7# run show ospf interface ge-0/0/4.127 detail
Interface          State   Area        DR ID           BDR ID          Nbrs
ge-0/0/4.127       DR      0.0.0.0     172.30.5.7      0.0.0.0            0
  Type: LAN, Address: 172.30.1.18, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  DR addr: 172.30.1.18, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 01:00:00 CET
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Using again the rename feature you can change the interface IP address on R7 with just one command.

```
[edit]
lab@R7# rename interface ge-0/0/4.127 family inet address 172.30.1.18/30 to
address 172.30.0.18/30

[edit]
lab@R7# commit
commit complete
```

Fixing the interface address configuration puts the neighbors to *Full* state.

```
[edit]
lab@R2# run show ospf interface
Interface          State    Area            DR ID           BDR ID          Nbrs
ge-0/0/4.123       BDR      0.0.0.0         172.30.5.3      172.30.5.2         1
ge-0/0/4.127       DR       0.0.0.0         172.30.5.2      172.30.5.7         1
lo0.0              DR       0. 0.0.0         172.30.5.2       0.0.0.0           0
ae0.0              BDR      0.0.0.3         172.30.5.1      172.30.5.2         1
```

```
[edit]
lab@R7# run show ospf interface
Interface          State    Area            DR ID           BDR ID          Nbrs
ge-0/0/4.127       BDR      0.0.0.0         172.30.5.2      172.30.5.7         1
ge-0/0/4.167       DR       0.0.0.0         172.30.5.7      172.30.5.2         1
lo0.0              DR       0.0.0.0         172.30.5.7      0.0.0.0            0
ge-0/0/4.178       BDR      0.0.0.1         172.30.5.8      172.30.5.7         1
```

Note: Delete or deactivate the traceoptions configuration on R2 before moving on to troubleshooting R3.

```
[edit]
lab@R2# delete protocols ospf traceoptions

[edit]
lab@R2# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

d) R3 – R4 adjacency.

On R3 only the neighbor with R2 is in the *Full* state, which was fixed in the previous step. Fixing the neighbor establishment on the other two interfaces requires again enabling traceoptions.

Enable traceoptions on R3.

```
[edit]
lab@R3# set protocols ospf traceoptions file ospf.log

[edit]
lab@R3# set protocols ospf traceoptions flag error detail

[edit]
lab@R3# commit
commit complete
```

The interface connected to R4 reports a different type of Area configured on the other end.

```
[edit]
lab@R3# run show log ospf.log
Jun 25 19:47:56 trace_on: Tracing to "/var/log/ospf.log" started
Jun 25 19:47:59.581812 OSPF packet ignored: area nssaness mismatch from
172.30.0.22 on intf ge-0/0/4.134 area 0.0.0.4
```

Area 4 is configured as a *NSSA* on R3, but on R4 the same Area is configured as *STUB*.

```
[edit]
lab@R3# show protocols ospf area 4
nssa {
        area-range 172.30.32.0/20;
}
interface ge-0/0/4.134 {
        authentication {
            md5 1 key "$9$B6BElMg4ZDHmVw2aUH5TBIEyeW"; ## SECRET-DATA
        }
}
```

```
[edit]
lab@R4# show protocols ospf area 4
stub;
interface ge-0/0/4.134 {
        authentication {
            md5 1 key "$9$eGdWNbHkPTF/ZUi.5FAtevWx-w"; ## SECRET-DATA
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 }
interface ge-0/0/4.145 {
        authentication {
            md5 1 key "$9$GGiPQ1IcrK8tuBEyKx7GDi.fz"; ## SECRET-DATA
        }
    }
interface lo0.0;
```

Change the Area type on R4 to *NSSA*.

```
[edit]
lab@R4# set protocols ospf area 4 nssa

[edit]
lab@R4# commit
commit complete
```

This results in one more neighbor in state *Full* on R3.

```
[edit]
lab@R3# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.123       DR      0.0.0.0         172.30.5.3      172.30.5.2         1
ge-0/0/4. 136       DR      0.0.0.0         172.30.5.3      0.0.0.0            0
lo0.0              DR      0.0.0.0         172.30.5.3      0.0.0.0            0
ge-0/0/4.134       DR      0.0.0.4         172.30.5.3      172.30.5.4         1
```

e) R3 – R6 adjacency

Again, check the ospf.log file on R3. You can see the reason why the neighborship with R6 is not established.

```
[edit]
lab@R3# run show log ospf.log
Jun 25 19:47:56 trace on: Tracing to "/var/log/ospf.log" started
Jun 25 19:48:02.210493 OSPF packet ignored: authentication failure (bad cksum).
Jun 25 19:48:02.211206 OSPF packet ignored: authentication failure from
172.30.0.26
```

You probably noticed that all interfaces configured for OSPF are using MD5 authentication. This is also a requirement for successfully completion of the Task.

OSPF supports simple password, MD5 key hash and IPsec authentication. Messages are not accepted if not correctly authenticated.

© 2020 Juniper Networks, Inc. All rights reserved.

The first method is not very secured because the password is exchanged in clear text format in every message, but it is of course more than no authentication. MD5 authentication is a one-way generated hash based on a shared key. The receiving side generates the same hash and compares both. If they are the same the packet is accepted. The shared key is not exchanged, and it is practically impossible to calculate the key from the hash. IPsec authentication uses a manual created security association attached to the interface. This method is available for OSPFv3 using the embedded IPsec header in IPv6 packet.

It is obvious from the traceoptions file that the MD5 authentication shared key is wrong. Because you do not know the key used and it is not given in the beginning. You can copy the hashed key from R3 to R6 for interface `ge-0/0/4.136` directly how it is displayed with **`show protocols ospf`**.

```
[edit]
lab@R6# set protocols ospf area 0 interface ge-0/0/4.136 authentication md5 1 key
"$9$xe7-b2.mT3nCZUqf"

[edit]
lab@R6# commit
commit complete
```

Fixing the authentication issue resolves all neighbor adjacencies issues on R3 and R6.

```
[edit]
lab@R3# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.123       DR      0.0.0.0         172.30.5.3      172.30.5.2         1
ge-0/0/4.136       DR      0.0.0.0         172.30.5.3      172.30.5.6         1
lo0.0              DR      0.0.0.0         172.30.5.3      0.0.0.0            0
ge-0/0/4.134       DR      0.0.0.4         172.30.5.3      172.30.5.4         1

[edit]
lab@R6# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.136       BDR     0.0.0.0         172.30.5.3      172.30.5.6         1
ge-0/0/4.167       BDR     0.0.0.0         172.30.5.7      172.30.5.6         1
lo0.0              DR      0.0.0.0         172.30.5.2      0.0.0.0            0
ae0.0              BDR     0.0.0.4         172.30.5.5      172.30.5.6         1
```

We have reached the point that all neighbor adjacencies are established and in the *Full* state. Also, all routers should use MD5 authentication for these adjacencies. Great work so far!

Delete or deactivate the OSPF traceoptions before moving on in your troubleshooting.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# delete protocols ospf traceoptions

[edit]
lab@R3# commit
commit complete
```

4) Let's verify the LSDB.

With all neighbors established you can now focus to fix the other issues and configure the network according to the task requirements.

First, you must examine the content of the link state database on all devices.

a) R1.

```
[edit]
lab@R1# run show ospf database

    OSPF database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.1       172.30.5.1      0x80000013  1684  0x22 0xaa4a  48
Router   172.30.5.8       172.30.5.8      0x8000000b  2330  0x22 0xf2f3  48
Network  172.30.0.10      172.30.5.8      0x80000005  1634  0x22 0xe7fb  32
Summary *172.30.0.0       172.30.5.1      0x80000008  1720  0x22 0x1983  28
Summary  172.30.0.44      172.30.5.8      0x80000007   512  0x22 0x3733  28

    OSPF database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.1       172.30.5.1      0x80000011  1684  0x22 0xc91f  36
Router   172.30.5.2       172.30.5.2      0x8000001a    51  0x22 0xb527  36
Network *172.30.0.1       172.30.5.1      0x80000002   103  0x22 0x3ac2  32
Summary *172.30.0.8       172.30.5.1      0x80000008   952  0x22 0xc8cb  28
Summary  172.30.0.12      172.30.5.2      0x80000026  3600  0x22 0x5e13  28
Summary  172.30.0.16      172.30.5.2      0x80000026  3600  0x22 0x3637  28
Summary  172.30.0.20      172.30.5.2      0x80000013  3600  0x22 0x3e3d  28
Summary  172.30.0.24      172.30.5.2      0x80000013  3600  0x22 0x1661  28
Summary  172.30.0.28      172.30.5.2      0x80000013  3600  0x22 0xf77a  28
Summary  172.30.0.32      172.30.5.2      0x80000013  3600  0x22 0xd993  28
Summary  172.30.0.40      172.30.5.2      0x80000013  3600  0x22 0x75f1  28
Summary  172.30.0.44      172.30.5.2      0x80000013  3600  0x22 0x4d16  28
Summary *172.30.5.1       172.30.5.1      0x80000008   528  0x22 0xdfb4  28
Summary  172.30.5.2       172.30.5.2      0x80000026  3600  0x22 0x93e0  28
Summary  172.30.5.3       172.30.5.2      0x80000013  3600  0x22 0xb9cb  28
Summary  172.30.5.4       172.30.5.2      0x80000013  3600  0x22 0xb9c9  28
Summary  172.30.5.5       172.30.5.2      0x80000013  3600  0x22 0xb9c7  28
Summary  172.30.5.7       172.30.5.2      0x80000013  3600  0x22 0x91ef  28
Summary *172.30.5.8       172.30.5.1      0x80000006  1406  0x22 0xa7e6  28
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ASBRSum  172.30.5.3       172.30.5.2       0x80000013  3600  0x22 0xabd8  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   172.30.32.0      172.30.5.3       0x80000001   442  0x22 0xcf4   36
```

The command above displays the content of the link state database. R1 has interfaces in two Areas respectively. The content should be the same on all routers in one Area. The last entry in the output shown is one LSA which carries external information.

The first column displays the type of the LSA. The "*" represents LSAs originated on the local router.

- Type *"Router"* – LSA Type 1 (generated by every router)
- Type *"Network"* – LSA Type 2 (generated only by the DR on the segment)
- Type *"Summary"* – LSA Type 3 (generated by the ABR)
- Type *"ASBRSum"* – LSA Type 4 (generated by ABR to advertise ASBR)
- Type *"Extern"* LSA Type 5 (external prefixes generated by ASBR)
- Type *"NSSA"* LSA Type 7 (external prefixes inside NSSA Area)

"ID" represents the prefix carried by the LSA. The id used for Router LSAs is the router-id of the device. In case of Summary LSA, this id is the prefix in other Areas reachable through the ABR. For External LSAs, the ID is the external prefix advertised in OSPF domain.

With the usage of the *extensive* option you can display detailed content for each LSA.

```
[edit]
lab@R1# run show ospf database router lsa-id 172.30.5.1 extensive

    OSPF database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.1       172.30.5.1       0x80000014   902  0x22 0xa84b  48
  bits 0x1, link count 2
  id 172.30.0.10, data 172.30.0.9, Type Transit (2)
    Topology count: 0, Default metric: 1
  id 172.30.5.1, data 255.255.255.255, Type Stub (3)
    Topology count: 0, Default metric: 0
  Topology default (ID 0)
    Type: Transit, Node ID: 172.30.0.10
      Metric: 1, Bidirectional
  Gen timer 00:34:52
  Aging timer 00:44:57
  Installed 00:15:02 ago, expires in 00:44:58, sent 00:15:00 ago
  Last changed 03:56:05 ago, Change count: 3, Ours
```


© 2020 Juniper Networks, Inc. All rights reserved.

Router LSAs are generated on every OSPF router in every Area where the router has participating interfaces in. The bits field is a one octet field and represents the characteristics of the local router. The leading five bits are zero, whereas the last three lowest bits are known as the V/E/B bits.

- *V bit* – this bit is set when the router is an endpoint of a virtual link.
- *E bit* – this bit is set when the router is configured as an ASBR.
- *B bit* – this bit is set when the router is configured as an ABR.

In this case the router is and ABR as only the B bit is set. Please note that the bits are represented in hexadecimal format. 0x1 equals the B bit.

b) R2

```
[edit]
lab@R2# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.2      172.30.5.2      0x80001537     5  0x22 0xf8a0  60
Router   172.30.5.3      172.30.5.3      0x80000017   299  0x22 0xd136  60
Router   172.30.5.7      172.30.5.7      0x80000017  1354  0x22 0x9d37  60
Network  172.30.0.14     172.30.5.3      0x80000001   773  0x22 0xc12b  32
Network *172.30.0.17     172.30.5.2      0x8000030f     2  0x22 0xc015  32
Network  172.30.0.25     172.30.5.3      0x80000001   299  0x22 0x538e  32
Network  172.30.0.42     172.30.5.7      0x80000008   276  0x22 0xaa17  32
Summary *172.30.0.0      172.30.5.2      0x80000ae5     4  0x22 0x397a  28
Summary *172.30.0.20     172.30.5.2      0x80000123  3600  0x22 0x2544  28
Summary  172.30.0.20     172.30.5.3      0x8000000a    93  0x22 0x4044  28
Summary *172.30.0.28     172.30.5.2      0x80000ae2  3600  0x22 0x3069  28
Summary  172.30.0.28     172.30.5.3      0x80000001   497  0x22 0xc78   28
Summary *172.30.0.32     172.30.5.2      0x80000ae3  3600  0x22 0xfb99  28
Summary  172.30.0.32     172.30.5.3      0x80000001   497  0x22 0xed91  28
Summary  172.30.0.44     172.30.5.7      0x8000000a   490  0x22 0x3731  28
Summary *172.30.5.4      172.30.5.2      0x80000123  3600  0x22 0x96db  28
Summary  172.30.5.4      172.30.5.3      0x80000001   497  0x22 0xcdc7  28
Summary *172.30.5.5      172.30.5.2      0x80000ae2  3600  0x22 0xe7c1  28
Summary  172.30.5.5      172.30.5.3      0x80000001   497  0x22 0xcdc5  28

    OSPF database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.1      172.30.5.1      0x80000011  1741  0x22 0xc91f  36
Router  *172.30.5.2      172.30.5.2      0x8000001a   107  0x22 0xb527  36
Network  172.30.0.1      172.30.5.1      0x80000002   161  0x22 0x3ac2  32
Summary  172.30.0.8      172.30.5.1      0x80000008  1010  0x22 0xc8cb  28
Summary *172.30.0.12     172.30.5.2      0x80000030    10  0x22 0x4a1d  28
Summary *172.30.0.16     172.30.5.2      0x80000030    10  0x22 0x2241  28
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Summary *172.30.0.20      172.30.5.2      0x80000018    10  0x22 0x3442  28
Summary *172.30.0.24      172.30.5.2      0x80000018    10  0x22 0xc66   28
Summary *172.30.0.28      172.30.5.2      0x80000018    10  0x22 0xed7f  28
Summary *172.30.0.32      172.30.5.2      0x80000018    10  0x22 0xcf98  28
Summary *172.30.0.40      172.30.5.2      0x80000018    10  0x22 0x6bf6  28
Summary *172.30.0.44      172.30.5.2      0x80000018    10  0x22 0x431b  28
Summary  172.30.5.1       172.30.5.1      0x80000008   586  0x22 0xdfb4  28
Summary *172.30.5.2       172.30.5.2      0x80000030    10  0x22 0x7fea  28
Summary *172.30.5.3       172.30.5.2      0x80000018    10  0x22 0xafd0  28
Summary *172.30.5.4       172.30.5.2      0x80000018    10  0x22 0xafce  28
Summary *172.30.5.5       172.30.5.2      0x80000018    10  0x22 0xafcc  28
Summary *172.30.5.7       172.30.5.2      0x80000018    10  0x22 0x87f4  28
Summary  172.30.5.8       172.30.5.1      0x80000006  1463  0x22 0xa7e6  28
ASBRSum *172.30.5.3       172.30.5.2      0x80000018    10  0x22 0xa1dd  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Extern   172.30.32.0      172.30.5.3      0x80000001   497  0x22 0xcf4   36
```

Several things are important from the output on R2. R2 is an ABR having interfaces in Area 0 and Area 3. According to Figure 4, four routers should have interfaces in Area 0, but from the R2 output only 3 LSA Type 1 are shown in the database. Some of the LSAs in Area 0 have an age time of 3600. This is the maximum age for LSAs, meaning no refresh of this LSA arrived last 3600 seconds and they will be deleted. Usually the originator of LSA sends a copy of it LSA every 1800 seconds. The LSAs were still kept in the LSDB from the previous topology, before all neighbor adjacencies were in *Full* state. You should be careful and not get confused until the network is converged.

Note: An ASBRSum LSA is generated from R2 to Area 3 to advertise the ASBR for the domain. In this case, it is R3 with Router-id 172.30.5.3. As you will see below, R3 is aggregating the external routes learned by R4 and R5.

c) R3

```
[edit]
lab@R3# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Router   172.30.5.2       172.30.5.2      0x80001572     2  0x22 0x3479  60
Router  *172.30.5.3       172.30.5.3      0x80000017   348  0x22 0xd136  60
Router   172.30.5.7       172.30.5.7      0x80000017  1406  0x22 0x9d37  60
Network *172.30.0.14      172.30.5.3      0x80000001   822  0x22 0xc12b  32
Network  172.30.0.17      172.30.5.2      0x8000032d  3600  0x22 0x8433  32
Network *172.30.0.25      172.30.5.3      0x80000001   348  0x22 0x538e  32
Network  172.30.0.42      172.30.5.7      0x80000008   327  0x22 0xaa17  32
Summary  172.30.0.0       172.30.5.2      0x80000b03  3600  0x22 0xfb99  28
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Summary  172.30.0.20      172.30.5.2       0x80000141  3600  0x22 0xe862  28
Summary *172.30.0.20      172.30.5.3       0x8000000a   142  0x22 0x4044  28
Summary  172.30.0.28      172.30.5.2       0x80000afd  3600  0x22 0xf984  28
Summary *172.30.0.28      172.30.5.3       0x80000001   547  0x22 0xc78   28
Summary  172.30.0.32      172.30.5.2       0x80000afe  3600  0x22 0xc5b4  28
Summary *172.30.0.32      172.30.5.3       0x80000001   547  0x22 0xed91  28
Summary  172.30.0.44      172.30.5.7       0x8000000a   542  0x22 0x3731  28
Summary  172.30.5.4       172.30.5.2       0x8000013e  3600  0x22 0x60f6  28
Summary *172.30.5.4       172.30.5.3       0x80000001   547  0x22 0xcdc7  28
Summary  172.30.5.5       172.30.5.2       0x80000afd  3600  0x22 0xb1dc  28
Summary *172.30.5.5       172.30.5.3       0x80000001   547  0x22 0xcdc5  28

    OSPF database, Area 0.0.0.4
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.2       172.30.5.2       0x8000000c   352  0x20 0xf8b1  36
Router  *172.30.5.3       172.30.5.3       0x8000000f   465  0x20 0xb30b  36
Router   172.30.5.4       172.30.5.4       0x80000005   514  0x20 0xb051  60
Router   172.30.5.5       172.30.5.5       0x80000016   554  0x20 0xcd08  60
Network *172.30.0.21      172.30.5.3       0x80000003   547  0x20 0xb132  32
Network  172.30.0.30      172.30.5.5       0x80000003   547  0x20 0x5f77  32
Network  172.30.0.33      172.30.5.5       0x80000007   997  0x20 0x1db4  32
Summary  0.0.0.0          172.30.5.2       0x80000006  1112  0x20 0xcf8e  28
Summary *172.30.0.12      172.30.5.3       0x8000000b   822  0x20 0xace0  28
Summary *172.30.0.24      172.30.5.3       0x8000000a   348  0x20 0x364c  28
Summary *172.30.0.40      172.30.5.3       0x80000001  3600  0x20 0xb1c8  28
Summary *172.30.0.44      172.30.5.3       0x80000001  3600  0x20 0x93e1  28
Summary *172.30.5.2       172.30.5.3       0x8000001a     7  0x20 0xcdb2  28
Summary *172.30.5.3       172.30.5.3       0x80000009   299  0x20 0xdbb5  28
Summary *172.30.5.6       172.30.5.3       0x80000001  3600  0x20 0xd7bd  28
Summary *172.30.5.7       172.30.5.3       0x80000001  3600  0x20 0xd7bb  28
NSSA     172.30.32.0      172.30.5.4       0x80000002   309  0x28 0x1df7  36
NSSA     172.30.32.0      172.30.5.5       0x80000004   839  0x28 0x21ef  36
NSSA     172.30.33.0      172.30.5.4       0x80000002   135  0x28 0x1202  36
NSSA     172.30.33.0      172.30.5.5       0x80000004   681  0x28 0x16f9  36
NSSA     172.30.34.0      172.30.5.4       0x80000001   554  0x28 0x90b   36
NSSA     172.30.34.0      172.30.5.5       0x80000004   523  0x28 0xb04   36
NSSA     172.30.35.0      172.30.5.4       0x80000001   554  0x28 0xfd15  36
NSSA     172.30.35.0      172.30.5.5       0x80000004   365  0x28 0xff0e  36
NSSA     172.30.36.0      172.30.5.4       0x80000001   554  0x28 0xf21f  36
NSSA     172.30.36.0      172.30.5.5       0x80000004   207  0x28 0xf418  36
NSSA     172.30.37.0      172.30.5.4       0x80000001   554  0x28 0xe729  36
NSSA     172.30.37.0      172.30.5.5       0x80000004    50  0x28 0xe922  36
NSSA     172.30.38.0      172.30.5.4       0x80000001   554  0x28 0xdc33  36
NSSA     172.30.38.0      172.30.5.5       0x80000003  2576  0x28 0xe02b  36
NSSA     172.30.39.0      172.30.5.4       0x80000001   554  0x28 0xd13d  36
NSSA     172.30.39.0      172.30.5.5       0x80000003  2418  0x28 0xd535  36
NSSA     172.30.40.0      172.30.5.4       0x80000003   553  0x28 0xc249  36
NSSA     172.30.40.0      172.30.5.5       0x80000003  2260  0x28 0xca3f  36
NSSA     172.30.41.0      172.30.5.4       0x80000003   553  0x28 0xb753  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
NSSA     172.30.41.0      172.30.5.5       0x80000003  2103  0x28 0xbf49  36
NSSA     172.30.42.0      172.30.5.4       0x80000003   553  0x28 0xac5d  36
NSSA     172.30.42.0      172.30.5.5       0x80000003  1945  0x28 0xb453  36
NSSA     172.30.43.0      172.30.5.4       0x80000003   553  0x28 0xa167  36
NSSA     172.30.43.0      172.30.5.5       0x80000003  1787  0x28 0xa95d  36
NSSA     172.30.44.0      172.30.5.4       0x80000003   553  0x28 0x9671  36
NSSA     172.30.44.0      172.30.5.5       0x80000003  1629  0x28 0x9e67  36
NSSA     172.30.45.0      172.30.5.4       0x80000003   553  0x28 0x8b7b  36
NSSA     172.30.45.0      172.30.5.5       0x80000003  1471  0x28 0x9371  36
NSSA     172.30.46.0      172.30.5.4       0x80000003   553  0x28 0x8085  36
NSSA     172.30.46.0      172.30.5.5       0x80000003  1313  0x28 0x887b  36
NSSA     172.30.47.0      172.30.5.4       0x80000003   553  0x28 0x758f  36
NSSA     172.30.47.0      172.30.5.5       0x80000003  1155  0x28 0x7d85  36
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   *172.30.32.0      172.30.5.3       0x80000001   547  0x22 0xcf4   36
```

Here again, several things should be noticed. R3 has interfaces in Area 0 and Area 4, where Area 4 is a NSSA. You can see all RIP routes advertised from R4 as Type 7. There are only three Type 1 LSAs in Area 0 and some misleading LSAs with an Age 3600 are visible in the database. Another suspicious LSA is the *summary* LSA 0.0.0.0 in Area 4. The problem is that it is originated by 172.30.5.2, which should be the Router-id of R2, but R2 has no interface in Area 4.

An External LSA that is marked with "*" means that it is originated by the local router. This LSA aggregates all prefixes from NSSA LSAs originated from R4 and R5. This is according to the requirement. It looks like the requirement is accomplished without changing anything.

Let's check how the external prefixes are aggregated on R3.

Display the configuration of Area 4.

```
[edit]
lab@R3# show protocols ospf area 4
nssa {
        area-range 172.30.32.0/20;
    }
interface ge-0/0/4.134 {
        authentication {
            md5 1 key "$9$B6BElMg4ZDHmVw2aUH5TBIEyeW"; ## SECRET-DATA
        }
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

The `area-range` command aggregates more specific prefixes in the configured range. One more specific prefix must be available before the aggregate route may be originated by the ABR. If there are no more specific prefixes, then the aggregate is not generated. All more specific prefixes in the range of the aggregate are automatically suppressed and not leaked into the other area.

Note: The `area-range` command is configured under the `nssa` hierarchy of the area configuration. This way you aggregate only the NSSA, external routes for the domain.

d) R4

```
[edit]
lab@R4# run show ospf database

    OSPF database, Area 0.0.0.4
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.2       172.30.5.2      0x8000000c   424  0x20 0xf8b1  36
Router   172.30.5.3       172.30.5.3      0x8000000f   539  0x20 0xb30b  36
Router  *172.30.5.4       172.30.5.4      0x80000005   586  0x20 0xb051  60
Router   172.30.5.5       172.30.5.5      0x80000016   626  0x20 0xcd08  60
Network  172.30.0.21      172.30.5.3      0x80000003   621  0x20 0xb132  32
Network  172.30.0.30      172.30.5.5      0x80000003   619  0x20 0x5f77  32
Network  172.30.0.33      172.30.5.5      0x80000007  1069  0x20 0x1db4  32
Summary  0.0.0.0          172.30.5.2      0x80000006  1184  0x20 0xcf8e  28
Summary  172.30.0.12      172.30.5.3      0x8000000b   896  0x20 0xace0  28
Summary  172.30.0.16      172.30.5.3      0x80000001     2  0x20 0xa2ef  28
Summary  172.30.0.24      172.30.5.3      0x8000000a   423  0x20 0x364c  28
Summary  172.30.0.40      172.30.5.3      0x80000002     2  0x20 0xb9be  28
Summary  172.30.0.44      172.30.5.3      0x80000002     2  0x20 0x91e2  28
Summary  172.30.5.2       172.30.5.3      0x8000001c    11  0x20 0xc9b4  28
Summary  172.30.5.3       172.30.5.3      0x80000009   374  0x20 0xdbb5  28
Summary  172.30.5.7       172.30.5.3      0x80000002     2  0x20 0xd5bc  28
NSSA    *172.30.32.0      172.30.5.4      0x80000002   382  0x28 0x1df7  36
NSSA     172.30.32.0      172.30.5.5      0x80000004   911  0x28 0x21ef  36
NSSA    *172.30.33.0      172.30.5.4      0x80000002   207  0x28 0x1202  36
NSSA     172.30.33.0      172.30.5.5      0x80000004   753  0x28 0x16f9  36
NSSA    *172.30.34.0      172.30.5.4      0x80000002    32  0x28 0x70c   36
NSSA     172.30.34.0      172.30.5.5      0x80000004   595  0x28 0xb04   36
NSSA    *172.30.35.0      172.30.5.4      0x80000001   626  0x28 0xfd15  36
NSSA     172.30.35.0      172.30.5.5      0x80000004   438  0x28 0xff0e  36
NSSA    *172.30.36.0      172.30.5.4      0x80000001   626  0x28 0xf21f  36
NSSA     172.30.36.0      172.30.5.5      0x80000004   280  0x28 0xf418  36
NSSA    *172.30.37.0      172.30.5.4      0x80000001   626  0x28 0xe729  36
NSSA     172.30.37.0      172.30.5.5      0x80000004   123  0x28 0xe922  36
NSSA    *172.30.38.0      172.30.5.4      0x80000001   626  0x28 0xdc33  36
NSSA     172.30.38.0      172.30.5.5      0x80000003  2648  0x28 0xe02b  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
NSSA    *172.30.39.0      172.30.5.4      0x80000001   626  0x28 0xd13d  36
NSSA     172.30.39.0      172.30.5.5      0x80000003  2490  0x28 0xd535  36
NSSA    *172.30.40.0      172.30.5.4      0x80000003   625  0x28 0xc249  36
NSSA     172.30.40.0      172.30.5.5      0x80000003  2332  0x28 0xca3f  36
NSSA    *172.30.41.0      172.30.5.4      0x80000003   625  0x28 0xb753  36
NSSA     172.30.41.0      172.30.5.5      0x80000003  2175  0x28 0xbf49  36
NSSA    *172.30.42.0      172.30.5.4      0x80000003   625  0x28 0xac5d  36
NSSA     172.30.42.0      172.30.5.5      0x80000003  2017  0x28 0xb453  36
NSSA    *172.30.43.0      172.30.5.4      0x80000003   625  0x28 0xa167  36
NSSA     172.30.43.0      172.30.5.5      0x80000003  1859  0x28 0xa95d  36
NSSA    *172.30.44.0      172.30.5.4      0x80000003   625  0x28 0x9671  36
NSSA     172.30.44.0      172.30.5.5      0x80000003  1701  0x28 0x9e67  36
NSSA    *172.30.45.0      172.30.5.4      0x80000003   625  0x28 0x8b7b  36
NSSA     172.30.45.0      172.30.5.5      0x80000003  1543  0x28 0x9371  36
NSSA    *172.30.46.0      172.30.5.4      0x80000003   625  0x28 0x8085  36
NSSA     172.30.46.0      172.30.5.5      0x80000003  1385  0x28 0x887b  36
NSSA    *172.30.47.0      172.30.5.4      0x80000003   625  0x28 0x758f  36
NSSA     172.30.47.0      172.30.5.5      0x80000003  1227  0x28 0x7d85  36
```

R4 has interfaces only in Area 4, configured as NSSA. R4 is injecting RIP routes learned from the external VR device as NSSA LSA Type 7. Remember that type 5 LSAs are not allowed in NSSA. Again, here you can see the strange summary LSA 0.0.0.0 originated from 172.30.5.2.

e) R5

```
[edit]
lab@R5# run show ospf database

    OSPF database, Area 0.0.0.4
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Router   172.30.5.2       172.30.5.2      0x8000000c   681  0x20 0xf8b1  36
Router   172.30.5.3       172.30.5.3      0x8000000f   798  0x20 0xb30b  36
Router   172.30.5.4       172.30.5.4      0x80000005   845  0x20 0xb051  60
Router  *172.30.5.5       172.30.5.5      0x80000016   883  0x20 0xcd08  60
Network  172.30.0.21      172.30.5.3      0x80000003   880  0x20 0xb132  32
Network *172.30.0.30      172.30.5.5      0x80000003   876  0x20 0x5f77  32
Network *172.30.0.33      172.30.5.5      0x80000007  1326  0x20 0x1db4  32
Summary  0.0.0.0          172.30.5.2      0x80000006  1441  0x20 0xcf8e  28
Summary  172.30.0.12      172.30.5.3      0x8000000c   160  0x20 0xaae1  28
Summary  172.30.0.24      172.30.5.3      0x8000000a   682  0x20 0x364c  28
Summary  172.30.0.40      172.30.5.3      0x80000002  3600  0x20 0xafc9  28
Summary  172.30.0.44      172.30.5.3      0x80000002  3600  0x20 0x91e2  28
Summary  172.30.5.2       172.30.5.3      0x80000027     7  0x20 0xb3bf  28
Summary  172.30.5.3       172.30.5.3      0x80000009   633  0x20 0xdbb5  28
Summary  172.30.5.6       172.30.5.3      0x80000002  3600  0x20 0xd5be  28
Summary  172.30.5.7       172.30.5.3      0x80000002  3600  0x20 0xd5bc  28
NSSA     172.30.32.0      172.30.5.4      0x80000002   641  0x28 0x1df7  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
NSSA     *172.30.32.0       172.30.5.5       0x80000004  1168  0x28 0x21ef  36
NSSA      172.30.33.0       172.30.5.4       0x80000002   466  0x28 0x1202  36
NSSA     *172.30.33.0       172.30.5.5       0x80000004  1010  0x28 0x16f9  36
NSSA      172.30.34.0       172.30.5.4       0x80000002   291  0x28 0x70c   36
NSSA     *172.30.34.0       172.30.5.5       0x80000004   852  0x28 0xb04   36
NSSA      172.30.35.0       172.30.5.4       0x80000002   116  0x28 0xfb16  36
NSSA     *172.30.35.0       172.30.5.5       0x80000004   694  0x28 0xff0e  36
NSSA      172.30.36.0       172.30.5.4       0x80000001   885  0x28 0xf21f  36
NSSA     *172.30.36.0       172.30.5.5       0x80000004   536  0x28 0xf418  36
NSSA      172.30.37.0       172.30.5.4       0x80000001   885  0x28 0xe729  36
NSSA     *172.30.37.0       172.30.5.5       0x80000004   379  0x28 0xe922  36
NSSA      172.30.38.0       172.30.5.4       0x80000001   885  0x28 0xdc33  36
NSSA     *172.30.38.0       172.30.5.5       0x80000004   221  0x28 0xde2c  36
NSSA      172.30.39.0       172.30.5.4       0x80000001   885  0x28 0xd13d  36
NSSA     *172.30.39.0       172.30.5.5       0x80000004    63  0x28 0xd336  36
NSSA      172.30.40.0       172.30.5.4       0x80000003   883  0x28 0xc249  36
NSSA     *172.30.40.0       172.30.5.5       0x80000003  2589  0x28 0xca3f  36
NSSA      172.30.41.0       172.30.5.4       0x80000003   883  0x28 0xb753  36
NSSA     *172.30.41.0       172.30.5.5       0x80000003  2432  0x28 0xbf49  36
NSSA      172.30.42.0       172.30.5.4       0x80000003   883  0x28 0xac5d  36
NSSA     *172.30.42.0       172.30.5.5       0x80000003  2274  0x28 0xb453  36
NSSA      172.30.43.0       172.30.5.4       0x80000003   883  0x28 0xa167  36
NSSA     *172.30.43.0       172.30.5.5       0x80000003  2116  0x28 0xa95d  36
NSSA      172.30.44.0       172.30.5.4       0x80000003   883  0x28 0x9671  36
NSSA     *172.30.44.0       172.30.5.5       0x80000003  1958  0x28 0x9e67  36
NSSA      172.30.45.0       172.30.5.4       0x80000003   883  0x28 0x8b7b  36
NSSA     *172.30.45.0       172.30.5.5       0x80000003  1800  0x28 0x9371  36
NSSA      172.30.46.0       172.30.5.4       0x80000003   883  0x28 0x8085  36
NSSA     *172.30.46.0       172.30.5.5       0x80000003  1642  0x28 0x887b  36
NSSA      172.30.47.0       172.30.5.4       0x80000003   883  0x28 0x758f  36
NSSA     *172.30.47.0       172.30.5.5       0x80000003  1484  0x28 0x7d85  36
```

R5 is identical to R4. R5 has a connection to the VR and injects the same RIP routes into the OSPF domain.

f) R6

```
[edit]
lab@R6# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.2       172.30.5.2       0x80001710     1  0x22 0xf21b  60
Router   172.30.5.3       172.30.5.3       0x80000017   720  0x22 0xd136  60
Router   172.30.5.7       172.30.5.7       0x80000018   231  0x22 0x9b38  60
Network  172.30.0.14      172.30.5.3       0x80000002   355  0x22 0xbf2c  32
Network *172.30.0.17      172.30.5.2       0x80000402  3600  0x22 0xd70a  32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Network  172.30.0.25      172.30.5.3      0x80000001   720  0x22 0x538e  32
Network  172.30.0.42      172.30.5.7      0x80000008   697  0x22 0xaa17  32
Summary *172.30.0.0       172.30.5.2      0x80000bda  3600  0x22 0x4c71  28
Summary *172.30.0.20      172.30.5.2      0x80000216     1  0x22 0x3c39  28
Summary  172.30.0.20      172.30.5.3      0x8000000a   514  0x22 0x4044  28
Summary *172.30.0.28      172.30.5.2      0x80000bd2     3  0x22 0x4d5b  28
Summary  172.30.0.28      172.30.5.3      0x80000001   920  0x22 0xc78   28
Summary *172.30.0.32      172.30.5.2      0x80000bd3     3  0x22 0x198b  28
Summary  172.30.0.32      172.30.5.3      0x80000001   920  0x22 0xed91  28
Summary  172.30.0.44      172.30.5.7      0x8000000a   911  0x22 0x3731  28
Summary *172.30.5.4       172.30.5.2      0x80000219     1  0x22 0xa7d3  28
Summary  172.30.5.4       172.30.5.3      0x80000001   920  0x22 0xcdc7  28
Summary *172.30.5.5       172.30.5.2      0x80000bd2     3  0x22 0x5b3   28
Summary  172.30.5.5       172.30.5.3      0x80000001   920  0x22 0xcdc5  28

    OSPF database, Area 0.0.0.4
 Type       ID               Adv Rtr          Seq       Age  Opt  Cksum  Len
Router  *172.30.5.2       172.30.5.2      0x8000000c   719  0x20 0xf8b1  36
Router   172.30.5.3       172.30.5.3      0x8000000f   838  0x20 0xb30b  36
Router   172.30.5.4       172.30.5.4      0x80000005   886  0x20 0xb051  60
Router   172.30.5.5       172.30.5.5      0x80000016   924  0x20 0xcd08  60
Network  172.30.0.21      172.30.5.3      0x80000004    34  0x20 0xaf33  32
Network  172.30.0.30      172.30.5.5      0x80000003   916  0x20 0x5f77  32
Network  172.30.0.33      172.30.5.5      0x80000007  1366  0x20 0x1db4  32
Summary *0.0.0.0          172.30.5.2      0x80000007    39  0x20 0xcd8f  28
Summary  172.30.0.12      172.30.5.3      0x8000000c   200  0x20 0xaae1  28
Summary  172.30.0.16      172.30.5.3      0x80000001  3600  0x20 0xa2ef  28
Summary  172.30.0.24      172.30.5.3      0x8000000a   722  0x20 0x364c  28
Summary  172.30.0.40      172.30.5.3      0x80000005  3600  0x20 0xa9cc  28
Summary  172.30.0.44      172.30.5.3      0x80000005  3600  0x20 0x8be5  28
Summary  172.30.5.2       172.30.5.3      0x8000002d     5  0x20 0xa7c5  28
Summary  172.30.5.3       172.30.5.3      0x80000009   673  0x20 0xdbb5  28
Summary  172.30.5.6       172.30.5.3      0x80000004  3600  0x20 0xd1c0  28
Summary  172.30.5.7       172.30.5.3      0x80000005  3600  0x20 0xcfbf  28
NSSA     172.30.32.0      172.30.5.4      0x80000002   682  0x28 0x1df7  36
NSSA     172.30.32.0      172.30.5.5      0x80000004  1208  0x28 0x21ef  36
NSSA     172.30.33.0      172.30.5.4      0x80000002   506  0x28 0x1202  36
NSSA     172.30.33.0      172.30.5.5      0x80000004  1051  0x28 0x16f9  36
NSSA     172.30.34.0      172.30.5.4      0x80000002   331  0x28 0x70c   36
NSSA     172.30.34.0      172.30.5.5      0x80000004   893  0x28 0xb04   36
NSSA     172.30.35.0      172.30.5.4      0x80000002   157  0x28 0xfb16  36
NSSA     172.30.35.0      172.30.5.5      0x80000004   735  0x28 0xff0e  36
NSSA     172.30.36.0      172.30.5.4      0x80000001   926  0x28 0xf21f  36
NSSA     172.30.36.0      172.30.5.5      0x80000004   577  0x28 0xf418  36
NSSA     172.30.37.0      172.30.5.4      0x80000001   926  0x28 0xe729  36
NSSA     172.30.37.0      172.30.5.5      0x80000004   420  0x28 0xe922  36
NSSA     172.30.38.0      172.30.5.4      0x80000001   926  0x28 0xdc33  36
NSSA     172.30.38.0      172.30.5.5      0x80000004   261  0x28 0xde2c  36
NSSA     172.30.39.0      172.30.5.4      0x80000001   926  0x28 0xd13d  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
NSSA     172.30.39.0      172.30.5.5       0x80000004   103  0x28 0xd336  36
NSSA     172.30.40.0      172.30.5.4       0x80000003   924  0x28 0xc249  36
NSSA     172.30.40.0      172.30.5.5       0x80000003  2630  0x28 0xca3f  36
NSSA     172.30.41.0      172.30.5.4       0x80000003   924  0x28 0xb753  36
NSSA     172.30.41.0      172.30.5.5       0x80000003  2472  0x28 0xbf49  36
NSSA     172.30.42.0      172.30.5.4       0x80000003   924  0x28 0xac5d  36
NSSA     172.30.42.0      172.30.5.5       0x80000003  2314  0x28 0xb453  36
NSSA     172.30.43.0      172.30.5.4       0x80000003   924  0x28 0xa167  36
NSSA     172.30.43.0      172.30.5.5       0x80000003  2156  0x28 0xa95d  36
NSSA     172.30.44.0      172.30.5.4       0x80000003   924  0x28 0x9671  36
NSSA     172.30.44.0      172.30.5.5       0x80000003  1998  0x28 0x9e67  36
NSSA     172.30.45.0      172.30.5.4       0x80000003   924  0x28 0x8b7b  36
NSSA     172.30.45.0      172.30.5.5       0x80000003  1840  0x28 0x9371  36
NSSA     172.30.46.0      172.30.5.4       0x80000003   924  0x28 0x8085  36
NSSA     172.30.46.0      172.30.5.5       0x80000003  1682  0x28 0x887b  36
NSSA     172.30.47.0      172.30.5.4       0x80000003   924  0x28 0x758f  36
NSSA     172.30.47.0      172.30.5.5       0x80000003  1524  0x28 0x7d85  36
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   172.30.32.0      172.30.5.3       0x80000001   920  0x22 0xcf4   36
```

R6 is an ABR between Area 0 and Area 4, the same role as R3. However, looking at its LSAs you can see that it is originating an LSA Type 1 with ID `172.30.5.2` and three LSAs Type 1 in Area 0.

It is stated that the Router IDs of the routers must be unique throughout the network. Verify the configuration on R3 reveals that the Router ID is configured manually and is the same as the loopback IP address.

```
[edit]
lab@R3# show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.3/32;
    }
    family inet6 {
        address fd17:f0f4:f691:5::3/128;
    }
}

[edit]
lab@R3# show routing-options
static {
    route 10.10.10.0/24 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        next-hop 10.10.1.254;
        no-readvertise;
    }
}
router-id 172.30.5.3;
```

Note: The `router-id` is configured under the `routing-options` hierarchy and is used for whole router.

Check the same command outputs on R6.

```
[edit]
lab@R6# show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.6/32;
    }
    family inet6 {
        address fd17:f0f4:f691:5::6/128;
    }
}
```

```
[edit]
lab@R6# show routing-options
static {
    route 10.10.10.0/24 {
        next-hop 10.10.1.254;
        no-readvertise;
    }
}
router-id 172.30.5.2;
```

The `router-id` is different than the IP address configured on the loopback interface, which is a valid configuration. However, if we look in the database it could be a problem. The configured Router-id is the same as the loopback IP address (and possible the Router-ID) of router R2.

```
[edit]
lab@R2# show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        address 172.30.5.2/32;
    }
    family inet6 {
        address fd17:f0f4:f691:5::2/128;
    }
}

[edit]
lab@R2# show routing-options
static {
    route 10.10.10.0/24 {
        next-hop 10.10.1.254;
        no-readvertise;
    }
}
router-id 172.30.5.2;
```

It is obvious that the `router-id` on R2 is the same as the one configured on R6. Following the same paradigm for the other routers you should change the `router-id` on R6 to reference the IP address configured on loopback interface.

```
[edit]
lab@R6# set routing-options router-id 172.30.5.6

[edit]
lab@R6# commit
commit complete
```

Fixing this error changes the OSPF database significantly.

```
lab@R6# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.2       172.30.5.2      0x80001848    91  0x22 0x7f55  60
Router   172.30.5.3       172.30.5.3      0x80000019    86  0x22 0xcd38  60
Router  *172.30.5.6       172.30.5.6      0x80000005    45  0x22 0x5c7c  60
Router   172.30.5.7       172.30.5.7      0x8000001a    86  0x22 0x973a  60
Network  172.30.0.14      172.30.5.3      0x80000002   724  0x22 0xbf2c  32
Network  172.30.0.17      172.30.5.2      0x800004a5    84  0x22 0x90ad  32
Network  172.30.0.25      172.30.5.3      0x80000003    86  0x22 0x8754  32
Network  172.30.0.42      172.30.5.7      0x8000000a    86  0x22 0xdedc  32
Summary  172.30.0.0       172.30.5.2      0x80000c7b    87  0x22 0x814   28
Summary  172.30.0.20      172.30.5.3      0x8000000a   883  0x22 0x4044  28
Summary *172.30.0.20      172.30.5.6      0x80000001    80  0x22 0x5434  28
Summary  172.30.0.28      172.30.5.3      0x80000002   234  0x22 0xa79   28
Summary *172.30.0.28      172.30.5.6      0x80000001    80  0x22 0xf987  28
Summary  172.30.0.32      172.30.5.3      0x80000002    68  0x22 0xeb92  28
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Summary *172.30.0.32      172.30.5.6       0x80000002    80  0x22 0xc5b7  28
Summary  172.30.0.44      172.30.5.7       0x8000000b    44  0x22 0x3532  28
Summary  172.30.5.4       172.30.5.3       0x80000001  1287  0x22 0xcdc7  28
Summary *172.30.5.4       172.30.5.6       0x80000001    80  0x22 0xc5cb  28
Summary  172.30.5.5       172.30.5.3       0x80000001  1287  0x22 0xcdc5  28
Summary *172.30.5.5       172.30.5.6       0x80000001    80  0x22 0xb1df  28

    OSPF database, Area 0.0.0.4
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.2       172.30.5.2       0x8000000c  1088  0x20 0xf8b1  36
Router   172.30.5.3       172.30.5.3       0x8000000f  1205  0x20 0xb30b  36
Router   172.30.5.4       172.30.5.4       0x80000005  1252  0x20 0xb051  60
Router   172.30.5.5       172.30.5.5       0x80000018    86  0x20 0xc90a  60
Router  *172.30.5.6       172.30.5.6       0x80000004    80  0x20 0xc8e1  36
Network  172.30.0.21      172.30.5.3       0x80000004   400  0x20 0xaf33  32
Network  172.30.0.30      172.30.5.5       0x80000003  1283  0x20 0x5f77  32
Network  172.30.0.33      172.30.5.5       0x80000009    86  0x20 0x517a  32
Summary  0.0.0.0         172.30.5.2       0x80000007   407  0x20 0xcd8f  28
Summary *0.0.0.0         172.30.5.6       0x80000001    85  0x20 0xc19d  28
Summary  172.30.0.0      172.30.5.3       0x80000001    88  0x20 0x435f  28
Summary  172.30.0.12     172.30.5.3       0x8000000c   567  0x20 0xaae1  28
Summary  172.30.0.16     172.30.5.3       0x80000001    83  0x20 0xa2ef  28
Summary  172.30.0.24     172.30.5.3       0x8000000c    87  0x20 0x324e  28
Summary  172.30.0.40     172.30.5.3       0x80000002    83  0x20 0xafc9  28
Summary  172.30.0.44     172.30.5.3       0x80000002    83  0x20 0x91e2  28
Summary  172.30.5.2      172.30.5.3       0x80000036    95  0x20 0x95ce  28
Summary  172.30.5.3      172.30.5.3       0x80000009  1040  0x20 0xdbb5  28
Summary  172.30.5.6      172.30.5.3       0x80000002    83  0x20 0xd5be  28
Summary  172.30.5.7      172.30.5.3       0x80000002    83  0x20 0xd5bc  28
NSSA     172.30.32.0     172.30.5.4       0x80000002  1048  0x28 0x1df7  36
NSSA     172.30.32.0     172.30.5.5       0x80000004  1575  0x28 0x21ef  36
NSSA     172.30.33.0     172.30.5.4       0x80000002   873  0x28 0x1202  36
NSSA     172.30.33.0     172.30.5.5       0x80000004  1417  0x28 0x16f9  36
NSSA     172.30.34.0     172.30.5.4       0x80000002   698  0x28 0x70c   36
NSSA     172.30.34.0     172.30.5.5       0x80000004  1259  0x28 0xb04   36
NSSA     172.30.35.0     172.30.5.4       0x80000002   523  0x28 0xfb16  36
NSSA     172.30.35.0     172.30.5.5       0x80000004  1101  0x28 0xff0e  36
NSSA     172.30.36.0     172.30.5.4       0x80000002   349  0x28 0xf020  36
NSSA     172.30.36.0     172.30.5.5       0x80000004   943  0x28 0xf418  36
NSSA     172.30.37.0     172.30.5.4       0x80000002   174  0x28 0xe52a  36
NSSA     172.30.37.0     172.30.5.5       0x80000004   786  0x28 0xe922  36
NSSA     172.30.38.0     172.30.5.4       0x80000001  1292  0x28 0xdc33  36
NSSA     172.30.38.0     172.30.5.5       0x80000004   629  0x28 0xde2c  36
NSSA     172.30.39.0     172.30.5.4       0x80000001  1293  0x28 0xd13d  36
NSSA     172.30.39.0     172.30.5.5       0x80000004   471  0x28 0xd336  36
NSSA     172.30.40.0     172.30.5.4       0x80000003  1291  0x28 0xc249  36
NSSA     172.30.40.0     172.30.5.5       0x80000004   313  0x28 0xc840  36
NSSA     172.30.41.0     172.30.5.4       0x80000003  1291  0x28 0xb753  36
NSSA     172.30.41.0     172.30.5.5       0x80000004   155  0x28 0xbd4a  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
NSSA     172.30.42.0      172.30.5.4       0x80000003  1291  0x28 0xac5d  36
NSSA     172.30.42.0      172.30.5.5       0x80000003  2682  0x28 0xb453  36
NSSA     172.30.43.0      172.30.5.4       0x80000003  1291  0x28 0xa167  36
NSSA     172.30.43.0      172.30.5.5       0x80000003  2524  0x28 0xa95d  36
NSSA     172.30.44.0      172.30.5.4       0x80000003  1291  0x28 0x9671  36
NSSA     172.30.44.0      172.30.5.5       0x80000003  2366  0x28 0x9e67  36
NSSA     172.30.45.0      172.30.5.4       0x80000003  1291  0x28 0x8b7b  36
NSSA     172.30.45.0      172.30.5.5       0x80000003  2208  0x28 0x9371  36
NSSA     172.30.46.0      172.30.5.4       0x80000003  1291  0x28 0x8085  36
NSSA     172.30.46.0      172.30.5.5       0x80000003  2050  0x28 0x887b  36
NSSA     172.30.47.0      172.30.5.4       0x80000003  1291  0x28 0x758f  36
NSSA     172.30.47.0      172.30.5.5       0x80000003  1892  0x28 0x7d85  36
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Extern   *172.30.32.0      172.30.5.6       0x80000001   81  0x22 0x9783  36
Extern   *172.30.33.0      172.30.5.6       0x80000001   81  0x22 0x8c8d  36
Extern   *172.30.34.0      172.30.5.6       0x80000001   81  0x22 0x8197  36
Extern   *172.30.35.0      172.30.5.6       0x80000001   81  0x22 0x76a1  36
Extern   *172.30.36.0      172.30.5.6       0x80000001   81  0x22 0x6bab  36
Extern   *172.30.37.0      172.30.5.6       0x80000001   81  0x22 0x60b5  36
Extern   *172.30.38.0      172.30.5.6       0x80000001   81  0x22 0x55bf  36
Extern   *172.30.39.0      172.30.5.6       0x80000001   81  0x22 0x4ac9  36
Extern   *172.30.40.0      172.30.5.6       0x80000001   81  0x22 0x3fd3  36
Extern   *172.30.41.0      172.30.5.6       0x80000001   81  0x22 0x34dd  36
Extern   *172.30.42.0      172.30.5.6       0x80000001   82  0x22 0x29e7  36
Extern   *172.30.43.0      172.30.5.6       0x80000001   82  0x22 0x1ef1  36
Extern   *172.30.44.0      172.30.5.6       0x80000001   82  0x22 0x13fb  36
Extern   *172.30.45.0      172.30.5.6       0x80000001   82  0x22 0x806   36
Extern   *172.30.46.0      172.30.5.6       0x80000001   82  0x22 0xfc10  36
Extern   *172.30.47.0      172.30.5.6       0x80000001   82  0x22 0xf11a  36
```

The output looks different now. There are four LSA Type 1 in Area 0, reflecting the topology in Figure 4. However, the situation in Area 4 also changed. The aggregated prefix for all RIP routes is missing. Instead all separate prefixes are advertised as External. Later you will see why the area 4 database entries changed.

5) Fix Area 4 LSA Types.

One of the requirements is that only LSA Type 1 and LSA Type 7 are in the database of Area 4. Thus, no LSA Type 2, 3, 4, 5 may exist inside Area 4.

List R4's database as it is an internal router for Area 4 you can see that LSA Type 1, 2, 3, and 7 are present.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show ospf database

    OSPF database, Area 0.0.0.4
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Router   172.30.5.2       172.30.5.2      0x8000000c   424  0x20 0xf8b1  36
Router   172.30.5.3       172.30.5.3      0x8000000f   539  0x20 0xb30b  36
Router  *172.30.5.4       172.30.5.4      0x80000005   586  0x20 0xb051  60
Router   172.30.5.5       172.30.5.5      0x80000016   626  0x20 0xcd08  60
Network  172.30.0.21      172.30.5.3      0x80000003   621  0x20 0xb132  32
Network  172.30.0.30      172.30.5.5      0x80000003   619  0x20 0x5f77  32
Network  172.30.0.33      172.30.5.5      0x80000007  1069  0x20 0x1db4  32
Summary  0.0.0.0          172.30.5.2      0x80000006  1184  0x20 0xcf8e  28
Summary  172.30.0.12      172.30.5.3      0x8000000b   896  0x20 0xace0  28
Summary  172.30.0.16      172.30.5.3      0x80000001     2  0x20 0xa2ef  28
Summary  172.30.0.24      172.30.5.3      0x8000000a   423  0x20 0x364c  28
Summary  172.30.0.40      172.30.5.3      0x80000002     2  0x20 0xb9be  28
Summary  172.30.0.44      172.30.5.3      0x80000002     2  0x20 0x91e2  28
Summary  172.30.5.2       172.30.5.3      0x8000001c    11  0x20 0xc9b4  28
Summary  172.30.5.3       172.30.5.3      0x80000009   374  0x20 0xdbb5  28
Summary  172.30.5.7       172.30.5.3      0x80000002     2  0x20 0xd5bc  28
NSSA    *172.30.32.0      172.30.5.4      0x80000002   382  0x28 0x1df7  36
NSSA     172.30.32.0      172.30.5.5      0x80000004   911  0x28 0x21ef  36
NSSA    *172.30.33.0      172.30.5.4      0x80000002   207  0x28 0x1202  36
NSSA     172.30.33.0      172.30.5.5      0x80000004   753  0x28 0x16f9  36
NSSA    *172.30.34.0      172.30.5.4      0x80000002    32  0x28 0x70c   36
NSSA     172.30.34.0      172.30.5.5      0x80000004   595  0x28 0xb04   36
NSSA    *172.30.35.0      172.30.5.4      0x80000001   626  0x28 0xfd15  36
NSSA     172.30.35.0      172.30.5.5      0x80000004   438  0x28 0xff0e  36
NSSA    *172.30.36.0      172.30.5.4      0x80000001   626  0x28 0xf21f  36
NSSA     172.30.36.0      172.30.5.5      0x80000004   280  0x28 0xf418  36
NSSA    *172.30.37.0      172.30.5.4      0x80000001   626  0x28 0xe729  36
NSSA     172.30.37.0      172.30.5.5      0x80000004   123  0x28 0xe922  36
NSSA    *172.30.38.0      172.30.5.4      0x80000001   626  0x28 0xdc33  36
NSSA     172.30.38.0      172.30.5.5      0x80000003  2648  0x28 0xe02b  36
NSSA    *172.30.39.0      172.30.5.4      0x80000001   626  0x28 0xd13d  36
NSSA     172.30.39.0      172.30.5.5      0x80000003  2490  0x28 0xd535  36
NSSA    *172.30.40.0      172.30.5.4      0x80000003   625  0x28 0xc249  36
NSSA     172.30.40.0      172.30.5.5      0x80000003  2332  0x28 0xca3f  36
NSSA    *172.30.41.0      172.30.5.4      0x80000003   625  0x28 0xb753  36
NSSA     172.30.41.0      172.30.5.5      0x80000003  2175  0x28 0xbf49  36
NSSA    *172.30.42.0      172.30.5.4      0x80000003   625  0x28 0xac5d  36
NSSA     172.30.42.0      172.30.5.5      0x80000003  2017  0x28 0xb453  36
NSSA    *172.30.43.0      172.30.5.4      0x80000003   625  0x28 0xa167  36
NSSA     172.30.43.0      172.30.5.5      0x80000003  1859  0x28 0xa95d  36
NSSA    *172.30.44.0      172.30.5.4      0x80000003   625  0x28 0x9671  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
NSSA     172.30.44.0      172.30.5.5       0x80000003  1701  0x28 0x9e67  36
NSSA    *172.30.45.0      172.30.5.4       0x80000003   625  0x28 0x8b7b  36
NSSA     172.30.45.0      172.30.5.5       0x80000003  1543  0x28 0x9371  36
NSSA    *172.30.46.0      172.30.5.4       0x80000003   625  0x28 0x8085  36
NSSA     172.30.46.0      172.30.5.5       0x80000003  1385  0x28 0x887b  36
NSSA    *172.30.47.0      172.30.5.4       0x80000003   625  0x28 0x758f  36
NSSA     172.30.47.0      172.30.5.5       0x80000003  1227  0x28 0x7d85  36
```

As you learned in the overview, LSAs type 2 are only generated on multi-access broadcast segments. All the interfaces in our topology have only two peers on the segment. In other words, they are back-to-back connected. You can manually configure the interface media type to point-to-point under OSPF. This ensures that no LSAs Type 2 will be generated. It also speeds up neighbor establishment, as no DR and BRD need to be elected.

Change all interfaces in Area 4 on R3, R4, R5 and R6 to point-to-point.

R3

```
[edit]
lab@R3# set protocols ospf area 4 interface ge-0/0/4.134 interface-type p2p

[edit]
lab@R3# commit
commit complete
```

R4

```
[edit]
lab@R4# set protocols ospf area 4 interface ge-0/0/4.134 interface-type p2p

[edit]
lab@R4# set protocols ospf area 4 interface ge-0/0/4.145 interface-type p2p

[edit]
lab@R4# commit
commit complete
```

R5

```
[edit]
lab@R5# set protocols ospf area 4 interface ge-0/0/4.145 interface-type p2p

[edit]
lab@R5# set protocols ospf area 4 interface ae0.0 interface-type p2p
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# commit
commit complete
```

R6

```
[edit]
lab@R6# set protocols ospf area 4 interface ae0.0 interface-type p2p

[edit]
lab@R6# commit
commit complete
```

Check the results of the interface type change on R5.

```
[edit]
lab@R5# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ae0.0              PtToPt  0.0.0.4         0.0.0.0         0.0.0.0            1
ge-0/0/4.145       PtToPt  0.0.0.4         0.0.0.0         0.0.0.0            1
lo0.0              DR      0.0.0.4         172.30.5.5      0.0.0.0            0
```

Display the database on R4.

```
[edit]
lab@R4# run show ospf database area 4

    OSPF database, Area 0.0.0.4
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.3       172.30.5.3      0x8000001f   151  0x20 0x2ba3  48
Router  *172.30.5.4       172.30.5.4      0x80000010   149  0x20 0xf542  84
Router   172.30.5.5       172.30.5.5      0x8000005f     1  0x20 0x4784  84
Router   172.30.5.6       172.30.5.6      0x8000000d     2  0x20 0xe8d7  48
Summary  0.0.0.0          172.30.5.6      0x80000003   156  0x20 0xbd9f  28
Summary  172.30.0.0       172.30.5.3      0x80000003    61  0x20 0x3f61  28
Summary  172.30.0.12      172.30.5.3      0x8000001b   157  0x20 0x8cf0  28
Summary  172.30.0.16      172.30.5.3      0x80000002   157  0x20 0xa0f0  28
Summary  172.30.0.24      172.30.5.3      0x8000001c   157  0x20 0x125e  28
Summary  172.30.0.32      172.30.5.3      0x80000001  3600  0x20 0x280   28
Summary  172.30.0.40      172.30.5.3      0x80000009   157  0x20 0xa1d0  28
Summary  172.30.0.44      172.30.5.3      0x80000009   157  0x20 0x83e9  28
Summary  172.30.5.2       172.30.5.3      0x8000007a   157  0x20 0xd13   28
Summary  172.30.5.3       172.30.5.3      0x80000019   157  0x20 0xbbc5  28
Summary  172.30.5.6       172.30.5.3      0x80000008   157  0x20 0xc9c4  28
Summary  172.30.5.7       172.30.5.3      0x80000009   157  0x20 0xc7c3  28
NSSA      0.0.0.0         172.30.5.5      0x80000001     9  0x28 0xc03d  36
NSSA     *172.30.32.0     172.30.5.4      0x80000004   155  0x28 0x19f9  36
NSSA      172.30.32.0     172.30.5.5      0x80000018   156  0x28 0xf804  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
NSSA    *172.30.33.0      172.30.5.4      0x80000004   155  0x28 0xe04    36
NSSA     172.30.33.0      172.30.5.5      0x80000018   156  0x28 0xed0e   36
NSSA    *172.30.34.0      172.30.5.4      0x80000004   155  0x28 0x30e    36
NSSA     172.30.34.0      172.30.5.5      0x80000018   156  0x28 0xe218   36
NSSA    *172.30.35.0      172.30.5.4      0x80000004   155  0x28 0xf718   36
NSSA     172.30.35.0      172.30.5.5      0x80000018   156  0x28 0xd722   36
NSSA    *172.30.36.0      172.30.5.4      0x80000004   155  0x28 0xec22   36
NSSA     172.30.36.0      172.30.5.5      0x80000018   156  0x28 0xcc2c   36
NSSA    *172.30.37.0      172.30.5.4      0x80000004   155  0x28 0xe12c   36
NSSA     172.30.37.0      172.30.5.5      0x80000018   156  0x28 0xc136   36
NSSA    *172.30.38.0      172.30.5.4      0x80000004   155  0x28 0xd636   36
NSSA     172.30.38.0      172.30.5.5      0x80000018   156  0x28 0xb640   36
NSSA    *172.30.39.0      172.30.5.4      0x80000004   155  0x28 0xcb40   36
NSSA     172.30.39.0      172.30.5.5      0x80000018   156  0x28 0xab4a   36
NSSA    *172.30.40.0      172.30.5.4      0x80000004   155  0x28 0xc04a   36
NSSA     172.30.40.0      172.30.5.5      0x80000017   156  0x28 0xa253   36
NSSA    *172.30.41.0      172.30.5.4      0x80000004   155  0x28 0xb554   36
NSSA     172.30.41.0      172.30.5.5      0x80000017   156  0x28 0x975d   36
NSSA    *172.30.42.0      172.30.5.4      0x80000004   155  0x28 0xaa5e   36
NSSA     172.30.42.0      172.30.5.5      0x80000017   156  0x28 0x8c67   36
NSSA    *172.30.43.0      172.30.5.4      0x80000004   155  0x28 0x9f68   36
NSSA     172.30.43.0      172.30.5.5      0x80000017   156  0x28 0x8171   36
NSSA    *172.30.44.0      172.30.5.4      0x80000004   155  0x28 0x9472   36
NSSA     172.30.44.0      172.30.5.5      0x80000017   156  0x28 0x767b   36
NSSA    *172.30.45.0      172.30.5.4      0x80000003   155  0x28 0x8b7b   36
NSSA     172.30.45.0      172.30.5.5      0x80000017   156  0x28 0x6b85   36
NSSA    *172.30.46.0      172.30.5.4      0x80000003   155  0x28 0x8085   36
NSSA     172.30.46.0      172.30.5.5      0x80000017   156  0x28 0x608f   36
NSSA    *172.30.47.0      172.30.5.4      0x80000003   155  0x28 0x758f   36
NSSA     172.30.47.0      172.30.5.5      0x80000017   156  0x28 0x5599   36
```

```
[edit]
lab@R4# run show ospf database area 4 summary
Area 0.0.0.4:
    4 Router LSAs
    15 Summary LSAs
    33 NSSA LSAs
Externals:
Interface ge-0/0/4.134:
Area 0.0.0.4:
Interface ge-0/0/4.145:
Area 0.0.0.4:
Interface lo0.0:
Area 0.0.0.4:
```

There are no LSAs Type 2 present anymore in the link state database.

© 2020 Juniper Networks, Inc. All rights reserved.

Now you should remove Summary LSA Type 3 from the database. Not So Stubby Areas do not allow Type 5 and 4 LSAs to be flooded in the Area. Adding the `no-summary` option under the `nssa` hierarchy transforms the NSSA into a Totally NSSA. This option instructs the ABR to not flood LSAs Type 3 inside a totally NSSA except one and only one Summary LSA for the default route. The generation of this LSA Type 3 should be manually configured. Junos provides an option of originating that LSA as a Type 7 instead of LSA Type 3. This way only LSAs Type 1 and 7 will be flooded within Area 4.

```
[edit]
lab@R3# set protocols ospf area 4 nssa default-lsa default-metric 10 type-7

[edit]
lab@R3# set protocols ospf area 4 nssa no-summaries

[edit]
lab@R3# commit
commit complete
```

Note: The metric for the default LSA can be set with the `default-metric` option.

On R6, the NSSA is already configured as `no-summary` and a default LSA is generated. However, you need to change the type of the summary to NSSA Type 7.

```
[edit]
lab@R6# show protocols ospf area 4 nssa
default-lsa default-metric 10;
no-summaries;

[edit]
lab@R6# set protocols ospf area 4 nssa default-lsa type-7

[edit]
lab@R6# commit
commit complete
```

Checking the content of the database in NSSA 4.

```
[edit]
lab@R4# run show ospf database

    OSPF database, Area 0.0.0.4
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.3       172.30.5.3       0x80000025    75  0x20 0x1fa9  48
Router  *172.30.5.4       172.30.5.4       0x80000019    74  0x20 0xe34b  84
Router   172.30.5.5       172.30.5.5       0x80000065   448  0x20 0x3b8a  84
Router   172.30.5.6       172.30.5.6       0x80000012   448  0x20 0xdedc  48
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
NSSA      0.0.0.0          172.30.5.3       0x80000001    78  0x20 0xabaa  36
NSSA      0.0.0.0          172.30.5.5       0x80000001    13  0x28 0xc03d  36
NSSA      0.0.0.0          172.30.5.6       0x80000001    14  0x20 0x99b9  36
NSSA     *172.30.32.0      172.30.5.4       0x80000005  1232  0x28 0x17fa  36
NSSA      172.30.32.0      172.30.5.5       0x80000019  1237  0x28 0xf605  36
NSSA     *172.30.33.0      172.30.5.4       0x80000005  1067  0x28 0xc05   36
NSSA      172.30.33.0      172.30.5.5       0x80000019  1062  0x28 0xeb0f  36
NSSA     *172.30.34.0      172.30.5.4       0x80000005   901  0x28 0x10f   36
NSSA      172.30.34.0      172.30.5.5       0x80000019   887  0x28 0xe019  36
NSSA     *172.30.35.0      172.30.5.4       0x80000005   736  0x28 0xf519  36
NSSA      172.30.35.0      172.30.5.5       0x80000019   713  0x28 0xd523  36
NSSA     *172.30.36.0      172.30.5.4       0x80000005   571  0x28 0xea23  36
NSSA      172.30.36.0      172.30.5.5       0x80000019   538  0x28 0xca2d  36
NSSA     *172.30.37.0      172.30.5.4       0x80000005   406  0x28 0xdf2d  36
NSSA      172.30.37.0      172.30.5.5       0x80000019   363  0x28 0xbf37  36
NSSA     *172.30.38.0      172.30.5.4       0x80000005   240  0x28 0xd437  36
NSSA      172.30.38.0      172.30.5.5       0x80000019   188  0x28 0xb441  36
NSSA     *172.30.39.0      172.30.5.4       0x80000005    75  0x28 0xc941  36
NSSA      172.30.39.0      172.30.5.5       0x80000019    13  0x28 0xa94b  36
NSSA     *172.30.40.0      172.30.5.4       0x80000004  1593  0x28 0xc04a  36
NSSA      172.30.40.0      172.30.5.5       0x80000017  1594  0x28 0xa253  36
NSSA     *172.30.41.0      172.30.5.4       0x80000004  1593  0x28 0xb554  36
NSSA      172.30.41.0      172.30.5.5       0x80000017  1594  0x28 0x975d  36
NSSA     *172.30.42.0      172.30.5.4       0x80000004  1593  0x28 0xaa5e  36
NSSA      172.30.42.0      172.30.5.5       0x80000017  1594  0x28 0x8c67  36
NSSA     *172.30.43.0      172.30.5.4       0x80000004  1593  0x28 0x9f68  36
NSSA      172.30.43.0      172.30.5.5       0x80000017  1594  0x28 0x8171  36
NSSA     *172.30.44.0      172.30.5.4       0x80000004  1593  0x28 0x9472  36
NSSA      172.30.44.0      172.30.5.5       0x80000017  1594  0x28 0x767b  36
NSSA     *172.30.45.0      172.30.5.4       0x80000003  1593  0x28 0x8b7b  36
NSSA      172.30.45.0      172.30.5.5       0x80000017  1594  0x28 0x6b85  36
NSSA     *172.30.46.0      172.30.5.4       0x80000003  1593  0x28 0x8085  36
NSSA      172.30.46.0      172.30.5.5       0x80000017  1594  0x28 0x608f  36
NSSA     *172.30.47.0      172.30.5.4       0x80000003  1593  0x28 0x758f  36
NSSA      172.30.47.0      172.30.5.5       0x80000017  1594  0x28 0x5599  36

[edit]
lab@R4# run show ospf database summary
Area 0.0.0.4:
   4 Router LSAs
   35 NSSA LSAs
Externals:
Interface ge-0/0/4.134:
Area 0.0.0.4:
Interface ge-0/0/4.145:
Area 0.0.0.4:
Interface lo0.0:
Area 0.0.0.4:
```


© 2020 Juniper Networks, Inc. All rights reserved.

6) Verify the RIP configuration and OSPF – RIP redistribution.

R4 and R5 are redistributing RIP routes into the OSPF domain. R4 and R5 are internal routers in Area 4, which is configured as a NSSA. This means that the RIP routes are originated as Type 7 LSA.

RIP is a distance-vector Interior Gateway Protocol. It calculates the shortest path by evaluating the distance to reach the prefix (vector). The distance used by RIP is the hop count. Routing updates are sent at a regular basis, which is 30 seconds by default. There are two versions of the RIP protocol, where version 2 is the enhanced version. Version 2 enhancements includes multicast routing updates, variable length subnet masks support and message authentication.

In Junos by default all prefixes received from other RIP neighbors are automatically accepted but nothing is advertised. To advertise a route to another RIP neighbor you need to apply routing policy.

Let's display the current RIP configuration on R4 and R5.

```
[edit]
lab@R4# show protocols rip
group rip {
    export default-to-rip;
    neighbor ge-0/0/4.202;
}
```

```
[edit]
lab@R5# show protocols rip
group rip {
    export default-to-rip;
    neighbor ge-0/0/4.204;
}
```

The configuration is pretty simple. Now let's check the protocol operation on both routers.

a) R4

```
[edit]
lab@R4# run show rip neighbor
                Local  Source          Destination     Send   Receive   In
Neighbor        State  Address         Address         Mode   Mode      Met
--------        -----  -------         -----------     ----   -------   ---
ge-0/0/4.202     Up    172.30.0.49     224.0.0.9       mcast  both        1
```


© 2020 Juniper Networks, Inc. All rights reserved.

The RIP neighbor is Up on interface ge-0/0/4.202. The destination address is the default multicast address used by RIP.

```
[edit]
lab@R4# run show route advertising-protocol rip 172.30.0.49

inet.0: 39 destinations, 55 routes (39 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[OSPF/150] 00:31:12, metric 11, tag 0
                     > to 172.30.0.21 via ge-0/0/4.134
```

The command above is very useful and provides information on what is advertised by the local router to the RIP neighbors. As stated, you need routing policy to export routes from the routing table to RIP neighbors. The above command is not available for link state routing protocols like OSPF and IS-IS, because they do not advertise routes directly from the routing table. LSAs are originated from the link state database and stored there for processing.

Note: Because RIP does not maintain a neighbor state and uses multicasts to reach other neighbors, you should specify the IP address of the local RIP router when using the `show route advertising-protocol rip`.

```
[edit]
lab@R4# run show route 0/0 exact

inet.0: 37 destinations, 53 routes (39 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[OSPF/150] 00:09:10, metric 11, tag 0
                     > to 172.30.0.21 via ge-0/0/4.134
```

The default route is received by R4 from R3 via an OSPF Type 7 LSA. The output shown is from the local routing table.

```
[edit]
lab@R4# run show route receive-protocol rip 172.30.0.50

inet.0: 37 destinations, 53 routes (37 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                     > to 172.30.0.50 via ge-0/0/4.202
172.30.33.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                     > to 172.30.0.50 via ge-0/0/4.202
172.30.34.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.35.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.36.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.37.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.38.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.39.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.40.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.41.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.42.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.43.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.44.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.45.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.46.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.47.0/24      *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
```

The `receive-protocol` option provides information on what is received by the local router from the specified neighbor. The output contains all routes **before** import routing policy is applied.

Note: Here you need to specify the IP address of the neighbor sending the routes to local router. You should know the neighbor IP address as it is not seen from the neighbor table. You can find it using *traceoptions* for the RIP protocol.

```
[edit]
lab@R4# run show route protocol rip

inet.0: 37 destinations, 53 routes (37 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                     > to 172.30.0.50 via ge-0/0/4.202
172.30.33.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                     > to 172.30.0.50 via ge-0/0/4.202
172.30.34.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.35.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.36.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.37.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.38.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.39.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.40.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.41.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.42.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.43.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.44.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.45.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.46.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.47.0/24      *[RIP/100] 03:03:38, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
224.0.0.9/32        *[RIP/100] 00:33:31, metric 1
                       MultiRecv
```

The RIP routes in the OSPF database:

```
[edit]
lab@R4# run show ospf database

    OSPF database, Area 0.0.0.4
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.3       172.30.5.3      0x8000000a   490  0x20 0x558e  48
Router  *172.30.5.4       172.30.5.4      0x8000000e   489  0x20 0xf940  84
Router   172.30.5.5       172.30.5.5      0x80000014  1161  0x20 0xdd39  84
Router   172.30.5.6       172.30.5.6      0x80000009   150  0x20 0xf0d3  48
NSSA     0.0.0.0          172.30.5.3      0x80000004   495  0x20 0xa5ad  36
NSSA     0.0.0.0          172.30.5.6      0x80000004   434  0x20 0x93bc  36
NSSA    *172.30.32.0      172.30.5.4      0x80000004   872  0x28 0x19f9  36
NSSA     172.30.32.0      172.30.5.5      0x80000005   858  0x28 0x1ff0  36
NSSA    *172.30.33.0      172.30.5.4      0x80000004   709  0x28 0xe04   36
NSSA     172.30.33.0      172.30.5.5      0x80000005   690  0x28 0x14fa  36
NSSA    *172.30.34.0      172.30.5.4      0x80000004   545  0x28 0x30e   36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
NSSA     172.30.34.0      172.30.5.5       0x80000005   521  0x28 0x905    36
NSSA    *172.30.35.0      172.30.5.4       0x80000004   382  0x28 0xf718   36
NSSA     172.30.35.0      172.30.5.5       0x80000005   353  0x28 0xfd0f   36
NSSA    *172.30.36.0      172.30.5.4       0x80000004   218  0x28 0xec22   36
...
```

Routing policy is used to redistribute the RIP routes into OSPF (type 7).

Looking at R4 it seems like all requirements are accomplished.

b) R5

```
[edit]
lab@R5# run show rip neighbor
                Local  Source       Destination      Send   Receive   In
Neighbor        State  Address      Address          Mode   Mode      Met
--------        -----  -------      -----------      ----   -------   ---
ge-0/0/4.204       Up 172.30.0.57   224.0.0.9        mcast  both        1
```

R5 has one RIP neighbor.

```
[edit]
lab@R5# run show route advertising-protocol rip 172.30.0.57

[edit]
lab@R5#
```

No routes are advertised by R5. This contradicts with the requirement that an ASBR failure should not lead to connectivity problems. The reason of that will be explained below.

```
[edit]
lab@R5# run show route 0/0 exact

inet.0: 35 destinations, 52 routes (35 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0               *[RIP/100] 00:38:05, metric 3, tag 0
                         > to 172.30.0.58 via ge-0/0/4.204
                        [OSPF/150] 00:38:58, metric 11, tag 0
                         > to 172.30.0.34 via ae0.0
```

The default route on R5 is received from R6 via OSPF, but it is also received from the VR via RIP. Junos assigns a default route preference value to each route. Preference values depend on the source of the route, which routing protocol inserts them into the routing table. A lower preference value is preferred. The preference value and the protocol that imported the route in the routing table is displayed between brackets. The default

© 2020 Juniper Networks, Inc. All rights reserved.

preference value for RIP is 100, whereas external OSPF routes have value of 150. The RIP route is preferred and is marked as active with "*". Internal OSPF routes have a different preference value, which is 10. Keep in mind that the default route preference values can be changed in Junos.

```
[edit]
lab@R5# run show route receive-protocol rip 172.30.0.58

inet.0: 33 destinations, 50 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[RIP/100] 00:38:05, metric 3, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.32.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.33.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.34.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.35.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.36.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.37.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.38.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.39.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.40.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.41.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.42.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.43.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.44.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.45.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.46.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.47.0/24      *[RIP/100] 00:02:18, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
```

R5 receives the same routes from its RIP neighbor plus the default route.

© 2020 Juniper Networks, Inc. All rights reserved.

7) Fix suboptimal routing.

You saw that the default route on R5 is pointing to the VR router. R5 will send all traffic for which it has no specific route to the VR device, including the traffic to other Areas. R5 prefers the default route to the VR device, because of the lower default preference of the RIP routes compared to the default route preference of the OSPF external routes. The default route learned by OSPF is external, because it is learned as an LSA Type 7. One option for fixing the suboptimal routing is to manually configure the external OSPF preference, changing it to lower value than 100. But it is obvious that the default route received by R5 is the one that is advertised from R4 to the VR. You can prove this by removing the export policy from the RIP configuration on R4 and the default will not be advertised by the VR anymore. So, leaving that route in the routing table even with higher preference is not needed. Other option, which is recommended, in this case is to filter the received route by RIP with an import route policy.

Routing policies are a very powerful tool for manipulating information inserted in the routing table, modify default routing behavior and redistribute route information between different routing protocols. Routing policies are always performed from the perspective of the routing table. Import policy controls the routes that are imported to the routing table. Export policy is used for exporting routes from routing table to the routing protocols.

The policy structure is based on chained terms. You can define policy without terms, with just one single from-then sequence but using terms is recommended. Using different terms adds flexibility for building complex policies.

```
policy-options {
      policy-statement policy-name {
            term term-name {
                  from {
                        match-conditions;
                  }
                  then {
                        actions;
                  }
            }
            term term-name {
                  from {
                        match-conditions;
                  }
                        then {
                              actions;
                        }
                  }
            }
      }
```


© 2020 Juniper Networks, Inc. All rights reserved.

All active routes in the routing table are evaluated against every term and policy, if more than one policy is applied. At last the default policy is evaluated and the result is to *accept* or *reject* the routes. Please note that the default policies are different for each routing protocol.

The `from` statement is the match condition identifies the routes of interest. Junos provides numerous ways of matching routes. Throughout this book you will use many of them. The `then` statement defines the terminating action. An `accept`, `reject` and `next-policy` are the available terminating actions.

The `to` statement can be used as a match condition too. It is used to match routes which are send to one of the two IS-IS levels and to specific Routing Information Bases (RIB) (routing tables).

Note: Routing policies only evaluate the active routes in the routing table.

For distance vector protocols like RIP and BGP where the updates include the actual prefixes, policies can be applied in the import and export direction.

For link state routing protocols like OSPF where the database must be the same on all routers import policies cannot be applied to limiting the flooding of intra area routes. Import policies can be used for controlling inter area and external routes as they are exchanged in a distance vector fashion. Export policies can be used to redistribute routing information from other routing protocols into OSPF and to affect which inter area routes are flooded into the area by ABR.

The other link-state routing protocol IS-IS, accepts only export policies.

As stated, the default policy for RIP is to accept everything in the import direction, but nothing is exported by default. An additional export policy is needed for redistributing routes into RIP.

The default OSPF import policy accepts everything. The default OSPF export policy is to reject everything. Export policies are used to redistribute routes into OSPF. This can be done with manual configured export policies.

To fix the suboptimal routing you need to filter the default route received by RIP from VR. This is accomplished with an import routing policy applied to the RIP group. To accomplish the requirement that the default route must be available if any of the ABR fails, you have to apply this policy to both R4 and R5.

© 2020 Juniper Networks, Inc. All rights reserved.

R4

```
[edit]
lab@R4# set policy-options policy-statement rip-filter term 1 from protocol rip

[edit]
lab@R4# set policy-options policy-statement rip-filter term 1 from route-filter 0/0
exact

[edit]
lab@R4# set policy-options policy-statement rip-filter term 1 then reject

[edit]
lab@R4# set protocols rip group rip import rip-filter

[edit]
lab@R4# commit
commit complete
```

R5

```
[edit]
lab@R5# set policy-options policy-statement rip-filter term 1 from protocol rip

[edit]
lab@R5# set policy-options policy-statement rip-filter term 1 from route-filter 0/0
exact

[edit]
lab@R5# set policy-options policy-statement rip-filter term 1 then reject

[edit]
lab@R5# set protocols rip group rip import rip-filter

[edit]
lab@R5# commit
commit complete
```

The policy named `rip-filter` has one term matching exactly the `0/0` route learned from RIP. The policy is applied in the import direction to all neighbors in the RIP group. If a match is found the route is rejected and not imported into the routing table. The other routes are matched by the default routing policy and are accepted.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show route 0/0 exact

inet.0: 37 destinations, 53 routes (39 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 00:31:12, metric 11, tag 0
                    > to 172.30.0.21 via ge-0/0/4.134


[edit]
lab@R5# run show route 0/0 exact

inet.0: 33 destinations, 49 routes (35 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 00:56:23, metric 11, tag 0
                    > to 172.30.0.34 via ae0.0
```

Both routers have the OSPF default route received from the ABR.

```
lab@R4# run show route protocol rip

inet.0: 37 destinations, 53 routes (37 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.33.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.34.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.35.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.36.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.37.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.38.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.39.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.40.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.41.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.42.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.43.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.44.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.45.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.46.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.47.0/24     *[RIP/100] 00:01:37, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
```

```
[edit]
lab@R5# run show route protocol rip

inet.0: 33 destinations, 49 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.33.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.34.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.35.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.36.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.37.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.38.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.39.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.40.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.41.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.42.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.43.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.44.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.45.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.46.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.47.0/24     *[RIP/100] 00:02:18, metric 2, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
> to 172.30.0.58 via ge-0/0/4.204
```

The other routes received from the VR by RIP are accepted on both routers.

8) Advertise a default route to the VR-device (RIP router). R4 updates should be preferred.

Advertising routes to RIP neighbors requires export policy. Both R4 and R5 should advertise a default route to the VR.

```
[edit]
lab@R4# show protocols rip
group rip {
    export default-to-rip;
    import rip-filter;
    neighbor ge-0/0/4.202;
 }

[edit]
lab@R5# show protocols rip
group rip {
    export default-to-rip;
    import rip-filter;
    neighbor ge-0/0/4.204;
}
```

Export policies are already applied.

```
[edit]
lab@R4# show policy-options policy-statement default-to-rip
term 1 {
    from {
        protocol ospf;
        route-filter 0.0.0.0/0 exact;
    }
    then accept;
}
```

```
[edit]
lab@R5# show policy-options policy-statement default-to-rip
term 1 {
    from {
        protocol ospf;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        metric 5;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        accept;
    }
}
```

Both policies match the default route learned from OSPF and accept it for advertisement. The policy on R5 sets a higher metric to the default route. This will ensure that the VR will choose R4s advertisement until it is present as it is assigned a default metric (hop count) of 1.

9) Verify Area 4 summarization.

R4 and R5 are advertising the RIP routes, which they learned from the VR device into the OSPF domain.

```
[edit]
lab@R4# show protocols ospf
export rip-to-ospf;
area 0.0.0.4 {
    nssa;
    interface ge-0/0/4.134 {
        interface-type p2p;
        authentication {
                md5 1 key "$9$eGdWNbHkPTF/ZUi.5FAtevWx-w"; ## SECRET-DATA
            }
    }
    interface ge-0/0/4.145 {
        interface-type p2p;
        authentication {
        md5 1 key "$9$GGiPQ1IcrK8tuBEyKx7GDi.fz"; ## SECRET-DATA
                    }
              }
              interface lo0.0;
         }
```

```
[edit]
lab@R5# show protocols ospf
export rip-to-ospf;
        area 0.0.0.4 {
            nssa;
            interface ge-0/0/4.145 {
                interface-type p2p;
                authentication {
                    md5 1 key "$9$-PwoGF39t0IP5z6A0hc-VwgaU"; ## SECRET-DATA
                }
            }
            interface ae0.0 {
                interface-type p2p;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            authentication {
                md5 1 key "$9$gAJi.tpBRcln/A0IcvMgoJDHm"; ## SECRET-DATA
            }
        }
        interface lo0.0;
    }
```

R4 and R5 have the `rip-to-ospf` export policy applied for OSPF.

```
[edit]
lab@R4# show policy-options policy-statement rip-to-ospf
term 1 {
    from protocol rip;
    then accept;
}
```

The policy accepts all RIP routes for export into OSPF as Type 7 LSAs.

As you saw in Step 4 after fixing the duplicated Router IDs the summary route advertised from R3 disappeared. Instead, all more specific NSSA LSAs holding the RIP prefixes are advertised by R6 to Area 0.

```
[edit]
lab@R6# run show ospf database external

    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *172.30.32.0      172.30.5.6       0x80000001   116  0x22 0x9783  36
Extern  *172.30.33.0      172.30.5.6       0x80000001   116  0x22 0x8c8d  36
Extern  *172.30.34.0      172.30.5.6       0x80000001   116  0x22 0x8197  36
Extern  *172.30.35.0      172.30.5.6       0x80000001   116  0x22 0x76a1  36
Extern  *172.30.36.0      172.30.5.6       0x80000001   116  0x22 0x6bab  36
Extern  *172.30.37.0      172.30.5.6       0x80000001   116  0x22 0x60b5  36
Extern  *172.30.38.0      172.30.5.6       0x80000001   116  0x22 0x55bf  36
Extern  *172.30.39.0      172.30.5.6       0x80000001   116  0x22 0x4ac9  36
Extern  *172.30.40.0      172.30.5.6       0x80000001   116  0x22 0x3fd3  36
Extern  *172.30.41.0      172.30.5.6       0x80000001   116  0x22 0x34dd  36
Extern  *172.30.42.0      172.30.5.6       0x80000001   116  0x22 0x29e7  36
Extern  *172.30.43.0      172.30.5.6       0x80000001   116  0x22 0x1ef1  36
Extern  *172.30.44.0      172.30.5.6       0x80000001   116  0x22 0x13fb  36
Extern  *172.30.45.0      172.30.5.6       0x80000001   116  0x22 0x806   36
Extern  *172.30.46.0      172.30.5.6       0x80000001   116  0x22 0xfc10  36
Extern  *172.30.47.0      172.30.5.6       0x80000001   116  0x22 0xf11a  36
```

R3 and R6 are ABRs having interfaces in Area 0 and Area 4. OSPF elects one ABR, to translate NSSA Type 7 LSA into Type 5 and originates them in Area 0. OSPF choses the ABR with the highest Router ID. So R6 should be the ABR used for this purpose.

© 2020 Juniper Networks, Inc. All rights reserved.

Before you fixed the incorrect Router ID on R6, R3 had the highest Router ID and the correct summary route was advertised. As you saw R3 had the right `area-range` command configured doing that.

But as R6 has now the higher Router ID and thus the role of advertising the Type 7 LSA, the summary LSA Type 5 is not advertised.

10) Fix the Area 4 summarization.

Let's check R6 OSPF configuration if the area-range command is configured.

```
[edit]
lab@R6# show protocols ospf area 4
nssa {
    default-lsa {
        default-metric 10;
        type-7;
    }
    no-summaries;
}
area-range 172.30.32.0/20;
interface ae0.0 {
    authentication {
        md5 1 key "$9$PQ/teK8x-whSlMX-2gP5Qn9p"; ## SECRET-DATA
    }
}
```

The `area-range` command is configured, and the summary route looks correct. But as we stated before, to summarize the NSSA Type 7 LSAs to Type 5 LSAs the `area-range` command should be configured under the `nssa` hierarchy. Configured under the Area directly this command will summarize the only LSAs Type 5 but not LSAs Type 7 but applied under the `nssa` hierarchy only the NSSA Type 7 LSAs are summarized.

Fix the mistake.

First to remove the unneeded command.

```
[edit]
lab@R6# delete protocols ospf area 4 area-range 172.30.32.0/20

[edit]
lab@R6# commit
commit complete
```

Configure the `area-range` command under the correct `nssa` hierarchy.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols ospf area 4 nssa area-range 172.30.32.0/20

[edit]
lab@R6# commit
commit complete
```

```
[edit]
lab@R6# run show ospf database external

    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *172.30.32.0      172.30.5.6      0x80000003    14  0x22 0xf506  36
```

After the change, the summary route is originated by R6 and more specific routes are suppressed. You might need to wait a moment for this to take effect.

11) Verify loopback reachability.

a) R1

```
[edit]
lab@R1# run show route 172.30.5/24 terse

inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.30.5.1/32      D   0                        >lo0.0
* ? 172.30.5.2/32      O  10          1             >172.30.0.2
* ? 172.30.5.3/32      O  10          2             >172.30.0.2
* ? 172.30.5.4/32      O  10          3             >172.30.0.2
* ? 172.30.5.5/32      O  10          4             >172.30.0.2
* ? 172.30.5.6/32      O  10          3             >172.30.0.2
* ? 172.30.5.7/32      O  10          2             >172.30.0.2
* ? 172.30.5.8/32      O  10          1             >172.30.0.10
```

All loopback IP Addresses are present in the routing table.

b) R2

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show route 172.30.5/24 terse

inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.30.5.2/32      D   0                       >lo0.0
* ? 172.30.5.3/32      O  10          1            >172.30.0.14
* ? 172.30.5.4/32      O  10          2            >172.30.0.14
* ? 172.30.5.5/32      O  10          3            >172.30.0.14
                                                   172.30.0.18
* ? 172.30.5.6/32      O  10          2            >172.30.0.14
                                                   172.30.0.18
* ? 172.30.5.7/32      O  10          1            >172.30.0.18
```

The loopback IP addresses of R1 and R8 are missing from the R2 routing table.

c) R3

```
[edit]
lab@R3# run show route 172.30.5/24 terse

inet.0: 43 destinations, 43 routes (43 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.30.5.2/32      O  10          1            >172.30.0.13
* ? 172.30.5.3/32      D   0                       >lo0.0
* ? 172.30.5.4/32      O  10          1            >172.30.0.22
* ? 172.30.5.5/32      O  10          2            >172.30.0.22
* ? 172.30.5.6/32      O  10          1            >172.30.0.26
* ? 172.30.5.7/32      O  10          2            >172.30.0.13
                                                   172.30.0.26
```

R3 does not know these IP address also. This is obvious, because R2 and R7 do not advertise them.

d) R7

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show route 172.30.5/24 terse

inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.30.5.2/32      O  10          1             >172.30.0.17
* ? 172.30.5.3/32      O  10          2             >172.30.0.17
                                                    172.30.0.41
* ? 172.30.5.4/32      O  10          3             >172.30.0.17
                                                    172.30.0.41
* ? 172.30.5.5/32      O  10          2             >172.30.0.41
* ? 172.30.5.6/32      O  10          1             >172.30.0.41
* ? 172.30.5.7/32      D   0                        >lo0.0
```

On R7 the situation is the same as on R2, the loopback IP addresses from R1 and R8 are not present in the routing table.

e) R8

```
[edit]
lab@R8# run show route 172.30.5/24 terse

inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.30.5.1/32      O  10          1             >172.30.0.9
* ? 172.30.5.2/32      O  10          2             >172.30.0.45
* ? 172.30.5.3/32      O  10          3             >172.30.0.45
* ? 172.30.5.4/32      O  10          4             >172.30.0.45
* ? 172.30.5.5/32      O  10          3             >172.30.0.45
* ? 172.30.5.6/32      O  10          2             >172.30.0.45
* ? 172.30.5.7/32      O  10          1             >172.30.0.45
* ? 172.30.5.8/32      D   0                        >lo0.0
```

R8 as R1 have visibility to all 8 loopback IP addresses.

12) Fix R1 and R8 loopback IP address reachability issues.

Looking at Figure 4 it is clear that Area 2 does not have an ABR with an interface in Area 0. This fact is in contradiction with OSPF loop prevention mechanism. All inter area traffic should pass Area 0. On Figure 4 you have a design with a discontinued backbone Area 0. The ABRs, which have an interface in Area 0 will not accept Summary LSAs received from

© 2020 Juniper Networks, Inc. All rights reserved.

the non-backbone Area 2. You can see that the Summary LSAs are received by R2 but not used. The same is valid for R7.

```
[edit]
lab@R2# run show ospf database netsummary lsa-id 172.30.5.1

    OSPF database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Summary  172.30.5.1       172.30.5.1       0x80000003  1043  0x22 0xe9af  28

[edit]
lab@R2# run show ospf database netsummary lsa-id 172.30.5.8

    OSPF database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Summary  172.30.5.8       172.30.5.1       0x80000003     8  0x22 0xade3  28
```

```
[edit]
lab@R7# run show ospf database netsummary lsa-id 172.30.5.1

    OSPF database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Summary  172.30.5.1       172.30.5.8       0x80000005   152  0x22 0xc5c9  28

[edit]
lab@R7# run show ospf database netsummary lsa-id 172.30.5.8

    OSPF database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Summary  172.30.5.8       172.30.5.8       0x8000000a   564  0x22 0xc5c9  28
```

The Summary routes are available in the link state database. The loopback IP addresses are received as Summary LSA, but not as Router LSAs, because the loopback interfaces of R1 and R8 are configured under Area 2.

```
[edit]
lab@R1# show protocols
ospf {
    traceoptions {
        file ospf.log;
        flag error detail;
    }
    area 0.0.0.3 {
        interface ae0.0 {
            authentication {
                md5 1 key "$9$75dwgmPQn6AGD.5"; ## SECRET-DATA
            }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        }
    }
    area 0.0.0.2 {
        interface ge-0/0/4.118 {
            authentication {
                md5 1 key "$9$TF6ArlMXxdBIyK"; ## SECRET-DATA
            }
        }
        interface lo0.0;
    }
}
```

Although is not a recommended design to have discontinuous areas, sometimes it is needed for migration purposes. OSPF provides a virtual connection mechanism (virtual links) to connect areas to the backbone area through other non-backbone areas.

You can create a virtual link from R1 to R2 and from R8 to R7. This way both R1 and R8 will have a virtual interface in Area 0 and will be able to send summary routes for Area 2 directly to Area 0 through the virtual link.

a) R1

```
[edit]
lab@R1# set protocols ospf area 0 virtual-link neighbor-id 172.30.5.2 transit-area 3

[edit]
lab@R1# commit
commit complete
```

Note: The virtual link is configured under Area 0. You have to specify the Router ID of the ABR and the Area, which is between the non-backbone Area and Area 0. The OSPF configuration on R1 looks like any other ABR, but under Area 0 the only interface is the virtual link itself.

b) R2

```
[edit]
lab@R2# set protocols ospf area 0 virtual-link neighbor-id 172.30.5.1 transit-area 3

[edit]
lab@R2# commit
commit complete
```

c) R8

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set protocols ospf area 0 virtual-link neighbor-id 172.30.5.7 transit-area 1

[edit]
lab@R8# commit
commit complete
```

d) R7

```
[edit]
lab@R7# set protocols ospf area 0 virtual-link neighbor-id 172.30.5.8 transit-area 1

[edit]
lab@R7# commit
commit complete
```

Verify the result of the virtual link configuration on R1.

```
[edit]
lab@R1# show protocols ospf
traceoptions {
    file ospf.log;
    flag error detail;
}
area 0.0.0.3 {
    interface ae0.0 {
        authentication {
            md5 1 key "$9$eGdWNbHkPTF/ZUi.5FAtevWx-w"; ## SECRET-DATA
        }
    }
}
area 0.0.0.2 {
    interface ge-0/0/4.118 {
        authentication {
            md5 1 key "$9$ecsWNbHkPTF/ZUi.5FAtevWx-w"; ## SECRET-DATA
        }
    }
    interface lo0.0;
}
area 0.0.0.0 {
    virtual-link neighbor-id 172.30.5.2 transit-area 0.0.0.3;
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.10      ge-0/0/4.118           Full      172.30.5.8       128    33
172.30.0.2       ae0.0                  Full      172.30.5.2       128    36
172.30.0.2       vl-172.30.5.2          Full      172.30.5.2         0    37
```

Below is the OSPF configuration with the virtual link on R2.

```
[edit]
lab@R2# show protocols ospf
traceoptions {
    file ospf.log;
    flag error detail;
}
area 0.0.0.3 {
    interface ae0.0 {
        authentication {
            md5 1 key "$9$Sy9eLNUDkm5F4aGi.56/SreWX-"; ## SECRET-DATA
        }
    }
}
area 0.0.0.0 {
    virtual-link neighbor-id 172.30.5.1 transit-area 0.0.0.3;
    interface lo0.0;
    interface ge-0/0/4.127 {
        authentication {
            md5 1 key "$9$38y990IXxdw2aKMLNb2GU369pOR"; ## SECRET-DATA
        }
    }
    interface ge-0/0/4.123 {
        authentication {
            md5 1 key "$9$QtsSntOW87dwgreMX-waJQFnCpB"; ## SECRET-DATA
        }
    }
}

[edit]
lab@R2# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.14      ge-0/0/4.123           Full      172.30.5.3       128    38
172.30.0.18      ge-0/0/4.127           Full      172.30.5.7       128    38
172.30.0.1       vl-172.30.5.1          Full      172.30.5.1         0    30
172.30.0.1       ae0.0                  Full      172.30.5.1       128    36
```

The virtual link is represented as an interface and a neighbor adjacency is established over it. R1 has two neighbor adjacencies with R2, one for Area 3 and one for Area 2.

© 2020 Juniper Networks, Inc. All rights reserved.

Verify the results of the virtual link configuration on R8.

```
[edit]
lab@R8# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.45      ge-0/0/4.178           Full      172.30.5.7       128    38
172.30.0.9       ge-0/0/4.118           Full      172.30.5.1       128    37
172.30.0.45      vl-172.30.5.7          Full      172.30.5.7         0    38
```

R8 has also two adjacencies with R7, one for Area 1 and one for Area 2.

Again, verify loopback IP address reachability.

```
[edit protocols ospf area 0.0.0.0]
lab@R7# run show route 172.30.5/24 terse

inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.30.5.1/32      O  10          2             >172.30.0.17
                                                    172.30.0.46
* ? 172.30.5.2/32      O  10          1             >172.30.0.17
* ? 172.30.5.3/32      O  10          2             >172.30.0.17
                                                    172.30.0.41
* ? 172.30.5.4/32      O  10          3             >172.30.0.17
                                                    172.30.0.41
* ? 172.30.5.5/32      O  10          2             >172.30.0.41
* ? 172.30.5.6/32      O  10          1             >172.30.0.41
* ? 172.30.5.7/32      D   0                        >lo0.0
* ? 172.30.5.8/32      O  10          1             >172.30.0.46
```

All loopback IP addresses are visible in R7 routing table.

```
[edit]
lab@R7# run show ospf database lsa-id 172.30.5.1

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.1       172.30.5.1      0x80000001   770  0x22 0x3fb1  36
Summary  172.30.5.1       172.30.5.1      0x80000001   825  0x22 0xedad  28
Summary  172.30.5.1       172.30.5.8      0x80000001   481  0x22 0xcdc5  28
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    OSPF database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Summary *172.30.5.1       172.30.5.7       0x80000002   352  0x22 0xdbb6  28
Summary  172.30.5.1       172.30.5.8       0x80000005   152  0x22 0xc5c9  28
```

R7 has a summary LSA in Area 0 received from R8.

13) Test the redundancy.

You must verify the redundancy of the network. This is because the task requirement states to not isolate any parts of the network if an ABR fails or loses RIP routes if R4 or R5 fails.

You can verify this by issuing a *traceroute* to the VR device from R8.

```
[edit]
lab@R8# run traceroute 172.30.32.1 no-resolve
traceroute to 172.30.32.1 (172.30.32.1), 30 hops max, 40 byte packets
 1  172.30.0.45  12.809 ms  13.559 ms  8.500 ms
 2  172.30.0.41  7.781 ms  7.841 ms  10.023 ms
 3  172.30.0.33  10.020 ms  9.993 ms  10.018 ms
 4  172.30.0.58  8.278 ms !N  5.237 ms !N  4.815 ms !N
```

Traffic is passing from R8-R7-R6-R5-VR.

Now you can deactivate OSPF protocol on R6 and again issue the same *traceroute* command. (You may need to wait a minute or two for the network to synchronize.)

```
[edit]
lab@R6# deactivate protocols ospf

[edit]
lab@R6# commit
commit complete
```

```
[edit protocols ospf area 0.0.0.0]
lab@R8# run traceroute 172.30.32.1 no-resolve
traceroute to 172.30.32.1 (172.30.32.1), 30 hops max, 40 byte packets
 1  172.30.0.45  11.395 ms  11.927 ms  10.039 ms
 2  172.30.0.17  13.501 ms  8.080 ms  8.002 ms
 3  172.30.0.14  7.935 ms  14.079 ms  8.019 ms
 4  172.30.0.22  7.716 ms  8.089 ms  7.933 ms
 5  172.30.0.50  5.273 ms !N  4.985 ms !N  4.791 ms !N
```

You can see that the traffic is passing R8-R7-R2-R3-R4-VR. The redundancy is working as expected!

© 2020 Juniper Networks, Inc. All rights reserved.

Remember to activate OSPF on R6.

```
[edit]
lab@R6# activate protocols ospf

[edit]
lab@R6# commit
commit complete
```

The last step is to write down the changes that you performed to fix the network.

Task 2: IS-IS Troubleshooting

IS-IS Overview

IS-IS is another Interior Gateway Protocol that uses link state information to build forwarding paths through the network. Like OSPF it uses the Shortest Path First (SPF) algorithm. IS-IS was originally developed for OSI routing and later adopted for routing IP. In fact, IS-IS is designed in a way that it is easily extendible. So, for example it's much easier to add new address families or protocol extensions then with OSPF.

As with OSPF a single IS-IS domain can be divided into smaller areas. Routing between areas is organized in a hierarchical fashion. Routers in IS-IS can be configured as Level 1, Level 2 or both. Level 1 routers route inside the area whereas Level 2 routers connect different areas. Level 1/Level 2 routers have interfaces participating in Level 1 and others in Level 2. Unlike OSPF, a router can be part of a single area only. L1/L2 routers send PDUs with the attached bit, announcing that they are a gateway to other areas. Areas affect the adjacency formation. Levels control the flooding of LSPs.

IS-IS uses hello PDUs (IIH PDUs). They are used to discover other IS-IS routers on the link. Once the adjacency is established IIH are used as keepalives between IS-IS neighbors. There are two types of IIH PDUs: LAN and point-to-point. Different LAN IIHs are used for Level 1 and Level 2 but are very similar. By default, hello PDUs are send every 9 seconds, except on broadcast networks where the DIS sends hellos every 3 seconds. Adjacency establishment is much simpler compared to OSPF. Neighbors start in the “Down” state, when the first message is received from a neighbor, they reach the “Initializing” state. When a bidirectional conversation is verified the neighbors move to state ‘UP’ state.

The following adjacency formation rules are used by IS-IS:

- Level 1 routers never form adjacencies with Level 2 routers.
- Level 2 routers never form adjacencies with Level 1 routers.

© 2020 Juniper Networks, Inc. All rights reserved.

- The area ID must match for Level 1 adjacencies
- For Level 2 adjacencies Area IDs might be different.

As mentioned, on broadcast multi-access networks IS-IS elects a Designated Intermediate Systems (DIS). For DIS election the tie breaking rule is the highest MAC address in case the DIS priorities area equal. DIS priorities range from 0 to 127. The device with the highest priority wins the election. The default is 64 on all interfaces. A separate DIS is elected for Level 1 and Level 2. The DIS represents the broadcast segment. All routers on the segment synchronize their link state database with the DIS. There is no backup DIS elected as with OSPF. The other difference with OSPF is that a full mesh of adjacencies is formed on broadcast networks. The DIS election process starts all over again if new router is connected to the segment. If a new router has a higher priority, it becomes the new DIS.

IS-IS Link State PDUs (LSP) are exchanged and are used to build the link-state database. They are sent when a new adjacency is established, network changes or in response to sequence number PDUs. Based on the LSP in the database the SPF algorithm is executed. Separate LSPs are exchanged for Level 1 and Level 2.

Complete sequence number PDUs (CSNPs) contain a full description of the IS-IS link state database. They are periodically sent to all interfaces. IS-IS routers use them to update and synchronize their databases. On broadcast segment only the DIS sends CSNPs. The DIS also uses them for the acknowledgment of the received LSPs.

Partial sequence number PDUs (PSNPs) are used to request missing or outdated LSPs, and for the acknowledgment on point-to-point links.

IS-IS supports two types of metrics (narrow and wide). Narrow metrics are in the range of 1 to 63 but are insufficient for modern networks. Wide metrics can range up to 16 million. To force the use of wide metrics you have configure the router to use them. It is recommended that all routers in the network have a consistent metric type. Like OSPF, the Junos IS-IS implementation supports a reference bandwidth. The configured value is divided by the bandwidth of the interface in bits per second. By default, all interfaces are assigned an IS-IS metric of 10. You can configure a manual metric on each IS-IS interface.

1) Let's start with loading the part 2 starting configuration on all eight devices from the files provided with the JNCIE workbook for Chapter Two.

```
[edit]
lab@R1# load override ch2-part2-start.config
```

2) Verify IS-IS adjacencies according to Figure 5. Pay attention that only one adjacency is configured on each IS-IS interface.

© 2020 Juniper Networks, Inc. All rights reserved.

a) R1

```
[edit]
lab@R1# show interfaces ae0
mtu 1400;
aggregated-ether-options {
    lacp {
        active;
    }
}
unit 0 {
    description "R2 connection";
    family inet {
        address 172.30.0.1/30;
    }
    family iso;
    family inet6;
}
```

```
[edit]
lab@R1# show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.1/32;
    }
    family iso {
        address 49.0001.1720.3000.5001.00;
    }
    family inet6 {
        address fd17:f0f4:f691:5::1/128;
    }
}
```

When enabling IS-IS on a router a network entity title (NET) address must be configured, which is a special ISO Network service access point (NSAP) address. The NET is a combination of the Area ID and the System ID (SysID). The NET is represented in hexadecimal. The NET format consists of an authority and format Identifier (AFI), a domain ID, an area ID, a system identifier and a selector. The simplest format omits the domain ID and is 10 bytes long. As an example, the NET address 49.0001.1921.6800.1002.00 consists of the following parts:

49—AFI
0001—Area ID

© 2020 Juniper Networks, Inc. All rights reserved.

1921.6800.1002—System identifier
00—Selector

Please note that the system identifier must be unique within the network. For IP networks, it's recommended to use the loopback IP address of the router. Configure a loopback NET address with the IP address is helpful when troubleshooting is required on the network.

The first part of the address is the area number, which is a variable number from 1 to 13 bytes. The first byte of the area number, 49, is the authority and format indicator (AFI). The next bytes are the assigned area identifier and can be 0 to 12 bytes. In the examples, 0001 is the area identifier.

The NET address is configured under *family iso* and usually under the loopback interface.

To include interfaces for IS-IS routing, you must configure *family iso* together with adding it under the *protocol isis* hierarchy.

```
[edit]
lab@R1# show protocols isis
interface ge-0/0/4.118 {
    level 2 disable;
    level 1 {
        hello-authentication-key "$9$bB2Zj6/t01hTznCu1yrbs2aGi"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface ae0.0 {
    point-to-point;
    level 1 disable;
    level 2 {
        hello-authentication-key "$9$c-/l87GUH.fzgoZjqfn6cylMLN"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface lo0.0;
```

As with OSPF, to configure IS-IS you must specify the interfaces on which IS-IS adjacencies should be established. By default, both Level 1 and Level 2 adjacency formation is enabled. You can deactivate a particular level if needed.

```
[edit]
lab@R1# run show isis adjacency
Interface             System           L  State       Hold (secs) SNPA
ae0.0                 1720.3000.5002   2  Initializing   22
ge-0/0/4.118          R8               1  Up             8  f8:c0:1:dc:2e:84
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR         L1/L2 Metric
ae0.0                 2   0x1 Disabled        Point to Point           10/10
ge-0/0/4.118          1   0x1 R8.02           Disabled                 10/10
lo0.0                 0   0x1 Passive         Passive                   0/0
```

R1 should have two neighbor adjacencies. The adjacency to R2 is in the *Initializing* state.

Note: The adjacency to R8 is easily recognized, because the hostname of the router is displayed. IS-IS is sending TLV 137, which dynamically maps the hostname of the IS-IS enabled routers.

b) R2

```
[edit]
lab@R2# show interface lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.2/32;
    }
    family iso {
        address 49.0002.1720.3000.5002.00;
    }
    family inet6 {
        address fd17:f0f4:f691:5::2/128;
    }
}
```

Note: R2 is configured with Area 49.0002.

```
[edit]
lab@R2# show protocols isis
level 2 {
    authentication-key "$9$prLBBclws4JUH7-b2aU.mp0BESe"; ## SECRET-DATA
    authentication-type md5;
}
interface ge-0/0/4.123 {
    point-to-point;
    level 1 disable;
    level 2 {
        hello-authentication-key "$9$5FCuvMXNVYSrK87V4o5QF/A0"; ## SECRET-DATA
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        hello-authentication-type md5;
    }
}
interface ge-0/0/4.127 {
    point-to-point;
    level 1 disable;
}
interface ae0.0 {
    point-to-point;
    level 1 disable;
    level 2 {
        hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface lo0.0;

[edit]
lab@R2# run show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR         L1/L2 Metric
ae0.0              2   0x1 Disabled        Point to Point          10/10
ge-0/0/4.123       2   0x1 Disabled        Point to Point          10/10
ge-0/0/4.127       2   0x1 Disabled        Point to Point          10/10
lo0.0              0   0x1 Passive         Passive                  0/0

[edit]
lab@R2# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.123          1720.3000.5003   2  Up                   24
ge-0/0/4.127          1720.3000.5001   2  Up                   23
```

R2 has two of the three neighbors in state ‘Up’.

Note: On R2 both established adjacencies have no hostnames displayed.

c) R3

```
[edit]
lab@R3# show interface lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.3/32;
    }
    family iso {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        address 49.0001.1720.3000.5003.00;
    }
    family inet6 {
        address fd17:f0f4:f691:5::3/128;
    }
}
```

```
[edit]
lab@R3# show protocols isis
export [ inject-externals l1-to-l2 ];
interface ge-0/0/4.123 {
    point-to-point;
    level 1 disable;
    level 2 {
        hello-authentication-key "$9$5FCuvMXNVYSrK87V4o5QF/A0"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface ge-0/0/4.134 {
    level 2 disable;
    level 1 {
        hello-authentication-key "$9$MusXVY.mTF6ADiqfz6u0M8X-b2"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface ge-0/0/4.136 {
    level 2 disable;
    level 1 {
        hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface lo0.0;
```

```
[edit]
lab@R3# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
ge-0/0/4.123          2   0x1 Disabled        Point to Point       10/10
ge-0/0/4.134          1   0x1 R3.00           Disabled             10/10
ge-0/0/4.136          1   0x1 R3.00           Disabled             10/10
lo0.0                 0   0x1 Passive         Passive               0/0

[edit]
lab@R3# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.123          R2               2  Up                    20
```


© 2020 Juniper Networks, Inc. All rights reserved.

R3 has just one of the three neighbor adjacencies established.

d) R4

```
[edit]
lab@R4# show interfaces lo0.0
family inet {
    filter {
        input protect-re;
    }
    address 172.30.5.4/32;
}
family iso {
    address 49.0002.1720.3000.5004.00;
}
family inet6 {
    address fd17:f0f4:f691:5::4/128;
}
```

```
[edit]
lab@R4# show protocols isis
export rip-to-isis;
level 1 wide-metrics-only;
interface ge-0/0/4.134 {
    level 2 disable;
    level 1 {
        hello-authentication-key "$9$MusXVY.mTF6ADiqfz6u0M8X-b2"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface ge-0/0/4.145 {
    level 2 disable;
    level 1 {
        hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface lo0.0;

[edit]
lab@R4# run show isis adjacency
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR         L1/L2 Metric
ge-0/0/4.134          1   0x1 R4.00           Disabled                10/10
ge-0/0/4.145          1   0x1 R4.00           Disabled                10/10
lo0.0                 0   0x1 Passive         Passive                  0/0
```

R4 has no established adjacencies.

e) R5

```
[edit]
lab@R5# show interface lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.5/32;
    }
    family iso {
        address 49.0002.1720.3000.5005.00;
    }
    family inet6 {
        address fd17:f0f4:f691:5::5/128;
    }
}

[edit]
lab@R5# show protocols isis
export rip-to-isis;
interface ge-0/0/4.145 {
            level 2 disable;
            level 1 {
                hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ##
SECRET-DATA
                hello-authentication-type md5;
            }
        }
interface ae0.0 {
            level 2 disable;
            level 1 {
                hello-authentication-key "$9$MusXVY.mTF6ADiqfz6u0M8X-b2"; ##
SECRET-DATA
                hello-authentication-type md5;
            }
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
interface lo0.0;

[edit]
lab@R5# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 R6               1  Up                   20  f8:c0:1:dc:2c:80

[edit]
lab@R5# run show isis interface
IS-IS interface database:
Interface      L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ae0.0          1   0x2 R5.02             Disabled               10/10
ge-0/0/4.14    1   0x1 R5.00             Disabled               10/10
lo0.0          0   0x1 Passive           Passive                 0/0
```

R5 has one adjacency established with R6.

f) R6

```
[edit]
lab@R6# show interface lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.6/32;
    }
    family iso {
        address 49.0002.1720.3000.5006.00;
    }
    family inet6 {
        address fd17:f0f4:f691:5::6/128;
    }
}
```

```
[edit]
lab@R6# show protocols isis
export [ inject-externals l1-to-l2 ];
interface ge-0/0/4.136 {
            level 2 disable;
            level 1 {
                hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ##
SECRET-DATA
                hello-authentication-type md5;
            }
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
interface ge-0/0/4.167 {
            level 1 disable;
            level 2 {
                hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ##
SECRET-DATA
                hello-authentication-type md5;
            }
        }
interface ae0.0 {
            level 2 disable;
            level 1 {
                hello-authentication-key "$9$MusXVY.mTF6ADiqfz6u0M8X-b2"; ##
SECRET-DATA
                hello-authentication-type md5;
            }
        }
interface lo0.0;

[edit]
lab@R6# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 R5               1  Up                 7  f8:c0:1:dd:4:0

[edit]
lab@R6# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR         L1/L2 Metric
ae0.0                 1   0x1 R5.02   Disabled                  10/10
ge-0/0/4.136          1   0x1 R6.00         Disabled                10/10
ge-0/0/4.167          2   0x1 Disabled        R6.00                 10/10
lo0.0                 0   0x1 Passive         Passive                  0/0
```

R6 has only one adjacency established to R5.

g) R7

```
[edit]
lab@R7# show interface lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.7/32;
    }
    family iso {
        address 49.0002.1720.3000.5001.00;
    }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    family inet6 {
        address fd17:f0f4:f691:5::7/128;
    }
}

[edit]
lab@R7# show protocols isis
interface ge-0/0/4.127 {
            point-to-point;
            level 1 disable;
            level 2 {
                hello-authentication-key "$9$MusXVY.mTF6ADiqfz6u0M8X-b2"; ##
SECRET-DATA
                hello-authentication-type md5;
            }
        }
interface ge-0/0/4.167 {
            point-to-point;
            level 1 disable;
            level 2 {
                hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ##
SECRET-DATA
                hello-authentication-type md5;
            }
        }
interface ge-0/0/4.178 {
            point-to-point;
            level 1 disable;
            level 2 {
                hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ##
SECRET-DATA
                hello-authentication-type md5;
            }
        }
interface lo0.0;
```

```
[edit]
lab@R7# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.127          R2               2  Up                   20

[edit]
lab@R7# run show isis interface
IS-IS interface database:
Interface         L CirID Level 1 DR      Level 2 DR         L1/L2 Metric
ge-0/0/4.127      2   0x1 Disabled        Point to Point           10/10
ge-0/0/4.167      2   0x1 Disabled        Point to Point           10/10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ge-0/0/4.178       2    0x1 Disabled       Point to Point          10/10
lo0.0              0    0x1 Passive        Passive                  0/0
```

R7s adjacency with R2 is established. The other two interfaces configured for IS-IS have no adjacencies established.

h) R8

```
[edit]
lab@R8# show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.8/32;
    }
    family iso {
        address 49.0001.1720.3000.5008.00;
    }
    family inet6 {
        address fd17:f0f4:f691:5::8/128;
    }
}
```

```
[edit]
lab@R8# show protocols isis
interface ge-0/0/4.118 {
    level 2 disable;
    level 1 {
        hello-authentication-key "$9$MusXVY.mTF6ADiqfz6u0M8X-b2"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface ge-0/0/4.178 {
    point-to-point;
    level 1 disable;
    level 2 {
        hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface lo0.0;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/4.118          R1             1  Up                    18  0:26:88:ef:75:84
```

```
[edit]
lab@R8# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR      L1/L2 Metric
ge-0/0/4.118          1   0x2 R8.02             Disabled              10/10
lo0.0                 0   0x1 Passive           Passive                0/0
```

R8 has one established adjacency with R1.

Note: The `run show isis interface` command displays only the interface to R1. But according the configuration and Figure 5, there should be two.

3) Fix the IS-IS adjacencies.

a) R1 – R2 adjacency.

The adjacency on R1 to R2 is in the 'Initializing' state, which means that IIHs are received from R2, but on R2 there is nothing in the adjacency table about R1.

IS-IS pads the hello messages to ensure that the MTU on both ends is the same. For the similar reason as OSPF, IS-IS uses this to ensure that large LSPs can be transmitted and received on the links.

Verifying the MTU size on both ends of Aggregated Ethernet interface between R1 and R2 reveals the problem.

```
[edit]
lab@R1# run show interfaces ae0.0
  Logical interface ae0.0 (Index 335) (SNMP ifIndex 520)
    Description: R2 connection
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           917          0         85254            0
        Output:           145          0          6928            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Protocol inet, MTU: 1386
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 172.30.0.0/30, Local: 172.30.0.1, Broadcast: 172.30.0.3
    Protocol iso, MTU: 1383
      Flags: Is-Primary
    Protocol inet6, MTU: 1386
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold
cnt: 0, NH drop cnt: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2e6b:f5ff:fe86:c6c0
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

```
[edit]
lab@R2# run show interfaces ae0.0
  Logical interface ae0.0 (Index 332) (SNMP ifIndex 548)
    Description: R1 connection
    Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :         10363          0       1516348           48
        Output:         11442          0       1752261            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Protocol inet, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 1, Curr new hold
cnt: 0, NH drop cnt: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 172.30.0.0/30, Local: 172.30.0.2, Broadcast: 172.30.0.3
    Protocol iso, MTU: 1497
      Flags: Is-Primary
    Protocol inet6, MTU: 1500
    Max nh cache: 75000, New hold nh limit: 75000, Curr nh cnt: 0, Curr new hold
cnt: 0, NH drop cnt: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2e6b:f5ff:fe1a:b8c6
    Protocol multiservice, MTU: Unlimited
      Flags: Is-Primary
```

R1 is not accepting the padded IIHs because they are too big.

© 2020 Juniper Networks, Inc. All rights reserved.

Note: There is no easy way to find that the problem is MTU related as with OSPF, where you can easily find this with traceoptions. By default, the MTU on all Ethernet interfaces is the same. If there is misconfiguration it should be configured manually. Searching the configuration for the `mtu` command might indicate the problem. You can use the match keyword like this – `show configuration | display set | match mtu`

Fix the MTU mismatch by deleting the manually configured MTU on R1 ae0 interface.

```
[edit]
lab@R1# show interfaces ae0
mtu 1400;
aggregated-ether-options {
    lacp {
        active;
    }
}
unit 0 {
    description "R2 connection";
    family inet {
        address 172.30.0.1/30;
    }
    family iso;
    family inet6;
}
```

```
[edit]
lab@R1# delete interfaces ae0 mtu

[edit]
lab@R1# commit
commit complete
```

After committing the changes R1 has two neighbors in the “Up” state.

```
[edit]
lab@R1# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ae0.0                 R2             2  Up                 24
ge-0/0/4.118          R8             1  Up                  7  f8:c0:1:dc:2e:84
```

b) R3 – R4 and R3 – R6 adjacencies.

In the previous step you saw that R3 has one established adjacency with R2, however according to Figure 5 and the loaded configuration it's clear that there should be two adjacencies, with R4 and R6.

© 2020 Juniper Networks, Inc. All rights reserved.

Activate IS-IS traceoptions on R4 for collecting errors generated by the protocol.

```
[edit]
lab@R4# set protocols isis traceoptions file isis.log

[edit]
lab@R4# set protocols isis traceoptions flag error detail

[edit]
lab@R4# commit
commit complete
```

Analyze the output in the isis-debug file.

```
[edit]
lab@R4# run show log isis.log
Jun 29 14:27:28 trace_on: Tracing to "/var/log/isis.log" started
Jun 29 14:27:30.315972 ERROR: IIH from 1720.3000.5005 without matching addresses,
interface ge-0/0/4.145
Jun 29 14:27:30.316452 Notification not sent due to 5-second throttle as per
rfc4444: isisRejectedAdjacency
Jun 29 14:27:33.902046 ERROR: IIH from R3 with no matching areas, interface ge-
0/0/4.134
Jun 29 14:27:33.902458     local area  49.0002
Jun 29 14:27:37.416246 ERROR: IIH from 1720.3000.5005 without matching addresses,
interface ge-0/0/4.145
Jun 29 14:27:37.416691 Notification not sent due to 5-second throttle as per
rfc4444: isisRejectedAdjacency
Jun 29 14:27:40.973357 ERROR: IIH from R3 with no matching areas, interface ge-
0/0/4.134
Jun 29 14:27:40.973976     local area  49.0002
Jun 29 14:27:45.462672 ERROR: IIH from 1720.3000.5005 without matching addresses,
interface ge-0/0/4.145
Jun 29 14:27:45.463302 Notification not sent due to 5-second throttle as per
rfc4444: isisRejectedAdjacency
```

Configure the same traceoptions on R6.

```
[edit]
lab@R6# set protocols isis traceoptions file isis.log

[edit]
lab@R6# set protocols isis traceoptions flag error detail

[edit]
lab@R6# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Analyze the output from the file.

```
[edit]
lab@R6# run show log isis.log
Jul  8 21:10:23.310252 ERROR: IIH from 1720.3000.5003 with no matching areas,
interface ge-0/0/4.136
Jul  8 21:10:23.310304      local area  49.0002
Jul  8 21:10:23.310421      area address 49.0001 (3 bytes)
```

It is obvious that R3 is configured with the Area ID 49.0001, but R4 and R6 with Area ID 49.0002. According to Figure 5 the adjacencies between R3 with R4 and R6 should be established as Level 1 adjacencies. To form a Level 1 adjacency the area IDs should match, as it was stated in the overview.

Change the Area ID with the rename command.

```
[edit]
lab@R3# show interfaces lo0.0
family inet {
    filter {
        input protect-re;
    }
    address 172.30.5.3/32;
}
family iso {
    address 49.0001.1720.3000.5003.00;
}
family inet6 {
    address fd17:f0f4:f691:5::3/128;
}
```

```
[edit]
lab@R3# rename interface lo0.0 family iso address 49.0001.1720.3000.5003.00 to address
49.0002.1720.3000.5003.00

[edit]
lab@R3# commit
commit complete
```

Verify the adjacency table on R3 after changing the Area ID reveals that R3 has established adjacencies with R4 and R6.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.123          R2               2  Up                25
ge-0/0/4.134          R4               1  Up                 6  f8:c0:1:dc:31:84
ge-0/0/4.136          R6               1  Up                 7  f8:c0:1:dc:2c:4
```

c) R4 – R5 adjacency

```
[edit]
lab@R4# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.134          R3               1  Up                24  2c:21:72:cd:26:84
```

R4 has only one established adjacency, the one fixed previous step with R3.

Looking at the traceoptions file configured for the fixing the adjacency with R3 you can see the problem with the adjacency with R5 also.

```
[edit]
lab@R4# run show log isis.log
Jun 29 14:27:28 trace_on: Tracing to "/var/log/isis.log" started
Jun 29 14:27:30.315972 ERROR: IIH from 1720.3000.5005 without matching addresses,
interface ge-0/0/4.145
Jun 29 14:27:30.316452 Notification not sent due to 5-second throttle as per
rfc4444: isisRejectedAdjacency
Jun 29 14:27:33.902046 ERROR: IIH from R3 with no matching areas, interface ge-
0/0/4.134
Jun 29 14:27:33.902458     local area  49.0002
Jun 29 14:27:37.416246 ERROR: IIH from 1720.3000.5005 without matching addresses,
interface ge-0/0/4.145
Jun 29 14:27:37.416691 Notification not sent due to 5-second throttle as per
rfc4444: isisRejectedAdjacency
Jun 29 14:27:40.973357 ERROR: IIH from R3 with no matching areas, interface ge-
0/0/4.134
Jun 29 14:27:40.973976     local area  49.0002
Jun 29 14:27:45.462672 ERROR: IIH from 1720.3000.5005 without matching addresses,
interface ge-0/0/4.145
Jun 29 14:27:45.463302 Notification not sent due to 5-second throttle as per
rfc4444: isisRejectedAdjacency
```

There is and address mismatch on interface `ge-0/0/4.145`. Verifying the interface configuration on both routers reveals the mistake.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# show interface ge-0/0/4.145
description "R5 connection";
vlan-id 145;
family inet {
    address 172.30.0.29/30;
    }
family iso;
family inet6;
```

```
[edit]
lab@R5# show interface ge-0/0/4.145
description "R4 connection";
vlan-id 145;
family inet {
    address 172.30.1.30/30;
    }
family iso;
family inet6;
```

Indeed, both interfaces are configured with IP addresses from different subnets. Let's change the IP address on R5 with the command rename.

```
[edit]
lab@R5# rename interfaces ge-0/0/4.145 family inet address 172.30.1.30/30 to address
172.30.0.30/30

[edit]
lab@R5# commit
commit complete
```

After this change, R4 has established an adjacency with R5 also. Before moving on, remember to delete or disable the IS-IS traceoptions configuration to avoid any unnecessary processing on R4.

```
[edit]
lab@R4# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.134          R3               1  Up              23  2c:21:72:cd:26:84
ge-0/0/4.145          R5               1  Up               8  f8:c0:1:dd:3:84
```

```
[edit]
lab@R4# delete protocols isis traceoptions
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# commit
commit complete
```

d) R6 – R7 adjacency

Activate traceoptions on R6 and R7 to flag `hello` communication.

```
[edit]
lab@R6# set protocols isis traceoptions flag hello detail

[edit]
lab@R6# commit
commit complete

[edit]
lab@R7# set protocols isis traceoptions file isis.log

[edit]
lab@R7# set protocols isis traceoptions flag hello detail


lab@R7# commit
commit complete
```

Verify the trace files on both routers.

```
[edit]
lab@R6# run show log isis.log
Jul  8 21:16:26 trace_on: Tracing to "/var/log/isis.log" started
Jul  8 21:16:26.889460 Sending L2 LAN IIH on ge-0/0/4.167
Jul  8 21:16:26.889576     max area 0, circuit type l2
Jul  8 21:16:26.889647     hold time 27, priority 64, circuit id R6.00
Jul  8 21:16:26.889698     speaks IP
Jul  8 21:16:26.889780     speaks IPv6
Jul  8 21:16:26.890048     IP address 172.30.0.41
Jul  8 21:16:26.890419     IPv6 address fe80::fac0:100:a7dc:2c04
Jul  8 21:16:26.890546     area address 49.0001 (3)
Jul  8 21:16:26.890600     restart RR reset RA reset holdtime 0
Jul  8 21:16:26.890708     packet length 85
Jul  8 21:16:26.890831 Sending L1 LAN IIH on ge-0/0/4.136
Jul  8 21:16:26.890898     max area 0, circuit type l1
Jul  8 21:16:26.890957     hold time 9, priority 64, circuit id R6.02
Jul  8 21:16:26.891043     neighbor 2c:21:72:cd:26:84
Jul  8 21:16:26.891091     speaks IP
Jul  8 21:16:26.891151     speaks IPv6
Jul  8 21:16:26.891343     IP address 172.30.0.26
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Jul  8 21:16:26.891647     area address 49.0001 (3)
Jul  8 21:16:26.891700     restart RR reset RA reset holdtime 0
Jul  8 21:16:26.891814     packet length 75
```

```
[edit]
lab@R7# run show log isis.log
Jul  8 21:24:08 trace_on: Tracing to "/var/log/isis.log" started
Jul  8 21:24:08.780722 Sending PTP IIH on ge-0/0/4.178
Jul  8 21:24:08.780840     max area 0, circuit type l2
Jul  8 21:24:08.780905     ptp adjacency tlv length 5
Jul  8 21:24:08.780956     neighbor state down
Jul  8 21:24:08.781037     our extended local circuit id 72
Jul  8 21:24:08.781089     speaks IP
Jul  8 21:24:08.781150     speaks IPv6
Jul  8 21:24:08.781422     IP address 172.30.0.45
Jul  8 21:24:08.781818     IPv6 address fe80::fac0:100:b2dc:3204
Jul  8 21:24:08.781925     area address 49.0002 (3)
Jul  8 21:24:08.781980     restart RR reset RA reset holdtime 0
Jul  8 21:24:08.782106     packet length 85
Jul  8 21:24:08.782213 Sending PTP IIH on ge-0/0/4.167
Jul  8 21:24:08.782300     max area 0, circuit type l2
Jul  8 21:24:08.782350     ptp adjacency tlv length 5
Jul  8 21:24:08.782413     neighbor state down
Jul  8 21:24:08.782461     our extended local circuit id 71
Jul  8 21:24:08.782539     speaks IP
Jul  8 21:24:08.782590     speaks IPv6
Jul  8 21:24:08.782840     IP address 172.30.0.42
Jul  8 21:24:08.783202     IPv6 address fe80::fac0:100:a7dc:3204
Jul  8 21:24:08.783330     area address 49.0002 (3)
```

From the output it is clear that R6 is sending LAN hellos to R7, whereas on the other end R7 is sending PTP hellos only.

This can be seen when looking at the IS-IS statistics on R6 and R7.

```
[edit]
lab@R6# run show isis statistics
IS-IS statistics for R6:
PDU type    Received    Processed       Drops        Sent       Rexmit
LSP               14           14           0          23            0
IIH             1150            9         888        1173            0
CSNP              57           57           0          57            0
PSNP               0            0           0           0            0
Unknown            0            0           0           0            0
Totals             0            0           0           0            0

Total packets received: 1221 Sent: 1253
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 27
Fragments rebuilt: 24
LSP regenerations: 4
Purges initiated: 2
```

```
[edit]
lab@R7# run show isis statistics
IS-IS statistics for R7:
PDU type    Received    Processed       Drops        Sent       Rexmit
LSP              522          522           0         181          354
IIH              673            5         337        1010            0
CSNP             175          175           0         524            0
PSNP               0            0           0         522            0
Unknown            0            0           0           0            0
Totals             0            0           0           0            0

Total packets received: 1370 Sent: 2591

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 14
Fragments rebuilt: 12
LSP regenerations: 6
Purges initiated: 0
```

On both routers the hello messages are dropped.

Like OSPF, IS-IS supports the configuration of back-to-back Ethernet broadcast interfaces (point-to-point interfaces). This prevents DIS election. One of the requirements is that on all L2 interfaces no DIS is elected. So, you have to change the interface on R6 to point-to-point.

```
[edit]
lab@R6# set protocols isis interface ge-0/0/4.167 point-to-point

[edit]
lab@R6# commit
commit complete
```

```
[edit]
lab@R6# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 R5               1  Up              6  f8:c0:1:dd:4:0
ge-0/0/4.136          R3               1  Up             26  2c:21:72:cd:26:84
ge-0/0/4.167          R7               2  Up             23
```


© 2020 Juniper Networks, Inc. All rights reserved.

After the change, R6 established an adjacency with R7. Before moving on, remember to delete or disable the IS-IS traceoptions configuration on R6 and R7 to avoid any unnecessary processing.

R6

```
[edit]
lab@R6# delete protocols isis traceoptions

[edit]
lab@R6# commit
commit complete
```

R7

```
[edit]
lab@R7# delete protocols isis traceoptions

[edit]
lab@R7# commit
commit complete
```

e) R7- R8 adjacency

Looking at the IS-IS interface table on R7 and R8 in the previous step reveals that on R8, the interface to R7 is missing.

```
[edit]
lab@R7# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR         L1/L2 Metric
ge-0/0/4.127          2   0x1 Disabled        Point to Point          10/10
ge-0/0/4.167          2   0x1 Disabled        Point to Point          10/10
ge-0/0/4.178          2   0x1 Disabled        Point to Point          10/10
lo0.0                 0   0x1 Passive         Passive                  0/0
```

```
[edit]
lab@R8# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR         L1/L2 Metric
ge-0/0/4.118          1   0x2 R8.02           Disabled                10/10
lo0.0                 0   0x1 Passive         Passive                  0/0
```

On R8 only the interface to R1 is displayed by the IS-IS interface output.

© 2020 Juniper Networks, Inc. All rights reserved.

This might be a problem with the configuration on R8. Verify if this interface is added to `protocols isis` hierarchy.

```
[edit]
lab@R8# show protocols isis
interface ge-0/0/4.118 {
    level 2 disable;
    level 1 {
        hello-authentication-key "$9$MusXVY.mTF6ADiqfz6u0M8X-b2"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface ge-0/0/4.178 {
    point-to-point;
    level 1 disable;
    level 2 {
        hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface lo0.0;
```

Interface `ge-0/0/4.178` is configured. So, another reason for this interface to be missing from the IS-IS interface table is if the interface is not configured with `family iso`.

```
[edit]
lab@R8# show interfaces ge-0/0/4.178
description "R7 connection";
vlan-id 178;
family inet {
    address 172.30.0.46/30;
}
family inet6;
```

Indeed, `family iso` is missing on the interface to R7.

```
[edit]
lab@R8# set interfaces ge-0/0/4.178 family iso

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.118          R7               1  Up                20  0:26:88:ef:75:84
ge-0/0/4.178          R7               2  Up                23
```

Enabling the interface for IS-IS fixed the adjacency.

Note: Before fixing the issue with missing `family iso` on R8 it had just one neighbor R1. In the previous step you saw that the adjacency table correctly displayed the hostname "R1" on R1. But after the change, the adjacency table shows R7's hostname "R7" on both interfaces. Later you will have to fix this issue as well.

4) Verify the IS-IS link state database.

a) R1

```
[edit]
lab@R1# run show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R1.00-00                    0xd   0x78ad      540 L1 L2 Attached
R8.00-00                    0x7   0x113a      758 L1 L2 Attached
R8.02-00                    0x5   0xe814      755 L1 L2
  3 LSPs

IS-IS level 2 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R1.00-00                   0x10   0x8690      929 L1 L2
R2.00-00                   0x11   0xe48d      446 L1 L2
  2 LSPs
```

As Level 1 and Level 2 router, it has separate databases for both Levels. LSP `R8.02-00` is sent from the DIS on the broadcast segment. Those LSPs can be recognized by the difference in the first digits displayed directly after the hostname. R8 is elected as the DIS on the segment between R1 and R8. R1 receives the L2 LSP from R2.

b) R2

```
[edit]
lab@R2# run show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R2.00-00                    0xb   0xf8ea     1055 L1 L2
  1 LSPs
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
R2.00-00                0x11   0xe48d      430 L1 L2
  1 LSPs
```

On R2 both databases have only the LSP generated by the local router. This is suspicious because R2 has adjacencies with three other routers.

c) R3

```
[edit]
lab@R3# run show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
R3.00-00                0xa   0x43a6     1134 L1 L2 Attached
R4.00-00                0xa    0x163     1195 L1 L2
R4.02-00                0x7   0x296c      867 L1 L2
R5.00-00                0x9   0xb32a     1185 L1 L2
R5.02-00                0x5   0x43bc      537 L1 L2
R5.03-00                0x3   0xcc30     1185 L1 L2
R6.00-00                0x6   0x2967      383 L1 L2 Attached
R6.02-00                0x5   0x25d9      383 L1 L2
  8 LSPs

IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
R2.00-00                0x12   0xc84b     1149 L1 L2
R3.00-00                0x9    0x6404      360 L1 L2
  2 LSPs
```

R3 correctly receives the LSP from R2 in its L2 database.

5) Check the link state database issue on R2

If you remember the adjacency table on R2, the neighbors were displayed without hostnames, but with their SysID instead. This means that the LSP with the appropriate TLVs are not received by R2. The adjacencies are established but the databases are not synchronized.

```
[edit]
lab@R2# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 1720.3000.5001   2  Up                 21
ge-0/0/4.123          1720.3000.5003   2  Up                 21
ge-0/0/4.127          1720.3000.5001   2  Up                 20
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify the IS-IS statistics for dropped PDUs.

```
[edit]
lab@R2# run show isis statistics
IS-IS statistics for R2:
PDU type     Received    Processed       Drops        Sent       Rexmit
LSP              3262            0        3262          24         2769
IIH              1783           13           0        2774            0
CSNP             2637            0        2637        1079            0
PSNP             2787            0        2787           0            0
Unknown             0            0           0           0            0
Totals              0            0           0           0            0

Total packets received: 10469 Sent: 6646

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0
SPF runs: 32
Fragments rebuilt: 29
LSP regenerations: 14
Purges initiated: 0
```

All LSPs, CSNPs and PSNPs are dropped by R2.

IS-IS supports simple and MD5 authentication. MD5 authentication uses an encrypted checksum for the PDUs. IS-IS authentication can be configured at different places in the configuration. IS-IS authentication is per Level. Adding authentication under the Level, authenticates every PDUs (Hellos, LSPs, CSNPs and PSNPs) generated by the router. Different types of authentication can be configured for Level 1 and Level 2. Adding authentication at the interface level under the `protocols isis` hierarchy, authenticates only the hello messages transmitted on that interface.

Verifying the authentication configured on R2 reveals that for Level 2 all PDUs are authenticated.

```
[edit]
lab@R2# run show isis authentication
Interface             Level IIH Auth   CSN Auth   PSN Auth
ae0.0                 2     MD5        MD5        MD5
ge-0/0/4.123          2     MD5        MD5        MD5
ge-0/0/4.127          2     MD5        MD5        MD5

L1 LSP Authentication: None
L2 LSP Authentication: MD5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# show protocols isis
level 2 {
    authentication-key "$9$prLBBclws4JUH7-b2aU.mp0BESe"; ## SECRET-DATA
    authentication-type md5;
}
interface ge-0/0/4.123 {
    point-to-point;
    level 1 disable;
    level 2 {
        hello-authentication-key "$9$5FCuvMXNVYSrK87V4o5QF/A0"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface ge-0/0/4.127 {
    point-to-point;
    level 1 disable;
}
interface ae0.0 {
    point-to-point;
    level 1 disable;
    level 2 {
        hello-authentication-key "$9$ROMSvLaJDH.5s2oGi.zFRhSeMX"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface lo0.0;
```

Level 2 is configured for MD5 authentication of all PDUs.

Note: Interface `ge-0/0/4.127` is not configured for hello authentication, but the adjacency with R7 is established. This is possible because L2 authentication is configured for all PDUs.

6) Fix the authentication problem on R2.

The requirement for this Task is to enable authentication for Hello PDUs only. To comply with this requirement and for fixing the authentication problem, you have to delete Level 2 authentication on R2.

```
[edit]
lab@R2# delete protocols isis level 2 authentication-type

[edit]
lab@R2# delete protocols isis level 2 authentication-key
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# commit
commit complete
```

Check the LSDB after this change.

```
[edit]
lab@R2# run show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R2.00-00                  0x1e   0x659a      993 L1 L2 Attached
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R1.00-00                  0x3b   0x9339      991 L1 L2
R2.00-00                  0x52   0x888c      993 L1 L2
R3.00-00                   0xe   0x528e     1134 L1 L2
  3 LSPs
```

Now the L2 database looks different. You can see that the R7 LSP is missing in the database. Checking the IS-IS adjacency shows the problem with R7.

```
[edit]
lab@R2# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 R1               2  Up                    26
ge-0/0/4.123          R3               2  Up                    22
ge-0/0/4.127          1720.3000.5007   2  Initializing          26
```

Deleting the authentication on R2 disables the authentication on the hellos transmitted on the interface to R7. Configuring authentication directly under interface `ge-0/0/4.127` establishes the adjacency again.

```
[edit]
lab@R2# set protocols isis interface ge-0/0/4.127 level 2 hello-authentication-type md5

[edit]
lab@R2# set protocols isis interface ge-0/0/4.127 level 2 hello-authentication-key
jncie

[edit]
lab@R2# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Configure hello authentication under the interface to R7 fixes part of the problem. Depending on when you issue the command you may see R7's entry in the database or you may see R1's entry. There is still something not quite right with the network.

```
[edit]
lab@R2# run show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R2.00-00                  0x4    0xfd3     1172 L1 L2 Attached
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R7.00-00                  0x9e   0x3c63     1196 L1 L2
R2.00-00                  0x6b   0xe74f     1108 L1 L2
R3.00-00                  0x10   0x4e90     1073 L1 L2
R6.00-00                  0x1c   0xfb5e      434 L1 L2
R8.00-00                   0xb   0xa679      560 L1 L2
  5 LSPs
```

Not all six LSPs of the L1 routers are in the database at the same time.

7) Verify that the L2 interfaces do not elect a DIS.

The last output from the previous task proves that the L2 segments do not elect a DIS. All LSPs are with IDs ending with 00. No pseudo node LSPs are into the L2 database. Great!

8) Verify the IS-IS hostname in the database.

a) R1.

```
[edit]
lab@R1# run show isis hostname
IS-IS hostname database:
System ID      Hostname                                  Type
1720.3000.5001 R1                                       Static
1720.3000.5002 R2                                       Dynamic
1720.3000.5003 R3                                       Dynamic
1720.3000.5006 R6                                       Dynamic
1720.3000.5008 R8                                       Dynamic
```

The output displays the SysID with the mapped hostname and the way it is learned. R7 is missing in the database.

b) R2

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show isis hostname
IS-IS hostname database:
System ID      Hostname                                    Type
1720.3000.5001 R7                                         Dynamic
1720.3000.5002 R2                                         Static
1720.3000.5003 R3                                         Dynamic
1720.3000.5006 R6                                         Dynamic
1720.3000.5008 R8                                         Dynamic
```

Depending on when you execute the command you will see R1's or R7's hostname in the list. The output indicates that the SysID of R1 is sometimes mapped to the hostname of R7.

c) R7

```
[edit]
lab@R7# show interfaces lo0.0
family inet {
    filter {
        input protect-re;
    }
    address 172.30.5.7/32;
}
family iso {
    address 49.0002.1720.3000.5001.00;
}
family inet6 {
    address fd17:f0f4:f691:5::7/128;
```

Indeed, R7 is configured with the same SysID as R1. This is the reason why both adjacencies on R8 are displayed as System 'R7'.

Because all SysID in the network are based on the IP address on the loopback interface, changing R7s NET address fixes the problem.

```
[edit]
lab@R7# rename interface lo0.0 family iso address 49.0002.1720.3000.5001.00 to address
49.0002.1720.3000.5007.00

[edit]
lab@R7# commit
commit complete
```

The new SysID on R7 is now displayed in the hostname database. And a quick verification of R2's database reveals that all devices are now present in the Level 2 database.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show isis hostname
IS-IS hostname database:
System ID      Hostname                                    Type
1720.3000.5001 R1                                        Dynamic
1720.3000.5002 R2                                        Dynamic
1720.3000.5003 R3                                        Dynamic
1720.3000.5006 R6                                        Dynamic
1720.3000.5007 R7                                        Static
1720.3000.5008 R8                                        Dynamic


[edit]
lab@R2# run show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R2.00-00                    0x5   0x9781      474 L1 L2 Attached
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R1.00-00                  0x10e   0xea0e     1181 L1 L2
R2.00-00                   0x76   0x41de     1170 L1 L2
R3.00-00                   0x12   0x4a92      704 L1 L2
R6.00-00                   0x1f   0x3e0d     1168 L1 L2
R7.00-00                    0x4   0x4343     1175 L1 L2
R8.00-00                    0xe   0x8b85     1168 L1 L2
  6 LSPs
```

9) Verify RIP routing and RIP – IS-IS redistribution.

Multilevel IS-IS networks have certain rules for the distribution of routing information between Levels. Flooding of LSPs is limited within the Level. L1/L2 routers control the distribution based on the following rules:

Internal routes:

- Level 1 routers flood LSPs to other L1 routers and L1/L2 routers.
- L1/L2 routers flood the L1 information to other L2 routers.
- L1/L2 routers do not flood L2 information to L1 routers.
- L2 routers send LSPs with the Attached bit set towards L1. L1 routers install a default route pointing to the L1/L2 routers.

```
[edit]
lab@R4# run show route protocol isis
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 49 destinations, 51 routes (49 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[IS-IS/15] 00:32:20, metric 10
                     > to 172.30.0.21 via ge-0/0/4.134
172.30.0.24/30      *[IS-IS/15] 00:32:20, metric 20
                     > to 172.30.0.21 via ge-0/0/4.134
172.30.0.32/30      *[IS-IS/15] 00:31:44, metric 20
                     > to 172.30.0.30 via ge-0/0/4.145
172.30.5.3/32       *[IS-IS/15] 00:32:20, metric 10
                     > to 172.30.0.21 via ge-0/0/4.134
172.30.5.5/32       *[IS-IS/15] 00:31:44, metric 10
                     > to 172.30.0.30 via ge-0/0/4.145
...
```

R4 installs a default route to R3.

External routes:

- L1 routers redistribute external information using policies. All L1 and L1/L2 routers receive the external information.
- L1/L2 routers are not distributing external routes outside L1. Except in one case, which is described later.
- L1/L2 routers do not distribute external routes to L1 routers. L1 routers use a default route to reach destinations outside of the L1 domain.

The default routing policy for IS-IS is similar to OSPF. As in OSPF, the information in the database should be identical on all routers in the same level, so import policies are not supported. There is difference with OSPF regarding the use of export policies. The default export policy permits local IS-IS information. This information is retrieved from the routing table.

Junos has different default route preferences for L1 internal, L1 external routes and for L2 internal and external routes.

- *IS-IS L1 internal route – 15*
- *IS-IS L2 internal route – 18*
- *IS-IS L1 external route – 160*
- *IS-IS L2 external route – 165*

R4 and R5 are talking RIP with the VR.

a) R4

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@@R4# show protocols rip
group rip {
    export isis-to-rip;
    neighbor ge-0/0/4.202;
}

[edit]
lab@R4# run show rip neighbor
                Local  Source          Destination      Send   Receive   In
Neighbor        State  Address         Address          Mode   Mode      Met
--------        -----  -------         -----------      ----   -------   ---
ge-0/0/4.202       Up 172.30.0.49     224.0.0.9        mcast  both        1
```

The routing policy `isis-to-rip` exports all IS-IS routes according to the task requirements.

```
[edit]
lab@R4# show policy-options policy-statement isis-to-rip
term 1 {
    from protocol isis;
    then {
        metric 1;
        accept;
    }
}
```

See below the RIP routes received from the VR device.

```
[edit]
lab@R4# run show route receive-protocol rip 172.30.0.50

inet.0: 49 destinations, 51 routes (49 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.0.20/30      [RIP/100] 02:30:58, metric 7, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.5.4/32       [RIP/100] 02:30:58, metric 7, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.32.0/24     *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.33.0/24     *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.34.0/24     *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.35.0/24     *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.36.0/24     *[RIP/100] 05:05:12, metric 2, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.37.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.38.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.39.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.40.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.41.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.42.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.43.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.44.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.45.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.46.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.47.0/24      *[RIP/100] 05:05:12, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202

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

inet6.0: 9 destinations, 11 routes (9 active, 0 holddown, 0 hidden)
```

The following RIP routes are advertised to the VR by R4.

```
[edit]
lab@R4# run show route advertising-protocol rip 172.30.0.49

inet.0: 49 destinations, 51 routes (49 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[IS-IS/15] 00:41:01, metric 10
                    > to 172.30.0.21 via ge-0/0/4.134
172.30.0.24/30      *[IS-IS/160] 00:06:23, metric 13
                    > to 172.30.0.30 via ge-0/0/4.145
172.30.0.32/30      *[IS-IS/15] 02:31:57, metric 20
                    > to 172.30.0.30 via ge-0/0/4.145
172.30.5.3/32       *[IS-IS/15] 02:48:12, metric 10
                    > to 172.30.0.21 via ge-0/0/4.134
172.30.5.5/32       *[IS-IS/15] 02:31:57, metric 10
                    > to 172.30.0.30 via ge-0/0/4.145
192.168.8.0/24      *[IS-IS/160] 02:48:12, metric 10
                    > to 172.30.0.21 via ge-0/0/4.134
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
192.168.9.0/24     *[IS-IS/160] 02:48:12, metric 10
                    > to 172.30.0.21 via ge-0/0/4.134
192.168.10.0/24    *[IS-IS/160] 02:48:12, metric 10
                    > to 172.30.0.21 via ge-0/0/4.134
192.168.11.0/24    *[IS-IS/160] 02:48:12, metric 10
                    > to 172.30.0.21 via ge-0/0/4.134
192.168.12.0/24    *[IS-IS/160] 00:06:23, metric 13
                    > to 172.30.0.30 via ge-0/0/4.145
192.168.13.0/24    *[IS-IS/160] 00:06:23, metric 13
                    > to 172.30.0.30 via ge-0/0/4.145
192.168.14.0/24    *[IS-IS/160] 00:06:23, metric 13
                    > to 172.30.0.30 via ge-0/0/4.145
192.168.15.0/24    *[IS-IS/160] 00:06:23, metric 13
                    > to 172.30.0.30 via ge-0/0/4.145
```

The external routes from the 192.168.8.0/21 range are redistributed by R3 and R6 with the use of the `inject-externals` export policy.

b) R5

```
[edit]
lab@R5# show protocols rip
group rip {
    export isis-to-rip;
    neighbor ge-0/0/4.204;
}

[edit]
lab@R5# run show rip neighbor
                Local  Source      Destination      Send   Receive   In
Neighbor        State  Address     Address          Mode   Mode     Met
--------        -----  -------     -----------      ----   -------  ---
ge-0/0/4.204       Up 172.30.0.57  224.0.0.9        mcast  both       1
```

The `isis-to-rip` policy is used to export IS-IS routes to RIP.

```
[edit]
lab@R5# show policy-options policy-statement isis-to-rip
term 1 {
    from protocol isis;
    then {
        metric 5;
        accept;
    }
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

The policy on R5 increases the metric from 1 to 5. This way the path to R4 is preferred by the VR as is required.

See below the RIP routes received by R5 from VR.

```
[edit]
lab@R5# run show route receive-protocol rip 172.30.0.58

inet.0: 45 destinations, 47 routes (45 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[RIP/100] 02:29:56, metric 3, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.0.24/30      *[RIP/100] 02:46:02, metric 3, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.0.32/30       [RIP/100] 02:29:46, metric 3, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.5.3/32       *[RIP/100] 02:46:02, metric 3, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.5.5/32        [RIP/100] 02:29:46, metric 3, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.32.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.33.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.34.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.35.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.36.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.37.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.38.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.39.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.40.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.41.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.42.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.43.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.44.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.45.0/24      *[RIP/100] 05:04:01, metric 2, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.46.0/24     *[RIP/100] 05:04:01, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.47.0/24     *[RIP/100] 05:04:01, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.8.0/24     *[RIP/100] 02:46:02, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.9.0/24     *[RIP/100] 02:46:02, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.10.0/24    *[RIP/100] 02:46:02, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.11.0/24    *[RIP/100] 02:46:02, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.12.0/24    *[RIP/100] 01:50:48, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.13.0/24    *[RIP/100] 01:50:48, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.14.0/24    *[RIP/100] 01:50:48, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.15.0/24    *[RIP/100] 01:50:48, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
```

RIP routes are advertised to VR by R5.

```
[edit]
lab@R5# run show route advertising-protocol rip 172.30.0.57

inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

172.30.0.20/30     *[IS-IS/15] 01:31:09, metric 20
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.5.4/32      *[IS-IS/15] 01:31:09, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
```


© 2020 Juniper Networks, Inc. All rights reserved.

R5 is not advertising prefixes in the range 192.168.8.0/21. Let's check the routing table for these prefixes.

```
[edit]
lab@R5# run show route 192.168/20

inet.0: 45 destinations, 47 routes (45 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.8.0/24     *[RIP/100] 02:30:23, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.9.0/24     *[RIP/100] 02:30:23, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.10.0/24    *[RIP/100] 02:30:23, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.11.0/24    *[RIP/100] 02:30:23, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.12.0/24    *[RIP/100] 01:35:09, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.13.0/24    *[RIP/100] 01:35:09, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.14.0/24    *[RIP/100] 01:35:09, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.15.0/24    *[RIP/100] 01:35:09, metric 3, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
```

The reason why the routes are not advertised by R5 is obvious. Routes in the range 192.168.8.0/21 are learned by the VR and have a default RIP preference of 100, while the external IS-IS prefixes have a route preference of 160. The active routes in routing table are those learned from RIP, so they are not examined by the RIP export RIP. Remember that only the active routes are processed by the routing policies.

Verify that R5 learns these routes from R3 and R6 also.

```
[edit]
lab@R5# run show isis database level 1 R3.00-00 extensive
IS-IS level 1 link-state database:

R3.00-00 Sequence: 0xc, Checksum: 0x9fc9, Lifetime: 1128 secs
   IS neighbor: R4.00                          Metric:       10
     Two-way fragment: R4.00-00, Two-way first fragment: R4.00-00
   IS neighbor: R6.00                          Metric:       10
     Two-way fragment: R6.00-00, Two-way first fragment: R6.00-00
   IP prefix: 172.30.0.20/30                   Metric:       10 Internal Up
   IP prefix: 172.30.0.24/30                   Metric:       10 Internal Up
   IP prefix: 172.30.5.3/32                    Metric:        0 Internal Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
   IP prefix: 192.168.8.0/24              Metric:        0 External Up
   IP prefix: 192.168.9.0/24              Metric:        0 External Up
   IP prefix: 192.168.10.0/24             Metric:        0 External Up
   IP prefix: 192.168.11.0/24             Metric:        0 External Up
   V6 prefix: fd17:f0f4:f691:5::3/128     Metric:        0 Internal Up
...
```

```
lab@R5# run show isis database level 1 R6.00-00 extensive
IS-IS level 1 link-state database:

R6.00-00 Sequence: 0xa, Checksum: 0x808e, Lifetime: 344 secs
   IS neighbor: R3.00                         Metric:       10
     Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00
   IS neighbor: R5.00                         Metric:       10
     Two-way fragment: R5.00-00, Two-way first fragment: R5.00-00
   IP prefix: 192.168.12.0/24                 Metric:        0 External Up
   IP prefix: 192.168.13.0/24                 Metric:        0 External Up
   IP prefix: 192.168.14.0/24                 Metric:        0 External Up
   IP prefix: 192.168.15.0/24                 Metric:        0 External Up
...
```

From the LSPs, it is clear that R5 receives the routes from R3 and R6. Because of the lower RIP route preference, R5 is using the suboptimal path via VR.

Note: The LSP from R6 carries only the external prefixes. You will see later why the internal prefixes are missing.

10) Fix suboptimal routing.

To fix the suboptimal routing, you can lower the external L1 preference of the IS-IS prefixes. For example, changing it to 90 will make the IS-IS routes to become active and eligible for policy processing.

Another option for fixing this issue is to filter the prefixes, which are advertised from the IS-IS domain to not be again imported by RIP. RIP supports *route tags*. Tagging the routes will help to recognize the routes, which are advertised back to IS-IS domain.

Modifying the policy, which exports the IS-IS prefixes by adding a tag.

```
[edit]
lab@R4# show policy-options policy-statement isis-to-rip
term 1 {
    from protocol isis;
    then {
        metric 1;
        accept;
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}
```

```
[edit]
lab@R4# set policy-options policy-statement isis-to-rip term 1 then tag 1234

[edit]
lab@R4# commit
commit complete
```

The routing policy `isis-to-rip` exports the IS-IS prefixes with metric 1 and tag of 1234.

```
[edit]
lab@R4# show policy-options policy-statement isis-to-rip
term 1 {
    from protocol isis;
    then {
        metric 1;
        tag 1234;
        accept;
    }
}
```

Modify the policy on R5 the same way.

```
[edit]
lab@R5# set policy-options policy-statement isis-to-rip term 1 then tag 1234

[edit]
lab@R5# commit
commit complete
```

The result is that R5 receives the RIP routes advertised by R4 with a tag of 1234.

```
[edit]
lab@R5# run show route 192.168/20

inet.0: 44 destinations, 46 routes (44 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.8.0/24     *[RIP/100] 02:46:02, metric 3, tag 1234
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.9.0/24     *[RIP/100] 02:46:02, metric 3, tag 1234
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.10.0/24    *[RIP/100] 02:46:02, metric 3, tag 1234
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.11.0/24    *[RIP/100] 02:46:02, metric 3, tag 1234
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.12.0/24     *[RIP/100] 01:50:48, metric 3, tag 1234
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.13.0/24     *[RIP/100] 01:50:48, metric 3, tag 1234
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.14.0/24     *[RIP/100] 01:50:48, metric 3, tag 1234
                    > to 172.30.0.58 via ge-0/0/4.204
192.168.15.0/24     *[RIP/100] 01:50:48, metric 3, tag 1234
                    > to 172.30.0.58 via ge-0/0/4.204
```

Now you can easily filter the routes by matching the tag on both R4 and R5.

R4

```
[edit]
lab@R4# set policy-options policy-statement rip-filter term 1 from protocol rip

[edit]
lab@R4# set policy-options policy-statement rip-filter term 1 from tag 1234

[edit]
lab@R4# set policy-options policy-statement rip-filter term 1 then reject

[edit]
lab@R4# set protocols rip group rip import rip-filter

[edit]
lab@R4# commit
commit complete
```

R5

```
[edit]
lab@R5# set policy-options policy-statement rip-filter term 1 from protocol rip

[edit]
lab@R5# set policy-options policy-statement rip-filter term 1 from tag 1234

[edit]
lab@R5# set policy-options policy-statement rip-filter term 1 then reject

[edit]
lab@R5# set protocols rip group rip import rip-filter

[edit]
lab@R5# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

After filtering the routes, R5 advertises the IS-IS routes.

```
[edit]
lab@R5# run show route advertising-protocol rip 172.30.0.57

inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[IS-IS/15] 00:32:06, metric 10
                     > to 172.30.0.34 via ae0.0
172.30.0.20/30      *[IS-IS/15] 00:32:44, metric 20
                     > to 172.30.0.29 via ge-0/0/4.145
172.30.0.24/30      *[IS-IS/15] 00:32:06, metric 30
                       to 172.30.0.34 via ae0.0
                     > to 172.30.0.29 via ge-0/0/4.145
172.30.5.3/32       *[IS-IS/15] 00:32:06, metric 20
                     > to 172.30.0.34 via ae0.0
                       to 172.30.0.29 via ge-0/0/4.145
172.30.5.4/32       *[IS-IS/15] 00:32:44, metric 10
                     > to 172.30.0.29 via ge-0/0/4.145
192.168.8.0/24      *[IS-IS/160] 00:32:06, metric 20
                     > to 172.30.0.34 via ae0.0
                       to 172.30.0.29 via ge-0/0/4.145
192.168.9.0/24      *[IS-IS/160] 00:32:06, metric 20
                       to 172.30.0.34 via ae0.0
                     > to 172.30.0.29 via ge-0/0/4.145
192.168.10.0/24     *[IS-IS/160] 00:32:06, metric 20
                     > to 172.30.0.34 via ae0.0
                       to 172.30.0.29 via ge-0/0/4.145
192.168.11.0/24     *[IS-IS/160] 00:32:06, metric 20
                     > to 172.30.0.34 via ae0.0
                       to 172.30.0.29 via ge-0/0/4.145
192.168.12.0/24     *[IS-IS/160] 00:32:06, metric 10
                     > to 172.30.0.34 via ae0.0
192.168.13.0/24     *[IS-IS/160] 00:32:06, metric 10
                     > to 172.30.0.34 via ae0.0
192.168.14.0/24     *[IS-IS/160] 00:32:06, metric 10
                     > to 172.30.0.34 via ae0.0
192.168.15.0/24     *[IS-IS/160] 00:32:06, metric 10
                     > to 172.30.0.34 via ae0.0
```

11) Verify L1/L2 summarization.

R4 and R5 are redistributing all RIP prefixes to IS-IS Level 1.

```
[edit]
lab@R4# show policy-options policy-statement rip-to-isis
term 1 {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    from protocol rip;
    then accept;
}
```

The default IS-IS behavior is to not distribute Level 1 external routes to Level 2. This way to meet the task requirement that only a summary route of all RIP routes should be advertised to Level 2, you have to manually advertise this summary route. To instruct R3 and R6 as L1/L2 routers to advertise summary route you have to use routing policy. Routing policy `l1-to-l2` is preconfigured on R3 and R6.

```
[edit]
lab@R3# show policy-options policy-statement l1-to-l2
term 1 {
    from {
        protocol aggregate;
        route-filter 172.30.32.0/20 exact;
    }
    to level 2;
    then accept;
}
```

This is an example of using the `to` matching statement for routing policy. The aggregate route is advertised only to Level 2 LSPs.

The aggregate route must be manually created.

```
[edit]
lab@R3# show routing-options aggregate
route 172.30.32.0/20;
```

Verify on R3 if the summary route is advertised and if more specific prefixes are suppressed.

```
[edit]
lab@R3# run show isis database level 2 R3.00-00 extensive
IS-IS level 2 link-state database:

R3.00-00 Sequence: 0xd, Checksum: 0xb8ab, Lifetime: 1092 secs
   IS neighbor: R2.00                         Metric:       10
     Two-way fragment: R2.00-00, Two-way first fragment: R2.00-00
   IP prefix: 172.30.0.12/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.20/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.24/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.28/30                  Metric:       20 Internal Up
   IP prefix: 172.30.0.32/30                  Metric:       30 Internal Up
   IP prefix: 172.30.5.3/32                   Metric:        0 Internal Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
   IP prefix: 172.30.5.4/32               Metric:       10 Internal Up
   IP prefix: 172.30.5.5/32               Metric:       20 Internal Up
   IP prefix: 172.30.32.0/20              Metric:       10 External Up
   IP prefix: 172.30.32.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.33.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.34.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.35.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.36.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.37.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.38.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.39.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.40.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.41.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.42.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.43.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.44.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.45.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.46.0/24              Metric:       12 Internal Up
   IP prefix: 172.30.47.0/24              Metric:       12 Internal Up
   IP prefix: 192.168.8.0/24              Metric:        0 External Up
   IP prefix: 192.168.9.0/24              Metric:        0 External Up
   IP prefix: 192.168.10.0/24             Metric:        0 External Up
   IP prefix: 192.168.11.0/24             Metric:        0 External Up
   V6 prefix: fd17:f0f4:f691:5::3/128     Metric:        0 Internal Up
   V6 prefix: fd17:f0f4:f691:5::4/128     Metric:       10 Internal Up
   V6 prefix: fd17:f0f4:f691:5::5/128     Metric:       20 Internal Up
...

[edit]
lab@R3# run show isis database level 2 R6.00-00 extensive
IS-IS level 2 link-state database:

R6.00-00 Sequence: 0xb, Checksum: 0xcb74, Lifetime: 1167 secs
   IS neighbor: R7.00                      Metric:       10
     Two-way fragment: R7.00-00, Two-way first fragment: R7.00-00
   IP prefix: 172.30.32.0/20               Metric:       10 External Up
   IP prefix: 192.168.12.0/24              Metric:        0 External Up
   IP prefix: 192.168.13.0/24              Metric:        0 External Up
   IP prefix: 192.168.14.0/24              Metric:        0 External Up
   IP prefix: 192.168.15.0/24              Metric:        0 External Up
...
```

Summary route is advertised by both R3 and R6, but all more specific routes are advertised also. The automatic suppressing of external routes is not working for some reason.

Note: R6 LSP has limited prefixes, as shown before. The reason of that will be revealed later in the troubleshooting steps.

© 2020 Juniper Networks, Inc. All rights reserved.

Verify on R3 the LSPs from R4 and R5, which they flood into Level 1. This will show you the reason of this behavior.

```
[edit]
lab@R3# run show isis database level 1 R5.00-00 extensive
IS-IS level 1 link-state database:

R5.00-00 Sequence: 0xc, Checksum: 0xb386, Lifetime: 1138 secs
   IS neighbor: R4.00                         Metric:       10
     Two-way fragment: R4.00-00, Two-way first fragment: R4.00-00
   IS neighbor: R6.00                         Metric:       10
     Two-way fragment: R6.00-00, Two-way first fragment: R6.00-00
   IP prefix: 172.30.0.28/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.32/30                  Metric:       10 Internal Up
   IP prefix: 172.30.5.5/32                   Metric:        0 Internal Up
   IP prefix: 172.30.32.0/24                  Metric:        2 External Up
   IP prefix: 172.30.33.0/24                  Metric:        2 External Up
   IP prefix: 172.30.34.0/24                  Metric:        2 External Up
   IP prefix: 172.30.35.0/24                  Metric:        2 External Up
   IP prefix: 172.30.36.0/24                  Metric:        2 External Up
   IP prefix: 172.30.37.0/24                  Metric:        2 External Up
   IP prefix: 172.30.38.0/24                  Metric:        2 External Up
   IP prefix: 172.30.39.0/24                  Metric:        2 External Up
   IP prefix: 172.30.40.0/24                  Metric:        2 External Up
   IP prefix: 172.30.41.0/24                  Metric:        2 External Up
   IP prefix: 172.30.42.0/24                  Metric:        2 External Up
   IP prefix: 172.30.43.0/24                  Metric:        2 External Up
   IP prefix: 172.30.44.0/24                  Metric:        2 External Up
   IP prefix: 172.30.45.0/24                  Metric:        2 External Up
   IP prefix: 172.30.46.0/24                  Metric:        2 External Up
   IP prefix: 172.30.47.0/24                  Metric:        2 External Up
   V6 prefix: fd17:f0f4:f691:5::5/128         Metric:        0 Internal Up
...
```

R5 is advertising the RIP prefixes as External. This is expected and should not lead to re-advertising those routes to Level 2 by the L1/L2 routers.

```
[edit]
lab@R3# run show isis database level 1 R4.00-00 extensive
IS-IS level 1 link-state database:

R4.00-00 Sequence: 0xb, Checksum: 0xfedf, Lifetime: 1101 secs
   IS neighbor: R3.00                         Metric:       10
     Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00
   IS neighbor: R5.00                         Metric:       10
     Two-way fragment: R5.00-00, Two-way first fragment: R5.00-00
   IP prefix: 172.30.0.20/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.28/30                  Metric:       10 Internal Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
IP prefix: 172.30.5.4/32              Metric:        0 Internal Up
IP prefix: 172.30.32.0/24             Metric:        2 Internal Up
IP prefix: 172.30.33.0/24             Metric:        2 Internal Up
IP prefix: 172.30.34.0/24             Metric:        2 Internal Up
IP prefix: 172.30.35.0/24             Metric:        2 Internal Up
IP prefix: 172.30.36.0/24             Metric:        2 Internal Up
IP prefix: 172.30.37.0/24             Metric:        2 Internal Up
IP prefix: 172.30.38.0/24             Metric:        2 Internal Up
IP prefix: 172.30.39.0/24             Metric:        2 Internal Up
IP prefix: 172.30.40.0/24             Metric:        2 Internal Up
IP prefix: 172.30.41.0/24             Metric:        2 Internal Up
IP prefix: 172.30.42.0/24             Metric:        2 Internal Up
IP prefix: 172.30.43.0/24             Metric:        2 Internal Up
IP prefix: 172.30.44.0/24             Metric:        2 Internal Up
IP prefix: 172.30.45.0/24             Metric:        2 Internal Up
IP prefix: 172.30.46.0/24             Metric:        2 Internal Up
IP prefix: 172.30.47.0/24             Metric:        2 Internal Up
V6 prefix: fd17:f0f4:f691:5::4/128    Metric:        0 Internal Up
...
```

The LSP of R4 shows the RIP prefixes as Internal and they are treated as other Internal Level 1 routes and are flooded into IS-IS Level 2 by the L1/L2 routers.

12) Fix the external route type advertised by R4.

R4 IS-IS is configured for using `wide-metrics-only` for Level 1.

```
[edit]
lab@R4# show protocols isis level 1
wide-metrics-only;
```

It was stated in this chapter overview that a command must be used to force IS-IS to use wide-metric for metric calculations. IS-IS uses a different TLV (135) for advertising IP prefixes with a wide-metric. TLV 135 “Extended IP Reachability TLV” does not contain the External/Internal bit, which is used by TLV 130 “IP External Reachability TLV”. All prefixes advertised by TLV 135 have wide metric and are seen as Internal by IS-IS.

Delete the `wide-metric-only` from the R4 configuration will resolve the issue.

```
[edit]
lab@R4# delete protocols isis level 1 wide-metrics-only

[edit]
lab@R4# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify the level 2 LSP of R3 if the more specific routes are suppressed.

```
[edit]
lab@R3# run show isis database level 2 R3.00-00 extensive | find 172.30.32.0
   IP prefix: 172.30.32.0/20              Metric:       10 External Up
   IP prefix: 192.168.8.0/24              Metric:        0 External Up
   IP prefix: 192.168.9.0/24              Metric:        0 External Up
   IP prefix: 192.168.10.0/24             Metric:        0 External Up
   IP prefix: 192.168.11.0/24             Metric:        0 External Up
   V6 prefix: fd17:f0f4:f691:5::3/128     Metric:        0 Internal Up
   V6 prefix: fd17:f0f4:f691:5::4/128     Metric:       10 Internal Up
   V6 prefix: fd17:f0f4:f691:5::5/128     Metric:       20 Internal Up
...
```

Only the aggregated route is advertised after deleting the `wide-metric-only` command.

Verify the LSP advertised by R4 after the above change.

```
[edit]
lab@R3# run show isis database level 1 R4.00-00 extensive | no-more
IS-IS level 1 link-state database:

R4.00-00 Sequence: 0xc, Checksum: 0x8240, Lifetime: 1136 secs
   IS neighbor: R3.00                     Metric:       10
     Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00
   IS neighbor: R5.00                     Metric:       10
     Two-way fragment: R5.00-00, Two-way first fragment: R5.00-00
   IP prefix: 172.30.0.20/30              Metric:       10 Internal Up
   IP prefix: 172.30.0.28/30              Metric:       10 Internal Up
   IP prefix: 172.30.5.4/32               Metric:        0 Internal Up
   IP prefix: 172.30.32.0/24              Metric:        2 External Up
   IP prefix: 172.30.33.0/24              Metric:        2 External Up
   IP prefix: 172.30.34.0/24              Metric:        2 External Up
   IP prefix: 172.30.35.0/24              Metric:        2 External Up
   IP prefix: 172.30.36.0/24              Metric:        2 External Up
   IP prefix: 172.30.37.0/24              Metric:        2 External Up
   IP prefix: 172.30.38.0/24              Metric:        2 External Up
   IP prefix: 172.30.39.0/24              Metric:        2 External Up
   IP prefix: 172.30.40.0/24              Metric:        2 External Up
   IP prefix: 172.30.41.0/24              Metric:        2 External Up
   IP prefix: 172.30.42.0/24              Metric:        2 External Up
   IP prefix: 172.30.43.0/24              Metric:        2 External Up
   IP prefix: 172.30.44.0/24              Metric:        2 External Up
   IP prefix: 172.30.45.0/24              Metric:        2 External Up
   IP prefix: 172.30.46.0/24              Metric:        2 External Up
   IP prefix: 172.30.47.0/24              Metric:        2 External Up
   V6 prefix: fd17:f0f4:f691:5::4/128     Metric:        0 Internal Up
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

The redistributed prefixes are now advertised as external.

13) Verify loopback reachability.

a) R1

```
[edit]
lab@R1# run show route 172.30.5/24

inet.0: 33 destinations, 33 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32      *[Direct/0] 05:38:44
                    > via lo0.0
172.30.5.2/32      *[IS-IS/18] 02:10:36, metric 10
                    > to 172.30.0.2 via ae0.0
172.30.5.3/32      *[IS-IS/18] 02:09:51, metric 20
                    > to 172.30.0.2 via ae0.0
172.30.5.4/32      *[IS-IS/18] 00:35:15, metric 30
                    > to 172.30.0.2 via ae0.0
172.30.5.5/32      *[IS-IS/18] 01:25:27, metric 40
                    > to 172.30.0.2 via ae0.0
172.30.5.7/32      *[IS-IS/18] 02:06:36, metric 20
                    > to 172.30.0.2 via ae0.0
172.30.5.8/32      *[IS-IS/15] 01:28:21, metric 10
                    > to 172.30.0.10 via ge-0/0/4.118
```

R1 receives all loopback IPs except the one of R6. Based on the preference you can recognize which routes are learned from the Level 1 LSP or Level 2 LSP. A preference of 15 is assigned to Level 1 Internal routes and a preference of 18 is for Level 2 Internal routes.

b) R3

```
[edit]
lab@R3# run show route 172.30.5/24

inet.0: 52 destinations, 52 routes (52 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32      *[IS-IS/18] 02:10:00, metric 20
                    > to 172.30.0.13 via ge-0/0/4.123
172.30.5.2/32      *[IS-IS/18] 02:10:10, metric 10
                    > to 172.30.0.13 via ge-0/0/4.123
172.30.5.3/32      *[Direct/0] 05:37:48
                    > via lo0.0
172.30.5.4/32      *[IS-IS/15] 00:35:35, metric 10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.22 via ge-0/0/4.134
172.30.5.5/32       *[IS-IS/15] 00:35:35, metric 20
                    > to 172.30.0.22 via ge-0/0/4.134
                      to 172.30.0.26 via ge-0/0/4.136
172.30.5.7/32       *[IS-IS/18] 02:06:55, metric 20
                    > to 172.30.0.13 via ge-0/0/4.123
172.30.5.8/32       *[IS-IS/18] 02:06:03, metric 30
                    > to 172.30.0.13 via ge-0/0/4.123
```

R3 has learned all loopback IP addresses except the loopback IP address of R6. The loopbacks of R4 and R5 are known as Level 1 Internal routes, others are known from Level 2.

c) R4

```
[edit]
lab@R4# run show route 172.30.5/24

inet.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.3/32       *[IS-IS/15] 00:35:52, metric 10
                    > to 172.30.0.21 via ge-0/0/4.134
172.30.5.4/32       *[Direct/0] 05:34:04
                    > via lo0.0
172.30.5.5/32       *[IS-IS/15] 00:36:18, metric 10
                    > to 172.30.0.30 via ge-0/0/4.145
```

R4 has only three loopback IP addresses. As a Level 1 router it should use a default route for routes outside the Area.

```
[edit]
lab@R4# run show route 0/0 exact

inet.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[IS-IS/15] 00:36:12, metric 10
                    > to 172.30.0.21 via ge-0/0/4.134
```

But again, R6s loopback IP address is missing. As a L1/L2 router, its' IS-IS interface routes should be sent in Level 1 and Level 2 LSPs.

Verify on R3 the Level 1 and Level 2 LSPs generated from R6.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show isis database level 1 R6.00-00 extensive
IS-IS level 1 link-state database:

R6.00-00 Sequence: 0xb, Checksum: 0x7e8f, Lifetime: 1187 secs
   IS neighbor: R3.00                         Metric:       10
     Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00
   IS neighbor: R5.00                         Metric:       10
     Two-way fragment: R5.00-00, Two-way first fragment: R5.00-00
   IP prefix: 192.168.12.0/24                 Metric:        0 External Up
   IP prefix: 192.168.13.0/24                 Metric:        0 External Up
   IP prefix: 192.168.14.0/24                 Metric:        0 External Up
   IP prefix: 192.168.15.0/24                 Metric:        0 External Up
...
```

R6 advertises four prefixes to Level 1 (see export policy).

```
[edit]
lab@R3# run show isis database level 2 R6.00-00 extensive
IS-IS level 2 link-state database:

R6.00-00 Sequence: 0xb, Checksum: 0xcb74, Lifetime: 1167 secs
   IS neighbor: R7.00                         Metric:       10
     Two-way fragment: R7.00-00, Two-way first fragment: R7.00-00
   IP prefix: 172.30.32.0/20                  Metric:       10 External Up
   IP prefix: 192.168.12.0/24                 Metric:        0 External Up
   IP prefix: 192.168.13.0/24                 Metric:        0 External Up
   IP prefix: 192.168.14.0/24                 Metric:        0 External Up
   IP prefix: 192.168.15.0/24                 Metric:        0 External Up
...
```

R6 announces the same four external routes plus an aggregate route to Level 2.

14) Fix R6s loopback reachability issue.

R6 announces only prefixes, which are exported by routing policies. Verifying these policies might reveal the reason of the missing interface routes (including the loopback IP address).

```
[edit]
lab@R6# show protocols isis export
export [ inject-externals l1-to-l2 ];

[edit]
lab@R6# show policy-options policy-statement inject-externals
term 1 {
    from {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        protocol static;
        route-filter 192.168.12.0/22 longer;
    }
    then accept;
}

[edit]
lab@R6# show policy-options policy-statement l1-to-l2
term 1 {
    from {
        protocol aggregate;
        route-filter 172.30.32.0/20 exact;
    }
    to level 2;
    then accept;
}
term 2 {
    then reject;
}
```

Both policies are advertising correctly the routes visible in the LSPs. However, policy `l1-to-l2` has a second term, which rejects every route that is not matched. This way all other routes that are matched by this term are rejected. In other words, they are not evaluated anymore by the default policy, which accepts all interface routes enabled for IS-IS.

Deleting `term 2` will allow IS-IS to advertise the missing interface routes.

```
[edit]
lab@R6# delete policy-options policy-statement l1-to-l2 term 2

[edit]
lab@R6# commit
commit complete
```

Verify if the change has effect and the loopback IP address of R6 is reachable.

On R5 Level 1.

```
[edit]
lab@R5# run show route 172.30.5/24

inet.0: 45 destinations, 45 routes (45 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.3/32      *[IS-IS/15] 00:44:48, metric 20
                      to 172.30.0.34 via ae0.0
                    > to 172.30.0.29 via ge-0/0/4.145
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.4/32      *[IS-IS/15] 00:45:13, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.5.5/32      *[Direct/0] 05:44:11
                    > via lo0.0
172.30.5.6/32      *[IS-IS/15] 00:01:01, metric 10
                    > to 172.30.0.34 via ae0.0
```

On R2 Level 2.

```
[edit]
lab@R2# run show route 172.30.5/24

inet.0: 33 destinations, 33 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32      *[IS-IS/18] 02:21:08, metric 10
                    > to 172.30.0.1 via ae0.0
172.30.5.2/32      *[Direct/0] 05:49:17
                    > via lo0.0
172.30.5.3/32      *[IS-IS/18] 02:20:23, metric 10
                    > to 172.30.0.14 via ge-0/0/4.123
172.30.5.4/32      *[IS-IS/18] 00:45:48, metric 20
                    > to 172.30.0.14 via ge-0/0/4.123
172.30.5.5/32      *[IS-IS/18] 00:02:01, metric 30
                      to 172.30.0.14 via ge-0/0/4.123
                    > to 172.30.0.18 via ge-0/0/4.127
172.30.5.6/32      *[IS-IS/18] 00:02:01, metric 20
                    > to 172.30.0.14 via ge-0/0/4.123
                      to 172.30.0.18 via ge-0/0/4.127
172.30.5.7/32      *[IS-IS/18] 02:17:08, metric 10
                    > to 172.30.0.18 via ge-0/0/4.127
172.30.5.8/32      *[IS-IS/18] 01:38:53, metric 20
                      to 172.30.0.1 via ae0.0
                    > to 172.30.0.18 via ge-0/0/4.127
```

The loopback IP address of R6 is reachable after fixing the `l1-to-l2` export policy. After this change, R6 LSP for Level 2 holds all other Internal prefixes that were missing before.

```
[edit]
lab@R3# run show isis database level 2 R6.00-00 extensive
IS-IS level 2 link-state database:

R6.00-00 Sequence: 0xc, Checksum: 0x8956, Lifetime: 1137 secs
   IS neighbor: R7.00                          Metric:       10
     Two-way fragment: R7.00-00, Two-way first fragment: R7.00-00
   IP prefix: 172.30.0.20/30                   Metric:       20 Internal Up
   IP prefix: 172.30.0.24/30                   Metric:       10 Internal Up
   IP prefix: 172.30.0.28/30                   Metric:       20 Internal Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
   IP prefix: 172.30.0.32/30              Metric:       10 Internal Up
   IP prefix: 172.30.0.40/30              Metric:       10 Internal Up
   IP prefix: 172.30.5.3/32               Metric:       10 Internal Up
   IP prefix: 172.30.5.4/32               Metric:       20 Internal Up
   IP prefix: 172.30.5.5/32               Metric:       10 Internal Up
   IP prefix: 172.30.5.6/32               Metric:        0 Internal Up
   IP prefix: 172.30.32.0/20              Metric:       10 External Up
   IP prefix: 192.168.12.0/24             Metric:        0 External Up
   IP prefix: 192.168.13.0/24             Metric:        0 External Up
   IP prefix: 192.168.14.0/24             Metric:        0 External Up
   IP prefix: 192.168.15.0/24             Metric:        0 External Up
   V6 prefix: fd17:f0f4:f691:5::4/128     Metric:       20 Internal Up
   V6 prefix: fd17:f0f4:f691:5::5/128     Metric:       10 Internal Up
   V6 prefix: fd17:f0f4:f691:5::6/128     Metric:        0 Internal Up
...
```

15) Test the redundancy.

From R8 test the connectivity to the RIP routes advertised by the VR with both R3 and R6 up and running. As required, R8 has only a summary route in the routing table for all RIP prefixes.

```
[edit]
lab@R8# run show route 172.30.32/20

inet.0: 34 destinations, 34 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/20     *[IS-IS/165] 00:12:03, metric 30
                    > to 172.30.0.45 via ge-0/0/4.178

[edit]
lab@R8# run traceroute 172.30.32.2 no-resolve
traceroute to 172.30.32.2 (172.30.32.2), 30 hops max, 40 byte packets
 1  172.30.0.45  7.515 ms  7.665 ms  7.973 ms
 2  172.30.0.41  7.786 ms  7.863 ms  10.565 ms
 3  172.30.0.33  8.161 ms  8.185 ms  7.696 ms
 4  172.30.0.58  11.884 ms !N  4.461 ms !N  4.909 ms !N
```

The traceroute shows that the path goes via R8-R7-R6-R5-VR.

Deactivate the `isis` hierarchy on R6.

```
[edit]
lab@R6# deactivate protocols isis
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# commit
commit complete
```

Again, perform a traceroute from R8. You may need to wait a minute or two for the path to synchronize.

```
[edit]
lab@R8# run traceroute 172.30.32.2 no-resolve
traceroute to 172.30.32.2 (172.30.32.2), 30 hops max, 40 byte packets
 1  172.30.0.45  8.112 ms  8.028 ms  7.905 ms
 2  172.30.0.17  8.085 ms  7.846 ms  7.708 ms
 3  172.30.0.14  7.657 ms  7.942 ms  7.777 ms
 4  172.30.0.22  8.574 ms  13.931 ms  7.949 ms
 5  172.30.0.50  5.298 ms !N  4.544 ms !N  5.117 ms !N
```

The path changed now via R3-R8-R7-R2-R3-R4-VR.

Reactivate the IS-IS protocol on R6.

```
[edit]
lab@R6# rollback 1
load complete

[edit]
lab@R6# commit
commit complete
```

As a last step, remember to write down the changes you performed to fix the issues.

Task 3: IGP Rollout

Task 3 contains several requirements and limitations for the IGP rollout. The configuration from the end of chapter one is used as the base configuration. Figure 6 describes the topology used for the IGP network.

1) Load the end configuration from chapter one on all eight devices.

```
[edit]
lab@R1# load override ch1-end.config

[edit]
lab@R1# commit
commit complete
```

2) Configure the additional interfaces on R4 and R5.

© 2020 Juniper Networks, Inc. All rights reserved.

Router	Interface Name	IP Address	IPv6 Address
R4	ge-0/0/4.202	172.30.0.49/30	
	ge-0/0/4.203	172.30.0.53/30	link-local
R5	ge-0/0/4.204	172.30.0.57/30	
	ge-0/0/4.205	172.30.0.61/30	link-local

Table 8

a) R4.

You have to configure two new subinterfaces connected to DC2 and DC3 routers according to Table 8.

```
[edit]
lab@R4# set interfaces ge-0/0/4.202 vlan-id 202

[edit]
lab@R4# set interfaces ge-0/0/4.202 description "DC2 connection"

[edit]
lab@R4# set interfaces ge-0/0/4.202 family inet address 172.30.0.49/30

[edit]
lab@R4# set interfaces ge-0/0/4.203 vlan-id 203

[edit]
lab@R4# set interfaces ge-0/0/4.203 description "DC3 connection"

[edit]
lab@R4# set interfaces ge-0/0/4.203 family inet address 172.30.0.53/30

[edit]
lab@R4# set interfaces ge-0/0/4.203 family inet6

[edit]
lab@R4# commit
commit complete
```

As the subnet mask of the IP address is 30 bits you can easily guess the IP address on other side. With ping you can verify the connection over the configured subinterfaces.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run ping 172.30.0.50 rapid
PING 172.30.0.50 (172.30.0.50): 56 data bytes
!!!!!
--- 172.30.0.50 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.440/1.782/2.766/0.494 ms

[edit]
lab@R4# run ping 172.30.0.54 rapid
PING 172.30.0.54 (172.30.0.54): 56 data bytes
!!!!!
--- 172.30.0.54 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.486/7.077/24.906/8.970 ms
```

b) R5

You also have to configure the interfaces on R5 to DC2 and DC3.

```
[edit]
lab@R5# set interfaces ge-0/0/4.204 vlan-id 204

[edit]
lab@R5# set interfaces ge-0/0/4.204 description "DC2 connection"

[edit]
lab@R5# set interfaces ge-0/0/4.204 family inet address 172.30.0.57/30

[edit]
lab@R5# set interfaces ge-0/0/4.205 vlan-id 205

[edit]
lab@R5# set interfaces ge-0/0/4.205 description "DC3 connection"

[edit]
lab@R5# set interfaces ge-0/0/4.205 family inet address 172.30.0.61/30

[edit]
lab@R5# set interfaces ge-0/0/4.205 family inet6

[edit]
lab@R5# commit
commit complete
```

Verify the connectivity with ping.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# run ping 172.30.0.58 rapid
PING 172.30.0.58 (172.30.0.58): 56 data bytes
!!!!!
--- 172.30.0.58 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.242/2.093/3.000/0.733 ms

[edit]
lab@R5# run ping 172.30.0.62 rapid
PING 172.30.0.62 (172.30.0.62): 56 data bytes
!!!!!
--- 172.30.0.62 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.373/6.166/24.120/8.982 ms
```

3) Configure the IS-IS protocol.

IS-IS is the IGP used in our backbone network. Table 9 contains the NET addresses of all 8 routers. According to Figure 6 all 8 routers should have only Level 1 adjacencies.

Because the configuration on all 8 routers is the same, only the configuration of R1 will be described in detail. The differences, if any, will be described for the appropriate routers.

a) Configure `family iso` on all core-facing interfaces.

To enable the IS-IS protocol you need to activate `family iso` on the interfaces where adjacencies have to be established.

Junos supports a very nice feature for applying common configurations on multiple places in the configuration called "apply-groups". Apply-groups support match conditions to define where the configuration from the apply-group must be applied.

In the particular step you have to configure several interfaces with `family iso`. You can use an apply-group to enable it on the core-facing interfaces. All core-facing interfaces are sub interfaces of ge-0/0/4 and aggregated Ethernet. So, you can easily match them.

Repeat this on all eight devices.

```
[edit]
lab@R1# set groups if-families interfaces ge-0/0/4 unit <*> family iso

[edit]
lab@R1# set groups if-families interfaces <ae*> unit <*> family iso
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: For match conditions you can use wildcards similar to regular expressions. You must use angle brackets "< >" to define the match condition with wildcards. "*" is a very commonly used wildcard which matches zero or more characters. If no wildcards are used, Junos matches exactly that part of the configuration which is specified. Apply-groups are configured under the `groups` hierarchy.

In the previously shown configuration the apply-group named *if-families* matches each unit on interface `ge-0/0/4` and applies `family iso` to them. The second line adds `family iso` to each unit of all configured aggregated interfaces.

You must apply the configured group and commit the configuration.

```
[edit]
lab@R1# set apply-groups if-families

[edit]
lab@R1# commit
commit complete
```

To verify the apply-groups, you must display the configuration with an additional option.

```
[edit]
lab@R1# show interfaces | display inheritance
ge-0/0/0 {
    unit 0 {
        description "OoB management";
        family inet {
            address 10.10.1.1/24;
        }
    }
}
ge-0/0/1 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/0/2 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/0/4 {
    vlan-tagging;
    unit 114 {
        description "R4 connection";
        vlan-id 114;
        family inet {
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        address 172.30.0.5/30;
    }
    ##
    ## 'iso' was inherited from group 'if-families'
    ##
    family iso;
}
unit 118 {
    description "R8 connection";
    vlan-id 118;
    family inet {
        address 172.30.0.9/30;
    }
    ##
    ## 'iso' was inherited from group 'if-families'
    ##
    family iso;
    family inet6;
}
}
ae0 {
    aggregated-ether-options {
        lacp {
            active;
        }
    }
    unit 0 {
        description "R2 connection";
        family inet {
            address 172.30.0.1/30;
        }
        ##
        ## 'iso' was inherited from group 'if-families'
        ##
        family iso;
        family inet6;
    }
}
lo0 {
    unit 0 {
        family inet {
            filter {
                input protect-re;
            }
            address 172.30.5.1/32;
        }
        family inet6 {
            address fd17:f0f4:f691:5::1/128;
        }
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    }
}
```

With the " | `display inheritance`" option, inherited configuration statements from the apply-groups are added in the show configuration output.

Note: Although some of the devices do not have aggregated Ethernet interfaces you still can commit the configuration with apply-groups.

Note: The loopback interface is not included in the apply-group configuration, as the NET address is usually configured under `family iso`. It will be configured manually in the next step.

b) Configure NET addresses.

The NET addresses are given in Table 9. All 8 routers are in the same AreaID. The SysID part from the NET address is built from the loopback IP address.

Router	NET
R1	49.0001.1720.3000.5001.00
R2	49.0001.1720.3000.5002.00
R3	49.0001.1720.3000.5003.00
R4	49.0001.1720.3000.5004.00
R5	49.0001.1720.3000.5005.00
R6	49.0001.1720.3000.5006.00
R7	49.0001.1720.3000.5007.00
R8	49.0001.1720.3000.5008.00

Table 9

R1

```
[edit]
lab@R1# set interface lo0.0 family iso address 49.0001.1720.3000.5001.00
```

R2

```
[edit]
lab@R2# set interface lo0.0 family iso address 49.0001.1720.3000.5002.00
```

R3

```
[edit]
lab@R3# set interface lo0.0 family iso address 49.0001.1720.3000.5003.00
```


© 2020 Juniper Networks, Inc. All rights reserved.

R4

```
[edit]
lab@R4# set interface lo0.0 family iso address 49.0001.1720.3000.5004.00
```

R5

```
[edit]
lab@R5# set interface lo0.0 family iso address 49.0001.1720.3000.5005.00
```

R6

```
[edit]
lab@R6# set interface lo0.0 family iso address 49.0001.1720.3000.5006.00
```

R7

```
[edit]
lab@R7# set interface lo0.0 family iso address 49.0001.1720.3000.5007.00
```

R8

```
[edit]
lab@R8# set interface lo0.0 family iso address 49.0001.1720.3000.5008.00
```

Note: The configuration is not committed after this change. The commit will be executed when all configurations has been performed in this step.

c) Configure the Router IDs.

The router ID of the routers should be configured explicitly and reflect the loopback IP address.

R1

```
[edit]
lab@R1# set routing-options router-id 172.30.5.1
```

R2

```
[edit]
lab@R2# set routing-options router-id 172.30.5.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

R3

```
[edit]
lab@R3# set routing-options router-id 172.30.5.3
```

R4

```
[edit]
lab@R4# set routing-options router-id 172.30.5.4
```

R5

```
[edit]
lab@R5# set routing-options router-id 172.30.5.5
```

R6

```
[edit]
lab@R6# set routing-options router-id 172.30.5.6
```

R7

```
[edit]
lab@R7# set routing-options router-id 172.30.5.7
```

R8

```
[edit]
lab@R8# set routing-options router-id 172.30.5.8
```

All other routers should be configured in the same way.

d) Configure the IS-IS protocol.

Area 49.0001 requires Level 1 adjacencies as shown in Figure 6. As Level 2 adjacencies are not needed you can disable Level 2 completely.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols isis level 2 disable
```

The interface metric should be calculated based on the bandwidth of the interfaces.

© 2020 Juniper Networks, Inc. All rights reserved.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols isis reference-bandwidth 10g
```

10G is used for reference bandwidth. As 10G interfaces are common today, this way 10G and 1G interfaces will have metric reflecting their speed correctly. Narrow metrics should not be used.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols isis level 1 wide-metrics-only
```

Next you have to configure MD5 authentication for all PDUs. Authentication configured under the interface authenticates only IIH PDUs. Configure authentication directly under the Level authenticates all PDUs for the Level.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols isis level 1 authentication-type md5

[edit]
lab@R1# set protocols isis level 1 authentication-key jncie
```

You must add the interfaces for which IS-IS should be enabled and establish adjacencies under the protocol configuration. Also ensure that no pseudo node exist in the database.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols isis interface all point-to-point
```

With just one single line you configure all IS-IS interfaces for `point-to-point`. This command will be applied to all IS-IS all interfaces, which have `family iso` configured.

You can add the configuration for the other interfaces in the same way. If some interface needs to have a different configuration, you can add its configuration separately only for that interface. More specific interface configuration is preferred by Junos over less specific configuration statements.

The task states that the IS-IS neighbors should detect adjacency failures less than 6 seconds. IS-IS internal hello mechanism supports hold timer of minimum 3 seconds with a default of 9 seconds. The hello interval can be a minimum of 1 second. Junos allows to a

© 2020 Juniper Networks, Inc. All rights reserved.

configured hold timer of 1 second, which means the hello interval is 333ms. However, the requirement states to you are not allowed to alter the IS-IS timers.

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism used for detecting failures in the network. BFD timers can be set to sub second values and can be adaptively negotiated between neighbors depending on the rate of failures.

To enable BFD, you must configure `bfd-liveness-detection` under the interface configuration.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols isis interface all family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R1# set protocols isis interface all family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R1# commit
commit complete
```

`minimum-interval` in specified ms. It specifies the minimum interval that the router send hello messages as well the minimum interval the router expects to receive a reply. The receive interval can be configured separately.

`multiplier` specifies the number of unreceived hellos that causes the neighbor to be declared down. It is set to 3 by default.

The configuration of BFD with a minimum-interval of 150 and a multiplier of 3 will meet the task requirements

e) Configure the VRRP subnets on R3 and R4.

R3 and R4 have two interfaces configured with VRRP groups for gateway redundancy. The requirement is to have those interface routes in IS-IS but no adjacencies to be formed on those interfaces.

Configure the interface as `passive` under `protocols isis`, advertises the direct interface addresses without actually running the IS-IS protocol on the interface.

© 2020 Juniper Networks, Inc. All rights reserved.

R3

```
[edit]
lab@R3# set protocols isis interface ge-0/0/4.200 passive

[edit]
lab@R3# set protocols isis interface ge-0/0/4.201 passive

[edit]
lab@R1# commit
commit complete
```

R4

```
[edit]
lab@R4# set protocols isis interface ge-0/0/4.200 passive

[edit]
lab@R4# set protocols isis interface ge-0/0/4.201 passive

[edit]
lab@R1# commit
commit complete
```

f) Verify the configuration.

Ensure that the IS-IS adjacencies are established.

The topology in Figure 6 shows 3 adjacencies per router. Verify the adjacency tables on R1, R3, R6 and R8.

```
[edit]
lab@R1# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 R2               1  Up                 20
ge-0/0/4.114          R4               1  Up                 26
ge-0/0/4.118          R8               1  Up                 25
```

```
[edit]
lab@R3# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.123          R2               1  Up                 23
ge-0/0/4.134          R4               1  Up                 25
ge-0/0/4.136          R6               1  Up                 26
```

```
[edit]
lab@R6# run show isis adjacency
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Interface             System         L State       Hold (secs) SNPA
ae0.0                 R5             1  Up                25
ge-0/0/4.136          R3             1  Up                21
ge-0/0/4.167          R7             1  Up                24
```

```
[edit]
lab@R8# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/4.118          R1             1  Up                22
ge-0/0/4.158          R5             1  Up                20
ge-0/0/4.178          R7             1  Up                23
```

The adjacency tables meet our expectations as every router has 3 Level 1 adjacencies.

Next you can verify the link state database on some of the routers. IS-IS is a link state protocol and link state databases should be the same across all routers in one Area.

```
[edit]
lab@R1# run show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence  Checksum  Lifetime  Attributes
R1.00-00                  0xd    0x16f1      1178   L1
R2.00-00                  0xe    0xbd4f      1034   L1
R3.00-00                  0xe    0xe1a2      1174   L1
R4.00-00                  0x8    0xefae      1176   L1
R5.00-00                  0xc    0xa426      1150   L1
R6.00-00                  0x8    0x45bb      1057   L1
R7.00-00                  0x6    0x614a      1073   L1
R8.00-00                  0x6    0x5fe1      1081   L1
  8 LSPs

IS-IS level 2 link-state database:
  0 LSPs
```

```
[edit]
lab@R4# run show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence  Checksum  Lifetime  Attributes
R1.00-00                  0x19   0xc3d5       962  L1
R2.00-00                  0x1b   0x8904      1015  L1
R3.00-00                  0x13   0x4b46      1008  L1
R4.00-00                 0x1b   0x8843      1035  L1
R5.00-00                 0x1d   0xbb26      1196  L1
R6.00-00                  0xc   0xd7dd       885  L1
R7.00-00                  0xb   0x5fca       871  L1
R8.00-00                  0xb   0x142b       824  L1
  8 LSPs
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
IS-IS level 2 link-state database:
  0 LSPs
```

The link state databases on R1 and R4 have the same 8 LSPs. Only Level 1 LSPs are present, no pseudo node LSPs are advertised. The attached bit is missing because all routers are in Level 1.

```
[edit]
lab@R1# run show isis database level 1 R2.00-00 extensive
IS-IS level 1 link-state database:

R2.00-00 Sequence: 0xe, Checksum: 0xbd4f, Lifetime: 955 secs
   IS neighbor: R1.00                          Metric:         5
     Two-way fragment: R1.00-00, Two-way first fragment: R1.00-00
   IS neighbor: R3.00                        Metric:        10
     Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00
   IS neighbor: R7.00                          Metric:        10
     Two-way fragment: R7.00-00, Two-way first fragment: R7.00-00
   IP prefix: 172.30.0.0/30                    Metric:         5 Internal Up
   IP prefix: 172.30.0.12/30                   Metric:        10 Internal Up
   IP prefix: 172.30.0.16/30                   Metric:        10 Internal Up
   IP prefix: 172.30.5.2/32                    Metric:         0 Internal Up
   V6 prefix: fd17:f0f4:f691:5::2/128          Metric:         0 Internal Up
...
```

You can see the metric is advertised by R2 are reflecting the interface bandwidth. Metric 10 is assigned for 1G interfaces, while metric value of 5 is used for Aggregated Ethernet interfaces.

Verify loopback reachability on R3 and R5.

```
[edit]
lab@R3# run show route 172.30.5/24

inet.0: 34 destinations, 34 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32      *[IS-IS/15] 00:21:58, metric 15
                    > to 172.30.0.13 via ge-0/0/4.123
172.30.5.2/32      *[IS-IS/15] 00:22:07, metric 10
                    > to 172.30.0.13 via ge-0/0/4.123
172.30.5.3/32      *[Direct/0] 00:32:22
                    > via lo0.0
172.30.5.4/32      *[IS-IS/15] 00:22:03, metric 10
                    > to 172.30.0.22 via ge-0/0/4.134
172.30.5.5/32      *[IS-IS/15] 00:18:58, metric 15
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.26 via ge-0/0/4.136
172.30.5.6/32       *[IS-IS/15] 00:18:58, metric 10
                    > to 172.30.0.26 via ge-0/0/4.136
172.30.5.7/32       *[IS-IS/15] 00:18:56, metric 20
                      to 172.30.0.13 via ge-0/0/4.123
                    > to 172.30.0.26 via ge-0/0/4.136
172.30.5.8/32       *[IS-IS/15] 00:17:10, metric 25
                      to 172.30.0.13 via ge-0/0/4.123
                    > to 172.30.0.26 via ge-0/0/4.136
```

```
[edit]
lab@R5# run show route 172.30.5/24

inet.0: 34 destinations, 34 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       *[IS-IS/15] 00:17:19, metric 20
                    > to 172.30.0.29 via ge-0/0/4.145
                      to 172.30.0.38 via ge-0/0/4.158
172.30.5.2/32       *[IS-IS/15] 00:17:19, metric 25
                    > to 172.30.0.34 via ae0.0
                      to 172.30.0.29 via ge-0/0/4.145
                      to 172.30.0.38 via ge-0/0/4.158
172.30.5.3/32       *[IS-IS/15] 00:19:07, metric 15
                    > to 172.30.0.34 via ae0.0
172.30.5.4/32       *[IS-IS/15] 00:22:02, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.5.5/32       *[Direct/0] 00:46:22
                    > via lo0.0
172.30.5.6/32       *[IS-IS/15] 00:19:07, metric 5
                    > to 172.30.0.34 via ae0.0
172.30.5.7/32       *[IS-IS/15] 00:19:05, metric 15
                    > to 172.30.0.34 via ae0.0
172.30.5.8/32       *[IS-IS/15] 00:17:21, metric 10
                    > to 172.30.0.38 via ge-0/0/4.158
```

R3 and R5 have routes in the routing table for all 8 loopback IP addresses.

Verify BFD sessions.

```
[edit]
lab@R2# run show bfd session extensive
                                             Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.30.0.1               Up        ae0.0          6.000     2.000         3
 Client ISIS L1, TX interval 2.000, RX interval 2.000
 Session up time 00:02:23
 Local diagnostic None, remote diagnostic None
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 Remote state Up, version 1
 Session type: Single hop BFD
 Min async interval 2.000, min slow interval 2.000
 Adaptive async TX interval 2.000, RX interval 2.000
 Local min TX interval 2.000, minimum RX interval 2.000, multiplier 3
 Remote min TX interval 2.000, min RX interval 2.000, multiplier 3
 Local discriminator 12033, remote discriminator 2055
 Echo mode disabled/inactive
 Remote is control-plane independent
  Session ID: 0x2a68

                                                Detect   Transmit
Address                 State     Interface     Time     Interval  Multiplier
172.30.0.14             Up        ge-0/0/4.123  6.000     2.000        3
 Client ISIS L1, TX interval 2.000, RX interval 2.000
 Session up time 00:02:19
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Session type: Single hop BFD
 Min async interval 2.000, min slow interval 2.000
 Adaptive async TX interval 2.000, RX interval 2.000
 Local min TX interval 2.000, minimum RX interval 2.000, multiplier 3
 Remote min TX interval 2.000, min RX interval 2.000, multiplier 3
 Local discriminator 12034, remote discriminator 12449
 Echo mode disabled/inactive
 Remote is control-plane independent
  Session ID: 0x2a69

                                                Detect   Transmit
Address                 State     Interface     Time     Interval  Multiplier
172.30.0.18             Up        ge-0/0/4.127  6.000     2.000        3
 Client ISIS L1, TX interval 2.000, RX interval 2.000
 Session up time 00:02:23
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Session type: Single hop BFD
 Min async interval 2.000, min slow interval 2.000
 Adaptive async TX interval 2.000, RX interval 2.000
 Local min TX interval 2.000, minimum RX interval 2.000, multiplier 3
 Remote min TX interval 2.000, min RX interval 2.000, multiplier 3
 Local discriminator 12032, remote discriminator 11096
 Echo mode disabled/inactive
 Remote is control-plane independent
  Session ID: 0x2a6a

3 sessions, 3 clients
Cumulative transmit rate 1.5 pps, cumulative receive rate 1.5 pps
```


© 2020 Juniper Networks, Inc. All rights reserved.

There are 3 up / active BFD sessions, one for each IS-IS enabled interface. From the extensive output of the command you can check that the client for the sessions is IS-IS L1, transmit and receive intervals are 2000ms. The multiplier is 3, which results in a protocol detection time of approximately of 6000ms.

Verify the advertisement of directly connected VRRP interfaces on R3 and R4.

```
lab@R1# run show isis database level 1 R3.00-00 extensive
IS-IS level 1 link-state database:

R3.00-00 Sequence: 0xe, Checksum: 0xe1a2, Lifetime: 1096 secs
   IS neighbor: R2.00                         Metric:       10
     Two-way fragment: R2.00-00, Two-way first fragment: R2.00-00
   IS neighbor: R4.00                         Metric:       10
     Two-way fragment: R4.00-00, Two-way first fragment: R4.00-00
   IS neighbor: R6.00                         Metric:       10
     Two-way fragment: R6.00-00, Two-way first fragment: R6.00-00
   IP prefix: 172.30.0.12/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.20/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.24/30                  Metric:       10 Internal Up
   IP prefix: 172.30.1.0/24                   Metric:       10 Internal Up
   IP prefix: 172.30.2.0/24                   Metric:       10 Internal Up
   IP prefix: 172.30.5.3/32                   Metric:        0 Internal Up
   V6 prefix: fd17:f0f4:f691:5::3/128         Metric:        0 Internal Up
...
```

```
[edit]
lab@R1# run show isis database level 1 R4.00-00 extensive
IS-IS level 1 link-state database:

R4.00-00 Sequence: 0x8, Checksum: 0xefae, Lifetime: 1098 secs
   IS neighbor: R1.00                         Metric:       10
     Two-way fragment: R1.00-00, Two-way first fragment: R1.00-00
   IS neighbor: R3.00                         Metric:       10
     Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00
   IS neighbor: R5.00                         Metric:       10
     Two-way fragment: R5.00-00, Two-way first fragment: R5.00-00
   IP prefix: 172.30.0.4/30                   Metric:       10 Internal Up
   IP prefix: 172.30.0.20/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.28/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.48/30                  Metric:       10 Internal Up
   IP prefix: 172.30.0.52/30                  Metric:       10 Internal Up
   IP prefix: 172.30.1.0/24                   Metric:       10 Internal Up
   IP prefix: 172.30.2.0/24                   Metric:       10 Internal Up
   IP prefix: 172.30.5.4/32                   Metric:        0 Internal Up
   V6 prefix: fd17:f0f4:f691:5::4/128         Metric:        0 Internal Up
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

R3 and R4 advertise the VRRP prefixes in their LSPs. The interfaces are listed as passive in the IS-IS interface table.

```
[edit]
lab@R3# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/4.123          1   0x1 Point to Point    Disabled               10/10
ge-0/0/4.134          1   0x1 Point to Point    Disabled               10/10
ge-0/0/4.136          1   0x1 Point to Point    Disabled               10/10
ge-0/0/4.200          1   0x1 Passive           Passive                10/10
ge-0/0/4.201          1   0x1 Passive           Passive                10/10
lo0.0                 1   0x1 Passive           Passive                 0/0
```

4) Configure RIP on R4 and R5.

R4 and R5 have a RIP connection to DC2.

a) Configure RIP on R4.

RIP configuration is simple. You need to add the neighbor under the group configuration. RIP uses multicast for advertisements and adjacencies are not formed, so adding the local interface is enough to initiate the RIP process.

```
[edit]
lab@R4# set protocols rip group dc2 neighbor ge-0/0/4.202

[edit]
lab@R4# commit
commit complete
```

Verify RIP neighbor.

```
[edit]
lab@R4# run show rip neighbor
                Local  Source          Destination     Send   Receive   In
Neighbor        State  Address         Address         Mode   Mode      Met
--------        -----  -------         -----------     ----   -------   ---
ge-0/0/4.202       Up 172.30.0.49     224.0.0.9       mcast  both        1
```

Junos accepts RIP routes by default but does not advertise anything without explicit policy.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show route advertising-protocol rip 172.30.0.49

[edit]
lab@R4#
```

```
[edit]
lab@R4# run show route receive-protocol rip 172.30.0.50

inet.0: 53 destinations, 53 routes (53 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.33.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.34.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.35.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.36.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.37.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.38.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.39.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.40.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.41.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.42.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.43.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.44.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.45.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.46.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.47.0/24     *[RIP/100] 00:03:29, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
```

b) R5

Configure RIP on R5 in the same way as on R4.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols rip group dc2 neighbor ge-0/0/4.204
[edit]
lab@R5# commit
commit complete
```

Verify RIP operation.

```
[edit]
lab@R5# run show rip neighbor
                Local  Source          Destination      Send   Receive   In
Neighbor        State  Address         Address          Mode   Mode     Met
--------        -----  -------         -----------      ----   -------  ---
ge-0/0/4.204       Up 172.30.0.57      224.0.0.9        mcast  both       1
```

```
[edit]
lab@R5# run show route receive-protocol rip 172.30.0.58

inet.0: 51 destinations, 51 routes (51 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.33.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.34.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.35.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.36.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.37.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.38.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.39.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.40.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.41.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.42.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.43.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.44.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
172.30.45.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.46.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.47.0/24     *[RIP/100] 00:01:48, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
```

R5 receives the same RIP routes as R4.

5) Configure the IS-IS to RIP redistribution policy at R4 and R5.

You only have to advertise a default route to the DC2 router using RIP. However, a default route is not present in the routing table. Routing policy only advertises active routes from the routing table. On most JNCIE exams, there is a limitation that you cannot use static routes if not explicitly allowed. In general, the creation of an aggregate route is not considered a static route and will not violate that limitation. If you are in doubt, please ask your exam proctor. The aggregate route will be redistributed to RIP on R4 and R5.

a) R4

First you must create the aggregate default route.

```
[edit]
lab@R4# set routing-options aggregate route 0.0.0.0/0

[edit]
lab@R4# commit
commit complete
```

Verify if the aggregate route exists in the routing table.

```
[edit]
lab@R4# run show route protocol aggregate

inet.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 03:22:03
                        Reject
```

Next you have to configure the policy exporting the aggregate default route to RIP. You must make sure that the default route is not received back from the DC2 router. Using a route tag to mark the route makes it easy to filter the route with import RIP policy.

```
[edit]
lab@R4# set policy-options policy-statement agg-to-rip term 1 from protocol aggregate
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set policy-options policy-statement agg-to-rip term 1 from route-filter
0.0.0.0/0 exact

[edit]
lab@R4# set policy-options policy-statement agg-to-rip term 1 then tag 123

[edit]
lab@R4# set policy-options policy-statement agg-to-rip term 1 then accept
```

The routing policy agg-to-rip matches the aggregate route 0.0.0.0/0, the tag 123 is added and the route is accepted for export.

Apply the policy to the RIP group in the export direction and commit.

```
[edit]
lab@R4# set protocols rip group dc2 export agg-to-rip

[edit]
lab@R4# commit
commit complete
```

Verify the export of the default route to DC2.

```
[edit]
lab@R4# run show route advertising-protocol rip 172.30.0.49

inet.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 03:22:03
                      Reject
```

b) R5

The requirement is to advertise a default route from both R4 and R5. This will ensure that a failure of one of these routers will not result in the default route disappearing from DC2 routing table.

You must configure the aggregate default route and the same export policy as on R4.

```
[edit]
lab@R5# set routing-options aggregate route 0.0.0.0/0

[edit]
lab@R5# set policy-options policy-statement agg-to-rip term 1 from protocol aggregate
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement agg-to-rip term 1 from route-filter
0.0.0.0/0 exact

[edit]
lab@R5# set policy-options policy-statement agg-to-rip term 1 then tag 123

[edit]
lab@R5# set policy-options policy-statement agg-to-rip term 1 then accept
```

Apply the policy to the RIP group in the export direction and commit.

```
[edit]
lab@R5# set protocols rip group dc2 export agg-to-rip

[edit]
lab@R5# commit
commit complete
```

Verify the result of the RIP export policy.

```
[edit]
lab@R5# run show route advertising-protocol rip 172.30.0.57

[edit]
lab@R5#
```

R5 is not advertising anything to DC2. Verify the routing table for the aggregate default route.

```
[edit]
lab@R5# run show route 0/0 exact | no-more

inet.0: 52 destinations, 53 routes (52 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[RIP/100] 00:05:56, metric 3, tag 123
                    > to 172.30.0.58 via ge-0/0/4.204
                   [Aggregate/130] 03:25:09
                       Reject
```

The aggregate default route is present in the routing table, but it is not active because of a less preferred route preference. The routing policy is not evaluating it for export to RIP. The reason for that is the received default route at R5 via RIP from DC2. This default route is the one advertised from R4. You can see the tag 123 attached to the route. You must filter the default route coming from RIP at R5 and R4.

© 2020 Juniper Networks, Inc. All rights reserved.

6) Configure RIP filtering policy.

To filter the default route received from RIP on R4 and R5 you can use import policy rejecting routes with tag 123.

a) R4

```
[edit]
lab@R4# set policy-options policy-statement filter-rip term 1 from protocol rip

[edit]
lab@R4# set policy-options policy-statement filter-rip term 1 from tag 123

[edit]
lab@R4# set policy-options policy-statement filter-rip term 1 then reject
```

Routing policy `filter-rip` matches the RIP routes with tag 123 and then rejects them. After applying the policy in the import direction, default route received from RIP is rejected.

```
[edit]
lab@R4# set protocols rip group dc2 import filter-rip

[edit]
lab@R4# commit
commit complete
```

b) R5

On R5 you must configure and apply the same policy as on R4. Here the result will be visible because R5 receives the default route readvertised from DC2 and it is active in the routing table. This makes the RIP export policy on R5 to not advertise a default route to DC2.

```
[edit]
lab@R5# set policy-options policy-statement filter-rip term 1 from protocol rip

[edit]
lab@R5# set policy-options policy-statement filter-rip term 1 from tag 123

[edit]
lab@R5# set policy-options policy-statement filter-rip term 1 then reject
```

Apply the policy and commit.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols rip group dc2 import filter-rip


[edit]
lab@R5# commit
commit complete
```

Verify if the default route from protocol RIP is filtered.

```
[edit]
lab@R5# run show route 0/0 exact | no-more

inet.0: 52 destinations, 52 routes (52 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 03:31:51
                       Reject
```

After applying the import policy to reject the default route received from RIP, only the aggregate default route is present in the routing table. Now that the aggregate route is active, the export policy matches it and its exported to DC2 via RIP.

```
[edit]
lab@R5# run show route advertising-protocol rip 172.30.0.57

inet.0: 52 destinations, 52 routes (52 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 03:31:41
                       Reject
```

Verify if on R4 that the default route from RIP is rejected, after when R5 is advertising successfully default route to DC2.

```
[edit]
lab@R4# run show route 0/0 exact

inet.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 03:35:47
                       Reject
```

Only the aggregate default route is present. The default route from RIP is filtered.

7) Configure the RIP to IS-IS redistribution policy on R4 and R5.

© 2020 Juniper Networks, Inc. All rights reserved.

On R4 and R5 you have to redistribute the other RIP routes into IS-IS.

a) R4

Use an IS-IS export policy to match all RIP routes.

```
[edit]
lab@R4# set policy-options policy-statement rip-to-isis term 1 from protocol rip

[edit]
lab@R4# set policy-options policy-statement rip-to-isis term 1 then accept

[edit]
lab@R4# set protocols isis export rip-to-isis

[edit]
lab@R4# commit
commit complete
```

b) R5

Implement the same export policy on R5.

```
[edit]
lab@R5# set policy-options policy-statement rip-to-isis term 1 from protocol rip

[edit]
lab@R5# set policy-options policy-statement rip-to-isis term 1 then accept

[edit]
lab@R5# set protocols isis export rip-to-isis

[edit]
lab@R5# commit
commit complete
```

Verify the redistribution of the RIP routes to IS-IS.

```
[edit]
lab@R4# run show isis database R4.00-00 extensive
IS-IS level 1 link-state database:

R4.00-00 Sequence: 0x1b, Checksum: 0x8843, Lifetime: 1017 secs
   IS neighbor: R1.00                          Metric:       10
     Two-way fragment: R1.00-00, Two-way first fragment: R1.00-00
   IS neighbor: R3.00                          Metric:       10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00
IS neighbor: R5.00                    Metric:       10
  Two-way fragment: R5.00-00, Two-way first fragment: R5.00-00
IP prefix: 172.30.0.4/30                Metric:       10 Internal Up
IP prefix: 172.30.0.20/30               Metric:       10 Internal Up
IP prefix: 172.30.0.28/30               Metric:       10 Internal Up
IP prefix: 172.30.0.48/30               Metric:       10 Internal Up
IP prefix: 172.30.0.52/30               Metric:       10 Internal Up
IP prefix: 172.30.1.0/24                Metric:       10 Internal Up
IP prefix: 172.30.2.0/24                Metric:       10 Internal Up
IP prefix: 172.30.5.4/32                Metric:        0 Internal Up
IP prefix: 172.30.32.0/24               Metric:        2 Internal Up
IP prefix: 172.30.33.0/24               Metric:        2 Internal Up
IP prefix: 172.30.34.0/24               Metric:        2 Internal Up
IP prefix: 172.30.35.0/24               Metric:        2 Internal Up
IP prefix: 172.30.36.0/24               Metric:        2 Internal Up
IP prefix: 172.30.37.0/24               Metric:        2 Internal Up
IP prefix: 172.30.38.0/24               Metric:        2 Internal Up
IP prefix: 172.30.39.0/24               Metric:        2 Internal Up
IP prefix: 172.30.40.0/24               Metric:        2 Internal Up
IP prefix: 172.30.41.0/24               Metric:        2 Internal Up
IP prefix: 172.30.42.0/24               Metric:        2 Internal Up
IP prefix: 172.30.43.0/24               Metric:        2 Internal Up
IP prefix: 172.30.44.0/24               Metric:        2 Internal Up
IP prefix: 172.30.45.0/24               Metric:        2 Internal Up
IP prefix: 172.30.46.0/24               Metric:        2 Internal Up
IP prefix: 172.30.47.0/24               Metric:        2 Internal Up
V6 prefix: fd17:f0f4:f691:5::4/128      Metric:        0 Internal Up
...
```

R4 is announcing the prefixes in the LSP advertised to other IS-IS routers. The routes are announced as Internal.

```
[edit]
lab@R5# run show isis database R5.00-00 extensive
IS-IS level 1 link-state database:

R5.00-00 Sequence: 0x1e, Checksum: 0x3a9a, Lifetime: 1144 secs
  IS neighbor: R4.00                    Metric:       10
    Two-way fragment: R4.00-00, Two-way first fragment: R4.00-00
  IS neighbor: R6.00                       Metric:        5
    Two-way fragment: R6.00-00, Two-way first fragment: R6.00-00
  IS neighbor: R8.00                      Metric:       10
    Two-way fragment: R8.00-00, Two-way first fragment: R8.00-00
  IP prefix: 172.30.0.28/30                  Metric:       10 Internal Up
  IP prefix: 172.30.0.32/30                  Metric:        5 Internal Up
  IP prefix: 172.30.0.36/30                  Metric:       10 Internal Up
  IP prefix: 172.30.0.56/30                  Metric:       10 Internal Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
   IP prefix: 172.30.0.60/30                Metric:       10 Internal Up
   IP prefix: 172.30.5.5/32                 Metric:        0 Internal Up
   V6 prefix: fd17:f0f4:f691:5::5/128       Metric:        0 Internal Up
...
```

The LSP announced by R5 does not include the RIP routes. Verify the routing table for these routes.

```
[edit]
lab@R5# run show route 172.30.32.0/24 exact

inet.0: 52 destinations, 68 routes (52 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[IS-IS/15] 00:06:20, metric 12
                    > to 172.30.0.29 via ge-0/0/4.145
                    [RIP/100] 00:31:38, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
```

The RIP routes are not active in routing table, so they are not eligible for export by the routing policy. The same routes are received from IS-IS as internal with a better route preference. The routes are advertised from R4 as Internal IS-IS routes because of the `wide-metrics-only` configuration.

8) Change the RIP preference on R4 and R5.

To fix the issue with non-advertising RIP routes on R5 you can change the preference of the RIP routes. You cannot reject the routes in IS-IS. IS-IS does not support import policies. Level 1 IS-IS routes have a preference of 15. So, changing the RIP preference to 14 on R4 and R5 will fix the issue.

R4

```
[edit]
lab@R4# set protocols rip group dc2 preference 14

[edit]
lab@R4# commit
commit complete
```

R5

```
[edit]
lab@R5# set protocols rip group dc2 preference 14
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# commit
commit complete
```

Verify the results of the change.

```
[edit]
lab@R5# run show route 172.30.32.0/24 exact

inet.0: 52 destinations, 52 routes (52 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[RIP/14] 00:34:37, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
```

R5's routing table now has the route only from RIP with a route preference of 14. The route from IS-IS is missing because R5 is also advertising the same RIP routes to the IS-IS domain and is rejecting the same routes received from R4.

```
[edit]
lab@R5# run show isis database R5.00-00 extensive
IS-IS level 1 link-state database:

R5.00-00 Sequence: 0x22, Checksum: 0xa3b0, Lifetime: 1134 secs
   IS neighbor: R4.00                    Metric:       10
     Two-way fragment: R4.00-00, Two-way first fragment: R4.00-00
   IS neighbor: R6.00                          Metric:        5
     Two-way fragment: R6.00-00, Two-way first fragment: R6.00-00
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
IS neighbor: R8.00                          Metric:       10
  Two-way fragment: R8.00-00, Two-way first fragment: R8.00-00
IP prefix: 172.30.0.28/30                        Metric:       10 Internal Up
IP prefix: 172.30.0.32/30                        Metric:        5 Internal Up
IP prefix: 172.30.0.36/30                        Metric:       10 Internal Up
IP prefix: 172.30.0.56/30                        Metric:       10 Internal Up
IP prefix: 172.30.0.60/30                        Metric:       10 Internal Up
IP prefix: 172.30.5.5/32                         Metric:        0 Internal Up
IP prefix: 172.30.32.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.33.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.34.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.35.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.36.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.37.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.38.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.39.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.40.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.41.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.42.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.43.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.44.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.45.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.46.0/24                        Metric:        2 Internal Up
IP prefix: 172.30.47.0/24                        Metric:        2 Internal Up
V6 prefix: fd17:f0f4:f691:5::5/128               Metric:        0 Internal Up
```

R5 is now advertising the RIP routes to IS-IS domain.

9) Ensuring IPv6 loopback reachability.

IS-IS routes IPv6 by default. The configuration is the also the same as for IPv4. The same operational commands are used as well.

© 2020 Juniper Networks, Inc. All rights reserved.

Check the IPv6 routing table on R1.

```
[edit]
lab@R1# run show route table inet6 | no-more

inet6.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691:5::1/128
                   *[Direct/0] 1d 04:39:01
                    > via lo0.0
fd17:f0f4:f691:5::2/128
                   *[IS-IS/15] 00:00:43, metric 5
                    > to fe80::2e6b:f5ff:fe1a:b8c6 via ae0.0
fd17:f0f4:f691:5::3/128
                   *[IS-IS/15] 00:00:43, metric 15
                    > to fe80::2e6b:f5ff:fe1a:b8c6 via ae0.0
fd17:f0f4:f691:5::5/128
                   *[IS-IS/15] 02:18:52, metric 20
                    > to fe80::250:5600:76a9:775f via ge-0/0/4.118
fd17:f0f4:f691:5::6/128
                   *[IS-IS/15] 00:00:43, metric 25
                      to fe80::2e6b:f5ff:fe1a:b8c6 via ae0.0
                    > to fe80::250:5600:76a9:775f via ge-0/0/4.118
fd17:f0f4:f691:5::7/128
                   *[IS-IS/15] 00:00:43, metric 15
                    > to fe80::2e6b:f5ff:fe1a:b8c6 via ae0.0
fd17:f0f4:f691:5::8/128
                   *[IS-IS/15] 02:18:52, metric 10
                    > to fe80::250:5600:76a9:775f via ge-0/0/4.118
fe80::250:5600:76a9:278c/128
                   *[Local/0] 1d 03:09:47
                      Local via ge-0/0/4.118
fe80::250:560f:fca9:4858/128
                   *[Direct/0] 1d 04:39:01
                    > via lo0.0
fe80::2e6b:f5ff:fe86:c6c0/128
                   *[Local/0] 00:11:53
                      Local via ae0.0
ff02::2/128        *[INET6/0] 3d 05:56:44
                      MultiRecv
```

R1 has routes for some of the IPv6 loopback interfaces but not for all. For example, the R4 loopback IPv6 address is not present.

© 2020 Juniper Networks, Inc. All rights reserved.

Even some of the routes that are present in routing table are not reachable.

```
[edit]
lab@R1# run ping fd17:f0f4:f691:5::5
PING6(56=40+8+8 bytes) fd17:f0f4:f691:5::1 --> fd17:f0f4:f691:5::5
64 bytes from fd17:f0f4:f691:5::8: Destination Host Unreachable
Vr TC  Flow Plen Nxt Hlim
 6 00 00000 0010  3a   40
fd17:f0f4:f691:5::1->fd17:f0f4:f691:5::5
ICMP6: type = 128, code = 0

64 bytes from fd17:f0f4:f691:5::8: Destination Host Unreachable
Vr TC  Flow Plen Nxt Hlim
 6 00 00000 0010  3a   40
fd17:f0f4:f691:5::1->fd17:f0f4:f691:5::5
ICMP6: type = 128, code = 0
```

Pinging R5's loopback IPv6 address from R1 results in an error received from R8 that the host is unreachable. This means that R8 has no route in the IPv6 routing table for R5 IPv6 loopback address.

Let's troubleshoot why R4's IPv6 loopback address is missing in the routing table. First check if the IPv6 prefix is present in the IS-IS LSP sent from R4.

```
[edit]
lab@R1# run show isis database R4.00-00 extensive | match V6
   V6 prefix: fd17:f0f4:f691:5::4/128         Metric:        0 Internal Up
    Protocols: IP, IPv6
    Speaks: IPV6
    IPv6 prefix: fd17:f0f4:f691:5::4/128 Metric 0 Up
```

The IPv6 route is present in the LSP sent from R4.

```
[edit]
lab@R1# run show route 172.30.5.4

inet.0: 32 destinations, 32 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.4/32      *[IS-IS/15] 00:05:25, metric 10
                    > to 172.30.0.6 via ge-0/0/4.114
```

The IPv4 loopback address of R4 is present in routing table and reachable. The route is reachable via interface ge-0/0/4.114.

Let's check the configuration of interface ge-0/0/4.114.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# show interface ge-0/0/4.114
description "R4 connection";
vlan-id 114;
family inet {
    address 172.30.0.5/30;
}
```

Now let's check the IPv4 and IPv6 loopback addresses in R5's routing table.

```
[edit]
lab@R1# run show route 172.30.5.5

inet.0: 32 destinations, 32 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.5/32      *[IS-IS/15] 00:05:25, metric 20
                    > to 172.30.0.6 via ge-0/0/4.118

[edit]
lab@R1# run show route table inet6 fd17:f0f4:f691:5::5

inet6.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691:5::5/128
                   *[IS-IS/15] 00:07:00, metric 20
                    > to fe80::fac0:100:76dc:2e84 via ge-0/0/4.118
```

R5 loopback IPv4 and IPv6 addresses are reachable via interface ge-0/0/4.118.

```
[edit]
lab@R1# show interface ge-0/0/4.118
description "R8 connection";
vlan-id 118;
family inet {
     address 172.30.0.9/30;
}
family inet6;
```

Interface ge-0/0/4.118 is configured with family inet6, while interface ge-0/0/4.114 is not.

By default, IS-IS calculates one topology for IPv4 and IPv6. In other words, one instance of SPF is calculated. However, because not every interface is enabled to use IPv6 in the network, some of the IPv6 routes are not reachable. The route towards the IPv6 loopback

© 2020 Juniper Networks, Inc. All rights reserved.

address of R4 is not resolved. The interface used for reaching IPv4 loopback address is not enabled for IPv6. The network has a different topology for IPv6 routing with less interfaces.

To resolve the issue, you have to enable IS-IS to calculate a different a topology for IPv6 on all routers. The following command ensures that a different SPF instance is used for the IPv6 topology.

```
[edit]
lab@R1# set protocols isis topologies ipv6-unicast

[edit]
lab@R1# commit
commit complete
```

Enable `topologies ipv6-unicast` on all routers in the network.

Verify the result after enabling the different topology for IPv6.

```
[edit]
lab@R1# run show route table inet6 | no-more

inet6.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691:5::1/128
                   *[Direct/0] 1d 04:48:38
                    > via lo0.0
fd17:f0f4:f691:5::2/128
                   *[IS-IS/15] 00:00:56, metric 5
                    > to fe80::2e6b:f5ff:fe1a:b8c6 via ae0.0
fd17:f0f4:f691:5::3/128
                   *[IS-IS/15] 00:00:50, metric 15
                    > to fe80::2e6b:f5ff:fe1a:b8c6 via ae0.0
fd17:f0f4:f691:5::4/128
                   *[IS-IS/15] 00:00:39, metric 25
                    > to fe80::2e6b:f5ff:fe1a:b8c6 via ae0.0
fd17:f0f4:f691:5::5/128
                   *[IS-IS/15] 00:00:20, metric 35
                      to fe80::2e6b:f5ff:fe1a:b8c6 via ae0.0
                    > to fe80::250:5600:76a9:775f via ge-0/0/4.118
fd17:f0f4:f691:5::6/128
                   *[IS-IS/15] 00:00:20, metric 30
                    > to fe80::250:5600:76a9:775f via ge-0/0/4.118
fd17:f0f4:f691:5::7/128
                   *[IS-IS/15] 00:00:20, metric 20
                    > to fe80::250:5600:76a9:775f via ge-0/0/4.118
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
fd17:f0f4:f691:5::8/128
                   *[IS-IS/15] 00:00:20, metric 10
                    > to fe80::250:5600:76a9:775f via ge-0/0/4.118
fe80::250:5600:76a9:278c/128
                   *[Local/0] 1d 03:19:24
                      Local via ge-0/0/4.118
fe80::250:560f:fca9:4858/128
                   *[Direct/0] 1d 04:48:38
                    > via lo0.0
fe80::2e6b:f5ff:fe86:c6c0/128
                   *[Local/0] 00:04:13
                      Local via ae0.0
ff02::2/128        *[INET6/0] 3d 06:06:21
                      MultiRecv
```

R1 now has an active route for all IPv6 loopback addresses in all 8 routers. All routes are resolved via IPv6 enabled interfaces.

The output below shows that IS-IS is using two instances of SPF.

```
[edit]
lab@R1# run show isis spf brief
IPV4 Unicast IS-IS level 1 SPF results:
Node             Metric     Interface       NH   Via            SNPA
R6.00            25         ae0.0           IPV4 R2             2c:6b:f5:1a:b8:c6
                            ge-0/0/4.114    IPV4 R4             0:50:56:a9:72:43
                            ge-0/0/4.118    IPV4 R8             0:50:56:a9:77:5f
R5.00            20         ge-0/0/4.114    IPV4 R4             0:50:56:a9:72:43
                            ge-0/0/4.118    IPV4 R8             0:50:56:a9:77:5f
R7.00            15         ae0.0           IPV4 R2             2c:6b:f5:1a:b8:c6
R3.00            15         ae0.0           IPV4 R2             2c:6b:f5:1a:b8:c6
R8.00            10         ge-0/0/4.118    IPV4 R8             0:50:56:a9:77:5f
R4.00            10         ge-0/0/4.114    IPV4 R4             0:50:56:a9:72:43
R2.00            5          ae0.0           IPV4 R2             2c:6b:f5:1a:b8:c6
R1.00            0
  8 nodes

IPV4 Unicast IS-IS level 2 SPF results:
  0 nodes

IPV6 Unicast IS-IS level 1 SPF results:
Node             Metric     Interface       NH   Via            SNPA
R5.00            35         ae0.0           IPV6 R2             2c:6b:f5:1a:b8:c6
                            ge-0/0/4.118    IPV6 R8             0:50:56:a9:77:5f
R6.00            30         ge-0/0/4.118    IPV6 R8             0:50:56:a9:77:5f
R4.00            25         ae0.0           IPV6 R2             2c:6b:f5:1a:b8:c6
R7.00            20         ge-0/0/4.118    IPV6 R8             0:50:56:a9:77:5f
R3.00            15         ae0.0           IPV6 R2             2c:6b:f5:1a:b8:c6
R8.00            10         ge-0/0/4.118    IPV6 R8             0:50:56:a9:77:5f
R2.00            5          ae0.0           IPV6 R2             2c:6b:f5:1a:b8:c6
R1.00            0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  8 nodes

IPV6 Unicast IS-IS level 2 SPF results:
  0 nodes
```

Note: Interface ge-0/0/4.114 is missing as next-hop in the IPv6 topology. SPF calculates the topology for IPv6 without considering interfaces that are not enabled for IPv6.

```
[edit]
lab@R1# run ping fd17:f0f4:f691:5::5
PING6(56=40+8+8 bytes) fd17:f0f4:f691:5::1 --> fd17:f0f4:f691:5::5
16 bytes from fd17:f0f4:f691:5::5, icmp_seq=0 hlim=61 time=3.034 ms
16 bytes from fd17:f0f4:f691:5::5, icmp_seq=1 hlim=61 time=2.532 ms
16 bytes from fd17:f0f4:f691:5::5, icmp_seq=2 hlim=61 time=7.306 ms
16 bytes from fd17:f0f4:f691:5::5, icmp_seq=3 hlim=61 time=8.927 ms
^C
--- fd17:f0f4:f691:5::5 ping6 statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 2.532/5.450/8.927/2.733 ms

[edit]
lab@R1# run traceroute fd17:f0f4:f691:5::5
traceroute6 to fd17:f0f4:f691:5::5 (fd17:f0f4:f691:5::5) from
fd17:f0f4:f691:5::1, 64 hops max, 12 byte packets
 1  fd17:f0f4:f691:5::8 (fd17:f0f4:f691:5::8)  17.906 ms  8.129 ms  5.532 ms
 2  fd17:f0f4:f691:5::7 (fd17:f0f4:f691:5::7)  4.461 ms  8.680 ms  8.032 ms
 3  fd17:f0f4:f691:5::6 (fd17:f0f4:f691:5::6)  8.580 ms  6.572 ms  5.662 ms
 4  fd17:f0f4:f691:5::5 (fd17:f0f4:f691:5::5)  3.275 ms  2.597 ms  7.898 ms
```

Ping and traceroute are successful with the new calculated topology for IPv6.

10) Configure OSPFv3 on R4 and R5.

Unlike IS-IS, for supporting IPv6 a new version of OSPF had to be developed. It was too difficult to extend OSPFv2 to support IPv6. OSPFv3 supports both IPv6 and IPv4, and in the same time it is an improved version of OSPFv2.

RouterID, AreaID and LSA IDs are kept in the same 32bits format. You can use the same RouterID for OSPFv2 and OSPFv3 if they are configured in the same time.

The neighbor adjacency process is also the same. With OSPFv3 link local addresses are used for communication on the link. The other difference is that several instances of OSPFv3 can form neighbors on single link.

The biggest improvement is with the LSAs added and the processing of these LSAs. Also, the names of some LSA are changed and for others the information carried is different.

© 2020 Juniper Networks, Inc. All rights reserved.

- **Router LSA (Type 0x2001)** – although the same name as with OSPFv2 the purpose of this LSA is different. In OSPFv3 it represents the router itself in the network and does not carry prefixes anymore.
- **Network LSA (Type 0x2002)** – this LSA has the same name and is used for the same purpose. Originated by the DR on the broadcast segment.
- **Inter-Area Prefix LSA (Type 0x2003)** – this type of LSA is in fact Summary Type 3 LSA used in OSPFv2. The name is changed to more accurately represent the function of the LSA. ABRs originate it into an area to advertise networks that are outside the area but internal to the OSPF domain.
- **Inter-Area Router LSA (Type 0x2004)** – this LSA is the same as Type 4 in OSPFv4. It is originated by ABR to advertise ASBR that resides outside the area.
- **AS-External LSA (Type 0x2005)** – External LSA from OSPFv2. It is used to advertise prefixes external to OSPF domain.
- **Type 7 LSA (Type 0x2007)** – this LSA is the same as NSSA Type 7 and used for same purpose.
- **Link LSA (Type 0x2008)** – this is totally new LSA for OSPFv3. Its scope is on a single link. It is not flooded to other routers in the area. It is used to advertise a link-local address to the other routers on the link and provides a list of all IPv6 prefixes on the link.
- **Intra-Area Prefix LSA (Type 0x2009)** – this new LSA carries the Intra Area prefixes instead of the Router LSA in OSPFv2.

11) Configure OSPFv3 on R4.

```
lab@R4# set protocols ospf3 area 0 interface ge-0/0/4.203

[edit]
lab@R4# commit
commit complete
```

Note: OSPFv3 is configured under the `protocols ospf3` hierarchy.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show ospf3 neighbor
ID                  Interface              State     Pri   Dead
172.30.63.1         ge-0/0/4.203           Full      128     38
  Neighbor-address fe80::200:800:cb00:1
```

The neighbor with DC3 is in the Full state.

Note: The `ospf3` keyword is used in operational commands to display OSPFv3 information. The neighbor-address is the IPv6 link-local address of the neighbor.

```
[edit]
lab@R4# run show route protocol ospf3

inet.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)

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

inet6.0: 32 destinations, 34 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691::3c/126
                   *[OSPF3/10] 00:02:02, metric 2
                    > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f::/80
                   *[OSPF3/150] 00:02:02, metric 0, tag 0
                    > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:1::/80
                   *[OSPF3/150] 00:02:02, metric 0, tag 0
                    > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:2::/80
                   *[OSPF3/150] 00:02:02, metric 0, tag 0
                    > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:3::/80
                   *[OSPF3/150] 00:02:02, metric 0, tag 0
                    > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:4::/80
                   *[OSPF3/150] 00:02:02, metric 0, tag 0
                    > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:5::/80
                   *[OSPF3/150] 00:02:02, metric 0, tag 0
                    > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:6::/80
                   *[OSPF3/150] 00:02:02, metric 0, tag 0
                    > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:7::/80
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    *[OSPF3/150] 00:02:02, metric 0, tag 0
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:8::/80
                    *[OSPF3/150] 00:02:02, metric 0, tag 0
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:9::/80
                    *[OSPF3/150] 00:02:02, metric 0, tag 0
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:a::/80
                    *[OSPF3/150] 00:02:02, metric 0, tag 0
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:b::/80
                    *[OSPF3/150] 00:02:02, metric 0, tag 0
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:c::/80
                    *[OSPF3/150] 00:02:02, metric 0, tag 0
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:d::/80
                    *[OSPF3/150] 00:02:02, metric 0, tag 0
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:e::/80
                    *[OSPF3/150] 00:02:02, metric 0, tag 0
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:f::/80
                    *[OSPF3/150] 00:02:02, metric 0, tag 0
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
fd17:f0f4:f691:f:f::1/128
                    *[OSPF3/10] 00:02:02, metric
                     > to fe80::200:800:cb00:1 via ge-0/0/4.203
ff02::5/128         *[OSPF3/10] 00:02:12, metric 1
                       MultiRecv
```

R4 is receiving IPv6 routes from DC3 via OSPFv3.

```
[edit]
lab@R4# run show ospf3 database

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.0          172.30.5.4        0x80000001  209  0x7aa3  40
Router      0.0.0.0          172.30.63.1       0x80000001  210  0xfae9  40
Network     0.0.0.1          172.30.63.1       0x80000001  210  0x5ebf  32
IntraArPfx  0.0.0.1          172.30.63.1       0x80000123  210  0x98b0  72
IntraArPfx  0.0.0.2          172.30.63.1       0x80000001  210  0x2688  52
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern      0.0.0.1          172.30.63.1       0x80000111 1594  0x50bc  40
Extern      0.0.0.2          172.30.63.1       0x80000111 1294  0x5ab0  40
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Extern      0.0.0.3          172.30.63.1      0x80000111  1144  0x64a4  40
Extern      0.0.0.4          172.30.63.1      0x80000111   994  0x6e98  40
Extern      0.0.0.5          172.30.63.1      0x80000111   844  0x788c  40
Extern      0.0.0.6          172.30.63.1      0x80000111   694  0x8280  40
Extern      0.0.0.7          172.30.63.1      0x80000111   544  0x8c74  40
Extern      0.0.0.8          172.30.63.1      0x80000111   394  0x9668  40
Extern      0.0.0.9          172.30.63.1      0x80000111   244  0xa05c  40
Extern      0.0.0.10         172.30.63.1      0x80000111    95  0xaa50  40
Extern      0.0.0.11         172.30.63.1      0x80000110  2944  0xb643  40
Extern      0.0.0.12         172.30.63.1      0x80000110  2794  0xc037  40
Extern      0.0.0.13         172.30.63.1      0x80000110  2644  0xca2b  40
Extern      0.0.0.14         172.30.63.1      0x80000110  2494  0xd41f  40
Extern      0.0.0.15         172.30.63.1      0x80000110  2344  0xde13  40
Extern      0.0.0.16         172.30.63.1      0x80000110  2194  0xe807  40

    OSPF3 Link-Local database, interface ge-0/0/4.203 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          172.30.5.4       0x80000002  207  0x3637  44
Link        0.0.0.1          172.30.63.1      0x80000111  1894 0xa995  64
```

The OSPFv3 database contains the new LSA types together with the External LSAs received from DC3. Only DC3 sends an Intra Area Prefix LSA. R4 is not advertising any Intra Area Prefixes, because the interface enabled for OSPFv3 has no IPv6 address different from the link-local address.

Note: Remember the LSAs marked with "*" are originated by the local router.

```
[edit]
lab@R4# run show ospf3 database intra-area-prefix lsa-id 0.0.0.1 extensive
    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
IntraArPfx  0.0.0.1          172.30.63.1      0x80000123   831  0x98b0  72
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 172.30.63.1
  Prefix-count 2
  Prefix fd17:f0f4:f691::3c/126
    Prefix-options 0x0, Metric 1
  Prefix fd17:f0f4:f691:f:f::1/128
    Prefix-options 0x2, Metric 0
  Aging timer 00:46:09
  Installed 00:13:50 ago, expires in 00:46:09
  Last changed 00:13:50 ago, Change count: 2
```

The intra Area Prefix LSA carries the IPv6 loopback address of DC3.

```
[edit]
lab@R4# run show ospf3 database external lsa-id 0.0.0.1 extensive
    OSPF3 AS SCOPE link state database
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern      0.0.0.1          172.30.63.1       0x80000111  2189  0x50bc  40
  Prefix fd17:f0f4:f691:f::/80
  Prefix-options 0x0, Metric 0, Type 2,
  Aging timer 00:23:31
  Installed 00:13:24 ago, expires in 00:23:31
  Last changed 00:13:24 ago, Change count: 1
```

External LSA advertised by DC3, carry one IPv6 prefix.

You have to make sure that OSPFv3 supports both IPv4 and IPv6. By default, OSPFv3 supports only IPv6 unicast routes but you can configure it to support multiple address families, including IPv4 unicast.

Configure OSPFv3 on R4 to support IPv4 unicast.

```
[edit]
lab@R4# set protocols ospf3 realm ipv4-unicast area 0 interface ge-0/0/4.203

[edit]
lab@R4# commit
commit complete
```

Adding `realm ipv4-unicast` enables a new instance of OSPFv3 on interface `ge-0/0/4.203`. Each realm maintains a separate set of neighbors and link state database.

```
[edit]
lab@R4# run show route protocol ospf3

inet.0: 71 destinations, 72 routes (71 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.0.60/30     *[OSPF3/10] 00:01:08, metric 2
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.63.1/32     *[OSPF3/10] 00:01:08, metric 1
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.128.0/24    *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.129.0/24    *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.130.0/24    *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.131.0/24    *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.132.0/24    *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.133.0/24    *[OSPF3/150] 00:01:08, metric 0, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.134.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.135.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.136.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.137.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.138.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.139.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.140.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.141.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.142.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
172.30.143.0/24     *[OSPF3/150] 00:01:08, metric 0, tag 0
                    > to 172.30.0.54 via ge-0/0/4.203
...
```

R4 is receiving IPv4 prefixes from DC3 after adding `realm ipv4-unicast`.

Display the link state database for `realm ipv4-unicast`. Specifying the `realm`, you can display information for different OSPFv3 realms.

```
[edit]
lab@R4# run show ospf3 database realm ipv4-unicast

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.0          172.30.5.4        0x80000001   38  0x79a4  40
Router      0.0.0.0          172.30.63.1       0x80000001   39  0xf9ea  40
Network     0.0.0.1          172.30.63.1       0x80000002   36  0x5bc1  32
IntraArPfx  0.0.0.1          172.30.63.1       0x80000123   39  0x650d  48
IntraArPfx  0.0.0.2          172.30.63.1       0x80000002   36  0xd002  40
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern      0.0.0.1          172.30.63.1       0x80000111  2300  0xd0c3  32
Extern      0.0.0.2          172.30.63.1       0x80000111  2000  0xd3be  32
Extern      0.0.0.3          172.30.63.1       0x80000111  1850  0xd6b9  32
Extern      0.0.0.4          172.30.63.1       0x80000111  1700  0xd9b4  32
Extern      0.0.0.5          172.30.63.1       0x80000111  1550  0xdcaf  32
Extern      0.0.0.6          172.30.63.1       0x80000111  1400  0xdfaa  32
Extern      0.0.0.7          172.30.63.1       0x80000111  1250  0xe2a5  32
Extern      0.0.0.8          172.30.63.1       0x80000111  1100  0xe5a0  32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Extern      0.0.0.9          172.30.63.1      0x80000111   950  0xe89b  32
Extern      0.0.0.10         172.30.63.1      0x80000111   800  0xeb96  32
Extern      0.0.0.11         172.30.63.1      0x80000111   650  0xee91  32
Extern      0.0.0.12         172.30.63.1      0x80000111   500  0xf18c  32
Extern      0.0.0.13         172.30.63.1      0x80000111   350  0xf487  32
Extern      0.0.0.14         172.30.63.1      0x80000111   200  0xf782  32
Extern      0.0.0.15         172.30.63.1      0x80000111    50  0xfa7d  32
Extern      0.0.0.16         172.30.63.1      0x80000110  2900  0xff77  32

    OSPF3 Link-Local database, interface ge-0/0/4.203 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          172.30.5.4       0x80000001   38  0xa1f8  52
Link        0.0.0.1          172.30.63.1      0x80000111  2600  0x252b  52
```

Display the external LSAs carrying the IPv4 prefix.

```
[edit]
lab@R4# run show ospf3 database realm ipv4-unicast external lsa-id 0.0.0.1
extensive
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern      0.0.0.1          172.30.63.1      0x80000004  1234  0xedb4  32
  Prefix 172.30.128.0/24
  Prefix-options 0x0, Metric 0, Type 2,
  Aging timer 00:39:26
  Installed 00:03:24 ago, expires in 00:39:26
  Last changed 00:03:24 ago, Change count: 1
```

12) Configure OSPFv3 on R5.

Configure OSPFv3 on R5 the same way as on R4.

```
[edit]
lab@R5# set protocols ospf3 area 0 interface ge-0/0/4.205

[edit]
lab@R5# set protocols ospf3 realm ipv4-unicast area 0 interface ge-0/0/4.205

[edit]
lab@R5# commit
commit complete
```

Verify the operation of OSPFv3 on R5.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# run show ospf3 neighbor
ID              Interface              State     Pri   Dead
172.30.63.1     ge-0/0/4.205           Full      128     30
  Neighbor-address fe80::200:800:cd00:1

[edit]
lab@R5# run show ospf3 database

    OSPF3 database, Area 0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.0          172.30.5.4      0x80000001  2289  0x7aa3  40
Router     *0.0.0.0          172.30.5.5      0x80000002    43  0x8496  40
Router      0.0.0.0          172.30.63.1     0x80000006    44  0xfbc1  56
Network     0.0.0.1          172.30.63.1     0x80000001  2289  0x5ebf  32
Network     0.0.0.2          172.30.63.1     0x80000001    44  0x62b9  32
IntraArPfx  0.0.0.1          172.30.63.1     0x80000126    44  0x712b  52
IntraArPfx  0.0.0.2          172.30.63.1     0x80000001  2289  0x2688  52
IntraArPfx  0.0.0.3          172.30.63.1     0x80000001    44  0x376d  52
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern      0.0.0.1          172.30.63.1     0x80000112   859  0x4ebd  40
Extern      0.0.0.2          172.30.63.1     0x80000112   740  0x58b1  40
Extern      0.0.0.3          172.30.63.1     0x80000112   620  0x62a5  40
Extern      0.0.0.4          172.30.63.1     0x80000112   501  0x6c99  40
Extern      0.0.0.5          172.30.63.1     0x80000112   381  0x768d  40
Extern      0.0.0.6          172.30.63.1     0x80000112   261  0x8081  40
Extern      0.0.0.7          172.30.63.1     0x80000112   142  0x8a75  40
Extern      0.0.0.8          172.30.63.1     0x80000112    23  0x9469  40
Extern      0.0.0.9          172.30.63.1     0x80000111  2323  0xa05c  40
Extern      0.0.0.10         172.30.63.1     0x80000111  2173  0xaa50  40
Extern      0.0.0.11         172.30.63.1     0x80000111  2023  0xb444  40
Extern      0.0.0.12         172.30.63.1     0x80000111  1877  0xbe38  40
Extern      0.0.0.13         172.30.63.1     0x80000111  1747  0xc82c  40
Extern      0.0.0.14         172.30.63.1     0x80000111  1617  0xd220  40
Extern      0.0.0.15         172.30.63.1     0x80000111  1488  0xdc14  40
Extern      0.0.0.16         172.30.63.1     0x80000111  1358  0xe608  40

    OSPF3 Link-Local database, interface ge-0/0/4.205 Area 0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          172.30.5.5      0x80000002    44  0x532   44
Link        0.0.0.2          172.30.63.1     0x80000112   979  0x35fd  64
```

R5 receives IPv4 and IPv6 prefixes from DC3.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# run show route protocol ospf3

inet.0: 69 destinations, 70 routes (69 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.0.52/30     *[OSPF3/10] 00:03:39, metric 2
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.63.1/32     *[OSPF3/10] 00:03:39, metric 1
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.128.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.129.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.130.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.131.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.132.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.133.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.134.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.135.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.136.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.137.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.138.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.139.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.140.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.141.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.142.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
172.30.143.0/24    *[OSPF3/150] 00:03:39, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205

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

inet6.0: 32 destinations, 34 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691::34/126
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    *[OSPF3/10] 00:38:31, metric 2
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:1::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:2::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:3::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:4::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:5::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:6::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:7::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:8::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:9::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:a::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:b::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:c::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:d::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:e::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:f::/80
                    *[OSPF3/150] 00:38:31, metric 0, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
fd17:f0f4:f691:f:f::1/128
                    *[OSPF3/10] 00:38:31, metric 1
                     > to fe80::200:800:cd00:1 via ge-0/0/4.205
ff02::5/128         *[OSPF3/10] 00:40:04, metric 1
                       MultiRecv
```

13) Configure the IS-IS to OSPFv3 redistribution policy on R4 and R5.

To achieve reachability with DC3 you have to configure mutual redistribution for IPv4 and IPv6 between IS-IS and OSPFv3. First you may start with IS-IS to OSPFv3 redistribution. Since the prefixes redistributed in OSPFv3 by R4 will be received at R5 you can mark the routes with a tag. This will help to restrict the routes and prevent suboptimal routing later.

a) R4

Configure the policy to export all IS-IS routes to OSPFv3.

```
[edit]
lab@R4# set policy-options policy-statement isis-to-ospf3 term 1 from protocol isis

[edit]
lab@R4# set policy-options policy-statement isis-to-ospf3 term 1 then tag 123

[edit]
lab@R4# set policy-options policy-statement isis-to-ospf3 term 1 then accept
```

Apply the policy for both IPv4 and IPv6 routes.

```
[edit]
lab@R4# set protocols ospf3 export isis-to-ospf3

[edit]
lab@R4# set protocols ospf3 realm ipv4-unicast export isis-to-ospf3

[edit]
lab@R4# commit
commit complete
```

Verify the redistribution of IPv6 and IPv4 prefixes to DC3.

```
[edit]
lab@R4# run show ospf3 database external advertising-router 172.30.5.4
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern     *0.0.0.1          172.30.5.4       0x80000001   54  0xd1be  48
Extern     *0.0.0.2          172.30.5.4       0x80000001   54  0xc3cf  48
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Extern     *0.0.0.3           172.30.5.4       0x80000001    54  0x9704  48
Extern     *0.0.0.4           172.30.5.4       0x80000001    54  0xc1d6  48
Extern     *0.0.0.5           172.30.5.4       0x80000001    54  0xbeb3  48
Extern     *0.0.0.6           172.30.5.4       0x80000001    54  0x92e7  48

[edit]
lab@R4# run show ospf3 database external  advertising-router 172.30.5.4 realm
ipv4-unicast
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern     *0.0.0.1           172.30.5.4       0x80000001   107  0x3196  36
Extern     *0.0.0.2           172.30.5.4       0x80000001   107  0xb504  36
Extern     *0.0.0.3           172.30.5.4       0x80000001   107  0xe3d0  36
Extern     *0.0.0.4           172.30.5.4       0x80000001   107  0x307a  36
Extern     *0.0.0.5           172.30.5.4       0x80000001   107  0x782e  36
Extern     *0.0.0.6           172.30.5.4       0x80000001   107  0x1782  36
Extern     *0.0.0.7           172.30.5.4       0x80000003    79  0x9be7  36
Extern     *0.0.0.8           172.30.5.4       0x80000001   107  0xafd5  36
Extern     *0.0.0.9           172.30.5.4       0x80000001   107  0x1270  36
Extern     *0.0.0.10          172.30.5.4       0x80000001   107  0x16a5  36
Extern     *0.0.0.11          172.30.5.4       0x80000001   107  0x387c  36
Extern     *0.0.0.12          172.30.5.4       0x80000001   107  0x1e99  36
Extern     *0.0.0.13          172.30.5.4       0x80000001   107  0x3084  36
Extern     *0.0.0.15          172.30.5.4       0x80000001   107  0x920f  36
Extern     *0.0.0.16          172.30.5.4       0x80000001   107  0x782c  36
```

R4 originates External LSAs for IPv6 and IPv4 prefixes after applying the policy in the export direction.

```
[edit]
lab@R5# run show route protocol ospf3 172.30.5.1/32

inet.0: 69 destinations, 87 routes (69 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       [OSPF3/150] 00:02:35, metric 10, tag 123
                    > to 172.30.0.62 via ge-0/0/4.205

[edit]
lab@R5# run show route protocol ospf3 fd17:f0f4:f691:5::1/128

inet6.0: 32 destinations, 41 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691:5::1/128
                    [OSPF3/150] 00:02:35, metric 25, tag 123
                    > to fe80::200:800:cd00:1 via ge-0/0/4.205
```


© 2020 Juniper Networks, Inc. All rights reserved.

R5 receives the re-advertised prefixes from DC3. Those prefixes are redistributed by R4 from IS-IS to OSPFv3. Both IPv4 and IPv6 prefixes are marked with the tag 123.

b) R5

Configure the same export policy on R5.

```
[edit]
lab@R5# set policy-options policy-statement isis-to-ospf3 term 1 from protocol isis

[edit]
lab@R5# set policy-options policy-statement isis-to-ospf3 term 1 then tag 123

[edit]
lab@R5# set policy-options policy-statement isis-to-ospf3 term 1 then accept

[edit]
lab@R5# set protocols ospf3 export isis-to-ospf3

[edit]
lab@R5# set protocols ospf3  realm ipv4-unicast export isis-to-ospf3

[edit]
lab@R5# commit
commit complete
```

Verify the redistribution of IPv6 and IPv4 prefixes to DC3.

```
[edit]
lab@R5# run show ospf3 database external advertising-router 172.30.5.5
    OSPF3 AS SCOPE link state database
 Type      ID               Adv Rtr           Seq         Age  Cksum  Len
Extern     *0.0.0.1             172.30.5.5       0x80000001    24  0x87d   48
Extern     *0.0.0.2             172.30.5.5       0x80000001    24  0xf98e  48
Extern     *0.0.0.3             172.30.5.5       0x80000001    24  0xcdc2  48
Extern     *0.0.0.4             172.30.5.5       0x80000001    24  0xa1f6  48
Extern     *0.0.0.5             172.30.5.5       0x80000001    24  0xadec  48
Extern     *0.0.0.6             172.30.5.5       0x80000001    24  0xf994  48
Extern     *0.0.0.7             172.30.5.5       0x80000001    24  0x463c  48
```

R5 originates External LSAs to DC3. The IPv6 loopback addresses are redistributed to OSPFv3.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# run show ospf3 database advertising-router 172.30.5.5 lsa-id 0.0.0.1
extensive
Extern     *0.0.0.1          172.30.5.5       0x80000001    44  0x87d   48
  Prefix fd17:f0f4:f691:5::1/128
  Prefix-options 0x0, Metric 35, Type 2,
   Tag 0.0.0.123
  Gen timer 00:49:16
  Aging timer 00:59:16
  Installed 00:00:44 ago, expires in 00:59:16, sent 00:00:44 ago
  Last changed 00:00:44 ago, Change count: 1, Ours
```

Note: There are 7 (not 8) IPv6 external LSAs, because the IPv6 loopback address of R5 is local and is active in the routing table from `protocol direct`, therefore not matched by the export policy.

```
[edit]
lab@R4# run show route 172.30.5.1/32

inet.0: 71 destinations, 91 routes (71 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32      *[IS-IS/15] 00:53:06, metric 10
                    > to 172.30.0.5 via ge-0/0/4.114
                    [OSPF3/150] 00:02:35, metric 20, tag 123
                    > to 172.30.0.54 via ge-0/0/4.203

[edit]
lab@R4# run show route table inet6 fd17:f0f4:f691:5::1/128

inet6.0: 32 destinations, 41 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691:5::1/128
                   *[IS-IS/15] 00:53:06, metric 25
                    > to fe80::2e21:7200:86cd:2684 via ge-0/0/4.134
                    [OSPF3/150] 00:02:39, metric 35, tag 123
                    > to fe80::200:800:cb00:1 via ge-0/0/4.203
```

R4 receives the IPv4 and IPv6 loopback IP addresses from DC3 with a tag of 123. The routing policy on R5 is advertising the IS-IS routes to DC3.

14) Configure the OSPFv3 to IS-IS redistribution policy on R4 and R5.

After IS-IS routes are redistributed to DC3, you now have to redistribute the routes in the other direction (from OSPFv3 to IS-IS).

© 2020 Juniper Networks, Inc. All rights reserved.

a) Configure OSPFv3 to IS-IS redistribution policy on R4.

Configure and apply the export policy matching all OSPFv3 routes.

```
[edit]
lab@R4# set policy-options policy-statement ospf3-to-isis term 1 from protocol ospf3

[edit]
lab@R4# set policy-options policy-statement ospf3-to-isis term 1 then accept

[edit]
lab@R4# set protocols isis export ospf3-to-isis

[edit]
lab@R4# commit
commit complete
```

Verify the redistribution of OSPFv3 IPv4 and IPv6 routes in the IS-IS domain.

```
[edit]
lab@R4# run show isis database level 1 R4.00-00 extensive
IS-IS level 1 link-state database:

R4.00-00 Sequence: 0x11, Checksum: 0x926f, Lifetime: 939 secs
   IPV4 Unicast IS neighbor: R1.00           Metric:       10
     Two-way fragment: R1.00-00, Two-way first fragment: R1.00-00
   IPV4 Unicast IS neighbor: R3.00       Metric:       10
     Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00
   IPV4 Unicast IS neighbor: R5.00    Metric:       10
     Two-way fragment: R5.00-00, Two-way first fragment: R5.00-00
   IPV6 Unicast IS neighbor: R3.00       Metric:       10
     Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00
   IPV6 Unicast IS neighbor: R5.00    Metric:       10
     Two-way fragment: R5.00-00, Two-way first fragment: R5.00-00
   IP IPV4 Unicast prefix: 172.30.0.4/30   Metric:      10 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.20/30  Metric:      10 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.28/30  Metric:      10 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.48/30  Metric:      10 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.52/30  Metric:      10 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.60/30  Metric:       2 Internal Up
   IP IPV4 Unicast prefix: 172.30.1.0/24   Metric:      10 Internal Up
   IP IPV4 Unicast prefix: 172.30.2.0/24   Metric:      10 Internal Up
   IP IPV4 Unicast prefix: 172.30.5.4/32   Metric:       0 Internal Up
   IP IPV4 Unicast prefix: 172.30.32.0/24  Metric:       2 Internal Up
   IP IPV4 Unicast prefix: 172.30.33.0/24  Metric:       2 Internal Up
   IP IPV4 Unicast prefix: 172.30.34.0/24  Metric:       2 Internal Up
   IP IPV4 Unicast prefix: 172.30.35.0/24  Metric:       2 Internal Up
   IP IPV4 Unicast prefix: 172.30.36.0/24  Metric:       2 Internal Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  IP IPV4 Unicast prefix: 172.30.37.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.38.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.39.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.40.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.41.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.42.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.43.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.44.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.45.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.46.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.47.0/24  Metric:       2 Internal Up
  IP IPV4 Unicast prefix: 172.30.63.1/32  Metric:       1 Internal Up
  IP IPV4 Unicast prefix: 172.30.128.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.129.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.130.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix:  172.30.131.0/24 Metric:      0 Internal Up
  IP IPV4 Unicast prefix: 172.30.132.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.133.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.134.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.135.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.136.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.137.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.138.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.139.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.140.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.141.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.142.0/24 Metric:       0 Internal Up
  IP IPV4 Unicast prefix: 172.30.143.0/24 Metric:       0 Internal Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691::3c/126 Metric:  2 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:5::4/128 Metric: 0 Internal Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f::/80 Metric:   0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:1::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:2::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:3::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:4::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:5::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:6::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:7::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:8::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:9::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:a::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:b::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:c::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:d::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:e::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:f::/80 Metric: 0 External Up
V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:f::1/128 Metric: 1 External
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: IPv4 prefixes are redistributed as *Internal* in IS-IS, because of the `wide-metrics-only` configuration. However, the IPv6 prefixes are redistributed as *External*. This is because the TLV 236 used for advertising IPv6 contains the Internal/External bit. External routes are redistributed with the External bit set. The `wide-metrics-only` configuration does not affect the route type for IPv6.

b) Configure OSPFv3 to IS-IS redistribution policy on R5.

Configure the same policy for exporting OSPFv3 routes to IS-IS on R5.

```
[edit]
lab@R5# set policy-options policy-statement ospf3-to-isis term 1 from protocol ospf3

[edit]
lab@R5# set policy-options policy-statement ospf3-to-isis term 1 then accept

[edit]
lab@R5# set protocols isis export ospf3-to-isis

[edit]
lab@R5# commit
commit complete
```

Verify the redistribution on R5 from OSPFv3 to IS-IS domain.

```
[edit]
lab@R5# run show isis database level 1 R5.00-00 extensive
IS-IS level 1 link-state database:

R5.00-00 Sequence: 0xc, Checksum: 0x94e1, Lifetime: 940 secs
   IPV4 Unicast IS neighbor: R4.00      Metric:       10
     Two-way fragment: R4.00-00, Two-way first fragment: R4.00-00
   IPV4 Unicast IS neighbor: R6.00          Metric:        5
     Two-way fragment: R6.00-00, Two-way first fragment: R6.00-00
   IPV6 Unicast IS neighbor: R4.00      Metric:       10
     Two-way fragment: R4.00-00, Two-way first fragment: R4.00-00
   IPV6 Unicast IS neighbor: R6.00          Metric:        5
     Two-way fragment: R6.00-00, Two-way first fragment: R6.00-00
   IP IPV4 Unicast prefix: 172.30.0.28/30   Metric:     10 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.32/30   Metric:      5 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.36/30   Metric:     10 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.52/30   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.56/30   Metric:     10 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.60/30   Metric:     10 Internal Up
   IP IPV4 Unicast prefix: 172.30.5.5/32    Metric:      0 Internal Up
   IP IPV4 Unicast prefix: 172.30.32.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.33.0/24   Metric:      2 Internal Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
   IP IPV4 Unicast prefix: 172.30.34.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.35.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.36.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.37.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.38.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.39.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.40.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.41.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.42.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.43.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.44.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.45.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.46.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.47.0/24   Metric:      2 Internal Up
   IP IPV4 Unicast prefix: 172.30.63.1/32   Metric:      1 Internal Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691::34/126  Metric: 2 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:5::5/128 Metric: 0 Internal Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f::/80   Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:1::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:2::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:3::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:4::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:5::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:6::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:7::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:8::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:9::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:a::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:b::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:c::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:d::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:e::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:f::/80 Metric: 0 External Up
 V6 IPV6 Unicast prefix: fd17:f0f4:f691:f:f::1/128 Metric: 1 External Up
...
```

Only IPv6 routes are advertised by R5 to the IS-IS domain. Verify OSPFv3 IPv4 and IPv6 routes in the routing table of R5.

```
[edit]
lab@R5# run show route 172.30.132.0/24

inet.0: 69 destinations, 87 routes (69 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.132.0/24    *[IS-IS/15] 00:18:36, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
                    [OSPF3/150] 00:32:23, metric 0, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.62 via ge-0/0/4.205

[edit]
lab@R5# run show route table inet6 fd17:f0f4:f691:f:f::1/128

inet6.0: 33 destinations, 52 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691:f:1::/80
                   *[OSPF3/150] 00:44:19, metric 0, tag 0
                   > to fe80::200:800:cd00:1 via ge-0/0/4.205
                   [IS-IS/160] 00:44:41, metric 11
                   > to fe80::fac0:100:91dc:3184 via ge-0/0/4.145
```

The IPv4 prefixes received from DC3 via OSPFv3 are not active in routing table. The active routes are known in IS-IS from R4. IS-IS is configured for `wide-metrics-only`, which represents external IPv4 routes as internal in the IS-IS domain. Remember that internal IS-IS routes have a lower route preference in the routing table. The export routing policy does not evaluate the non-active OSPFv3 routes. For IPv6 routes this is not true, as they are represented as External and their route preference is worse than for IS-IS routes with an internal route preference.

15) Set OSPFv3 external preference at R4 and R5.

Fix the issue in the previous step that requires that the OSPFv3 routes must be active in routing table of R4 and R5. Lowering the preference of the external OSPFv3 routes will resolve the issue. You already faced the same problem with the RIP routes from DC2 where you lower the RIP preference to 14 (one less than default internal Level 1 preference of 15). Now you have to the same with OSPFv3 routes. Because the same two routes are involved it is better to use value different of 14 otherwise both RIP and OSPFv3 routes will have the same preference. You can use 13 for OSPFv3. It will be lower than 15 and will do the job.

Lowering the external OSPFv3 preference of IPv4 and IPv6 routes to 13 on R4 and R5.

R4

```
[edit]
lab@R4# set protocols ospf3 external-preference 13

[edit]
lab@R4# set protocols ospf3 realm ipv4-unicast external-preference 13

[edit]
lab@R4# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

R5

```
[edit]
lab@R5# set protocols ospf3 external-preference 13

[edit]
lab@R5# set protocols ospf3 realm ipv4-unicast external-preference 13

[edit]
lab@R5# commit
commit complete
```

Note: The external preference is changed for IPv6 as well. It is not needed for fixing the redistribution from R4 and R5 to IS-IS domain. It is done for consistency so that IPv4 and IPv6 OSPFv3 routes have the same preference.

Verify the result of lowering the route preference of OSPFv3 routes.

```
[edit]
lab@R5# run show route 172.30.132.0/24

inet.0: 69 destinations, 87 routes (69 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.132.0/24    *[OSPF3/13] 00:00:49, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:20:35, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
```

Now the routes received from OSPFv3 are active and eligible for export by the IS-IS policy.

```
[edit]
lab@R5# run show isis database level 1 R5.00-00 extensive | match 172.30.132.0/24

   IP IPV4 Unicast prefix: 172.30.132.0/24    Metric:         0 Internal Up
    IP extended prefix: 172.30.132.0/24 metric 0 up
```

16) Configure RIP to OSPFv3 redistribution policy at R4 and R5.

You need to mutual redistribute between RIP and OSPFv3 routes.

a) R4

© 2020 Juniper Networks, Inc. All rights reserved.

Creating policy matching all RIP routes on R4 and apply it in the export direction to OSPFv3.

```
[edit]
lab@R4# set policy-options policy-statement rip-to-ospf3 term 1 from protocol rip

[edit]
lab@R4# set policy-options policy-statement rip-to-ospf3 term 1 then tag 123

[edit]
lab@R4# set policy-options policy-statement rip-to-ospf3 term 1 then accept
```

Apply the policy only to `realm ipv4-unicast`, because RIP only supports IPv4. For IPv6 a different protocol must be used (RIPng). Again, the routes are marked with tag 123 for easy filtering if needed later.

```
[edit]
lab@R4# set protocols ospf3 realm ipv4-unicast export rip-to-ospf3

[edit]
lab@R4# commit
commit complete
```

Verify the result of the export policy from RIP to OSPFv3. Now R4 is advertising RIP routes together with IS-IS IPv4 routes.

```
[edit]
lab@R4# run show ospf3 database external  advertising-router 172.30.5.4 realm
ipv4-unicast
    OSPF3 AS SCOPE link state database
 Type      ID               Adv Rtr           Seq         Age  Cksum  Len
Extern    *0.0.0.1         172.30.5.4       0x80000002   523  0x5861  36
Extern    *0.0.0.2         172.30.5.4       0x80000002   438  0x5b5c  36
Extern    *0.0.0.3         172.30.5.4       0x80000002   352  0x5e57  36
Extern    *0.0.0.4         172.30.5.4       0x80000002   266  0x6152  36
Extern    *0.0.0.5         172.30.5.4       0x80000002   181  0x644d  36
Extern    *0.0.0.6         172.30.5.4       0x80000002    95  0x6748  36
Extern    *0.0.0.7         172.30.5.4       0x80000002     9  0x6a43  36
Extern    *0.0.0.8         172.30.5.4       0x80000001   839  0x6f3d  36
Extern    *0.0.0.9         172.30.5.4       0x80000001   839  0x7238  36
Extern    *0.0.0.10        172.30.5.4       0x80000001   839  0x7533  36
Extern    *0.0.0.11        172.30.5.4       0x80000001   839  0x782e  36
Extern    *0.0.0.12        172.30.5.4       0x80000001   839  0x7b29  36
Extern    *0.0.0.13        172.30.5.4       0x80000001   839  0x7e24  36
Extern    *0.0.0.14        172.30.5.4       0x80000001   839  0x811f  36
Extern    *0.0.0.15        172.30.5.4       0x80000001   839  0x841a  36
Extern    *0.0.0.16        172.30.5.4       0x80000001   839  0x8715  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Extern    *0.0.0.17       172.30.5.4      0x80000001   817  0xa8e5  36
Extern    *0.0.0.18       172.30.5.4      0x80000001   817  0xb7c1  36
Extern    *0.0.0.20       172.30.5.4      0x80000001   817  0xe9c3  36
Extern    *0.0.0.21       172.30.5.4      0x80000001   793  0x684b  36
Extern    *0.0.0.22       172.30.5.4      0x80000001   793  0xecb8  36
Extern    *0.0.0.23       172.30.5.4      0x80000001   793  0x931b  36
Extern    *0.0.0.24       172.30.5.4      0x80000002   696  0xf8f   36
Extern    *0.0.0.25       172.30.5.4      0x80000001   783  0x5d38  36
Extern    *0.0.0.26       172.30.5.4      0x80000001   783  0xa104  36
Extern    *0.0.0.27       172.30.5.4      0x80000001   737  0xed9f  36
Extern    *0.0.0.28       172.30.5.4      0x80000002   696  0x8ffe  36
Extern    *0.0.0.29       172.30.5.4      0x80000001   722  0x86f2  36
Extern    *0.0.0.30       172.30.5.4      0x80000001   722  0xb1eb  36
Extern    *0.0.0.49       172.30.5.4      0x80000001   697  0x3225  36
Extern    *0.0.0.50       172.30.5.4      0x80000001   697  0x334b  36
Extern    *0.0.0.51       172.30.5.4      0x80000001   696  0x96f9  3
```

b) R5

The same policy has to applied to R5 as well for redundancy requirement.

```
[edit]
lab@R5# set policy-options policy-statement rip-to-ospf3 term 1 from protocol rip

[edit]
lab@R5# set policy-options policy-statement rip-to-ospf3 term 1 then tag 123

[edit]
lab@R5# set policy-options policy-statement rip-to-ospf3 term 1 then accept

[edit]
lab@R5# set protocols ospf3 realm ipv4-unicast export rip-to-ospf3

[edit]
lab@R5# commit
commit complete
```

Verify the advertisement of RIP routes to OSPFv3.

```
[edit]
lab@R5# run show ospf3 database realm ipv4-unicast advertising-router 172.30.5.5
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern    *0.0.0.1        172.30.5.5      0x80000002   280  0x5266  36
Extern    *0.0.0.2        172.30.5.5      0x80000002   199  0x5561  36
Extern    *0.0.0.3        172.30.5.5      0x80000002   118  0x585c  36
Extern    *0.0.0.4        172.30.5.5      0x80000002    36  0x5b57  36
Extern    *0.0.0.5        172.30.5.5      0x80000001   670  0x6051  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Extern     *0.0.0.6        172.30.5.5        0x80000001   670  0x634c   36
Extern     *0.0.0.7        172.30.5.5        0x80000001   670  0x6647   36
Extern     *0.0.0.8        172.30.5.5        0x80000001   670  0x6942   36
Extern     *0.0.0.9        172.30.5.5        0x80000001   670  0x6c3d   36
Extern     *0.0.0.10       172.30.5.5        0x80000001   670  0x6f38   36
Extern     *0.0.0.11       172.30.5.5        0x80000001   670  0x7233   36
Extern     *0.0.0.12       172.30.5.5        0x80000001   670  0x752e   36
Extern     *0.0.0.13       172.30.5.5        0x80000001   670  0x7829   36
Extern     *0.0.0.14       172.30.5.5        0x80000001   670  0x7b24   36
Extern     *0.0.0.15       172.30.5.5        0x80000001   670  0x7e1f   36
Extern     *0.0.0.16       172.30.5.5        0x80000001   670  0x811a   36
Extern     *0.0.0.17       172.30.5.5        0x80000001   656  0xe0cc   36
```

R5 is not advertising any RIP routes to DC3 via OSPFv3. Only IS-IS IPv4 routes are redistributed to DC3.

17) Configure a filtering policy on R4 and R5.

The reason why R5 is not redistributing RIP routes to OSPFv3 can be found in the routing table. Select a RIP route and displaying it will reveal the issue.

```
[edit]
lab@R5# run show route 172.30.32.0/24

inet.0: 69 destinations, 97 routes (69 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[OSPF3/13] 00:00:13, metric 2, tag 123
                    > to 172.30.0.62 via ge-0/0/4.205
                   [RIP/14] 02:11:28, metric 2, tag 0
                   > to 172.30.0.58 via ge-0/0/4.204
                   [IS-IS/15] 00:00:13, metric 12
                   > to 172.30.0.29 via ge-0/0/4.145
```

The RIP route is not active. In fact, by lowering the external preference of OSPFv3 routes, R5 is now using DC3 as gateway for all RIP routes. This is true also for all IS-IS routes in the network, because OSPFv3 external prefixes have the lowest preference compared to the IS-IS routes.

You already tagged all IPv4 and IPv6 routes redistributed to DC3 with tag 123. Creating a policy matching that tag and apply it in the import direction will filter all routes received back into the IS-IS domain from DC3 and will resolve the suboptimal routing.

a) Applying the filtering policy on R4.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set policy-options policy-statement ospf3-filter term 1 from protocol ospf3

[edit]
lab@R4# set policy-options policy-statement ospf3-filter term 1 from tag 123

[edit]
lab@R4# set policy-options policy-statement ospf3-filter term 1 then reject
```

Apply the policy in import direction for both IPv4 and IPv6 families and commit the configuration.

```
[edit]
lab@R4# set protocols ospf3 import ospf3-filter

[edit]
lab@R4# set protocols ospf3 realm ipv4-unicast import ospf3-filter

[edit]
lab@R4# commit
commit complete
```

Verify again the routing table of R5.

```
[edit]
lab@R5# run show route 172.30.32.0/24

inet.0: 69 destinations, 87 routes (69 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[RIP/14] 02:13:44, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
```

Verify if R5 is advertising RIP routes to OSPFv3.

```
[edit]
lab@R5# run show ospf3 database realm ipv4-unicast advertising-router 172.30.5.5
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern     *0.0.0.1          172.30.5.5      0x80000002   280  0x5266  36
Extern     *0.0.0.2          172.30.5.5      0x80000002   199  0x5561  36
Extern     *0.0.0.3          172.30.5.5      0x80000002   118  0x585c  36
Extern     *0.0.0.4          172.30.5.5      0x80000002    36  0x5b57  36
Extern     *0.0.0.5          172.30.5.5      0x80000001   670  0x6051  36
Extern     *0.0.0.6          172.30.5.5      0x80000001   670  0x634c  36
Extern     *0.0.0.7          172.30.5.5      0x80000001   670  0x6647  36
Extern     *0.0.0.8          172.30.5.5      0x80000001   670  0x6942  36
Extern     *0.0.0.9          172.30.5.5      0x80000001   670  0x6c3d  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Extern    *0.0.0.10       172.30.5.5      0x80000001   670  0x6f38  36
Extern    *0.0.0.11       172.30.5.5      0x80000001   670  0x7233  36
Extern    *0.0.0.12       172.30.5.5      0x80000001   670  0x752e  36
Extern    *0.0.0.13       172.30.5.5      0x80000001   670  0x7829  36
Extern    *0.0.0.14       172.30.5.5      0x80000001   670  0x7b24  36
Extern    *0.0.0.15       172.30.5.5      0x80000001   670  0x7e1f  36
Extern    *0.0.0.16       172.30.5.5      0x80000001   670  0x811a  36
Extern    *0.0.0.17       172.30.5.5      0x80000001   656  0xe0cc  36
Extern    *0.0.0.18       172.30.5.5      0x80000001   656  0xb7e4  36
Extern    *0.0.0.19       172.30.5.5      0x80000001   656  0x3748  36
Extern    *0.0.0.20       172.30.5.5      0x80000001   656  0x6d46  36
Extern    *0.0.0.21       172.30.5.5      0x80000001   656  0x7041  36
Extern    *0.0.0.22       172.30.5.5      0x80000001   656  0xc1e9  36
Extern    *0.0.0.23       172.30.5.5      0x80000001   632  0x8a1c  36
Extern    *0.0.0.24       172.30.5.5      0x80000001   632  0xf89   36
Extern    *0.0.0.25       172.30.5.5      0x80000001   632  0xb5eb  36
Extern    *0.0.0.26       172.30.5.5      0x80000003   520  0x1184  36
Extern    *0.0.0.27       172.30.5.5      0x80000002   513  0x4150  36
Extern    *0.0.0.28       172.30.5.5      0x80000001   622  0xc3d4  36
Extern    *0.0.0.29       172.30.5.5      0x80000001   561  0x632f  36
Extern    *0.0.0.30       172.30.5.5      0x80000001   561  0x3a47  36
Extern    *0.0.0.31       172.30.5.5      0x80000001   561  0x6540  36
Extern    *0.0.0.32       172.30.5.5      0x80000002   513  0x98d7  36
Extern    *0.0.0.33       172.30.5.5      0x80000002   513  0x99fd  36
Extern    *0.0.0.34       172.30.5.5      0x80000002   520  0x5743  36
```

Look for routes with tag 123 in the routing table. The output should indicate that no routes with tag 123 are found.

```
[edit]
lab@R4# run show route | match "tag 123"
```

b) Apply the filter policy on R5

Configure the same policy on R5.

```
[edit]
lab@R5# set policy-options policy-statement ospf3-filter term 1 from protocol ospf3

[edit]
lab@R5# set policy-options policy-statement ospf3-filter term 1 from tag 123

[edit]
lab@R5# set policy-options policy-statement ospf3-filter term 1 then reject

[edit]
lab@R5# set protocols ospf3 import ospf3-filter
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols ospf3 realm ipv4-unicast import ospf3-filter

[edit]
lab@R5# commit
commit complete
```

Verify the result of the import policy on R5.

```
[edit]
lab@R5# run show route | match "tag 123"
```

No routes with tag 123 are present in the routing table.

18) Configure OSPFv3 to RIP redistribution policy on R4 and R5.

Last step for reaching full reachability is to redistribute the IPv4 OSPFv3 prefixes to DC2 via RIP.

```
[edit]
lab@R4# set policy-options policy-statement ospf3-to-rip term 1 from protocol ospf3

[edit]
lab@R4# set policy-options policy-statement ospf3-to-rip term 1 then tag 123

[edit]
lab@R4# set policy-options policy-statement ospf3-to-rip term 1 then accept
```

The policy is matching all OSPFv3 routes. Tag 123 is added for filtering the routes in the import direction to the IS-IS domain. Such import policy is already in place implemented in step 7 where the aggregate default route had to be filtered.

```
[edit]
lab@R4# set protocols rip group dc2 export ospf3-to-rip

[edit]
lab@R4# commit
commit complete
```

Policy is applied in export direction to the RIP group on R4.

Verify the result of the applied policy.

```
[edit]
lab@R4# run show route advertising-protocol rip 172.30.0.49
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 71 destinations, 71 routes (71 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[Aggregate/130] 02:18:03
                       Reject
172.30.0.60/30      *[OSPF3/10] 00:53:42, metric 2
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.63.1/32      *[OSPF3/10] 00:53:42, metric 1
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.128.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.129.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.130.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.131.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.132.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.133.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.134.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.135.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.136.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.137.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.138.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.139.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.140.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.141.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.142.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
172.30.143.0/24     *[OSPF3/13] 00:18:39, metric 0, tag 0
                     > to 172.30.0.54 via ge-0/0/4.203
```

Together with the aggregate default route, all OSPFv3 routes are advertised to RIP.

Implement the same policy on R5 to achieve the required redundancy in the network.

```
[edit]
lab@R5# set policy-options policy-statement ospf3-to-rip term 1 from protocol ospf3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement ospf3-to-rip term 1 then tag 123

[edit]
lab@R5# set policy-options policy-statement ospf3-to-rip term 1 then accept

[edit]
lab@R5# set protocols rip group dc2 export ospf3-to-rip

[edit]
lab@R5# commit
commit complete
```

Verify the redistribution of the OSPFv3 routes to RIP on R5.

```
lab@R5# run show route advertising-protocol rip 172.30.0.57

inet.0: 69 destinations, 87 routes (69 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[Aggregate/130] 02:20:13
                       Reject
172.30.0.52/30      *[OSPF3/10] 00:55:11, metric 2
                     > to 172.30.0.62 via ge-0/0/4.205
                     [IS-IS/15] 01:38:19, metric 20
                     > to 172.30.0.29 via ge-0/0/4.145
172.30.63.1/32      *[OSPF3/10] 00:55:11, metric 1
                     > to 172.30.0.62 via ge-0/0/4.205
                     [IS-IS/15] 00:41:24, metric 11
                     > to 172.30.0.29 via ge-0/0/4.145
172.30.128.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                     > to 172.30.0.62 via ge-0/0/4.205
                     [IS-IS/15] 00:41:24, metric 10
                     > to 172.30.0.29 via ge-0/0/4.145
172.30.129.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                     > to 172.30.0.62 via ge-0/0/4.205
                     [IS-IS/15] 00:41:24, metric 10
                     > to 172.30.0.29 via ge-0/0/4.145
172.30.130.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                     > to 172.30.0.62 via ge-0/0/4.205
                     [IS-IS/15] 00:41:24, metric 10
                     > to 172.30.0.29 via ge-0/0/4.145
172.30.131.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                     > to 172.30.0.62 via ge-0/0/4.205
                     [IS-IS/15] 00:41:24, metric 10
                     > to 172.30.0.29 via ge-0/0/4.145
172.30.132.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.133.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.134.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.135.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.136.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.137.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.138.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.139.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.140.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.141.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.142.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.143.0/24     *[OSPF3/13] 00:21:38, metric 0, tag 0
                    > to 172.30.0.62 via ge-0/0/4.205
                    [IS-IS/15] 00:41:24, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
```


© 2020 Juniper Networks, Inc. All rights reserved.

R5 is successfully advertising the OSPFv3 routes to DC3 via RIP.

19) Verify the redundancy of the network.

To verify the redundancy of the network you can initiate a traceroute from R1 to IPv4 and IPv6 OSPFv3 routes, and IPv4 RIP routes.

First verify the traceroute from R1 when the network is converged with both R4 and R5 in operation.

Tracing RIP route 172.30.32.1.

```
[edit]
lab@R1# run traceroute 172.30.32.1 no-resolve
traceroute to 172.30.32.1 (172.30.32.1), 30 hops max, 40 byte packets
 1  172.30.0.6  7.983 ms  8.165 ms  7.976 ms
 2  172.30.0.50  18.481 ms !N  9.088 ms !N  34.812 ms !N
```

The path is from R1 and ends at DC2.

Tracing the OSPFv3 route 172.30.128.1.

```
[edit]
lab@R1# run traceroute 172.30.128.1 no-resolve
traceroute to 172.30.128.1 (172.30.128.1), 30 hops max, 40 byte packets
 1  172.30.0.6  7.907 ms  7.477 ms  6.497 ms
 2  172.30.0.54  5.517 ms !N  5.116 ms !N  9.935 ms !N
```

The path is from R1 to R4 and ends at DC3.

Finally trace the IPv6 OSPFv3 route fd17:f0f4:f691:f::1.

```
[edit]
lab@R1# run traceroute fd17:f0f4:f691:f::1 no-resolve
traceroute6 to fd17:f0f4:f691:f::1 (fd17:f0f4:f691:f::1) from
fd17:f0f4:f691:5::1, 64 hops max, 12 byte packets
 1  fd17:f0f4:f691:5::2  24.914 ms  8.498 ms  8.060 ms
 2  fd17:f0f4:f691:5::3  8.605 ms  4.222 ms  3.804 ms
 3  fd17:f0f4:f691:5::4  8.844 ms  8.637 ms  7.991 ms
 4  fd17:f0f4:f691::36  4.991 ms !A  3.626 ms !A  3.398 ms !A
```

Because the IPv6 IS-IS topology is different from the IPV4 topology, the path traces from R1, R2, R3, R4 and ends at DC3.

To simulate the failure of R4 you can deactivate the whole `protocols` hierarchy.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# deactivate protocols

[edit]
lab@R4# commit
commit complete
```

Now when R4 is out of the routing path, you can verify again the path to the same IP addresses. You should wait a couple minutes for the network to synchronize.

```
[edit]
lab@R1# run traceroute 172.30.32.1 no-resolve
traceroute to 172.30.32.1 (172.30.32.1), 30 hops max, 40 byte packets
 1  172.30.0.10  7.994 ms  2.116 ms  7.869 ms
 2  172.30.0.37  9.173 ms  8.160 ms  7.867 ms
 3  172.30.0.58  4.751 ms !N  11.560 ms !N  4.987 ms !N
```

The path to RIP route is now from R1, R8, R5 and ends at DC2.

```
[edit]
lab@R1# run traceroute 172.30.128.1 no-resolve
traceroute to 172.30.128.1 (172.30.128.1), 30 hops max, 40 byte packets
 1  172.30.0.10  8.292 ms  7.641 ms  8.207 ms
 2  172.30.0.37  3.276 ms  8.324 ms  8.118 ms
 3  172.30.0.62  4.803 ms !N  4.972 ms !N  4.713 ms !N
```

The path to the IPV4 OSPFv3 route is starting from E1, R8, R4 and ends at DC3.

OSPFv3 IPv6 route to DC3.

```
[edit]
lab@R1# run traceroute fd17:f0f4:f691:f::1 no-resolve
traceroute6 to fd17:f0f4:f691:f::1 (fd17:f0f4:f691:f::1) from
fd17:f0f4:f691:5::1, 64 hops max, 12 byte packets
 1  fd17:f0f4:f691:5::8  14.133 ms  7.891 ms  3.738 ms
 2  fd17:f0f4:f691:5::7  7.890 ms  7.677 ms  8.054 ms
 3  fd17:f0f4:f691:5::6  8.364 ms  4.933 ms  5.762 ms
 4  fd17:f0f4:f691:5::5  8.935 ms  9.121 ms  8.219 ms
 5  fd17:f0f4:f691::3e  5.473 ms !A  5.297 ms !A  3.597 ms !A
```

Redundancy of the network is operating as expected.

Remember to activate the `protocols` hierarchy on R4.

```
[edit]
lab@R4# rollback 1
load complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# commit
commit complete
```

Appendix - Chapter 3: BGP and Routing policy

This chapter continues with the upgrade of a networks of a typical Internet Service Provider. Please continue with the final configs of the previous chapter. You have to configure IBGP and EBGP for exchanging routing information with different (external) domains in the network.

Task 1: Creating Clusters – Initial Setup

BGP Overview

The Border Gateway Protocol (BGP) is an Exterior Gateway Protocol (EGP) used for exchanging routing information between Autonomous Systems (AS). An AS is a group of devices under common administration and each AS represents a different domain in the Internet. BGP is called a path vector routing protocol because it uses the AS path to evaluate the distance to the destination. The AS path is one of the attributes that BGP uses. It represents the path to follow to reach the destination prefix. This is also the method, which BGP uses to prevent routing loops. BGP routers examine if the local AS is already present in the path and if it is, this means that there is routing loop and the route is not used.

BGP uses TCP (port 179), which ensures for reliable communication. BGP routers listen on TCP port 179 for possible new connections. Network Layer Reachability Information (NLRI) is exchanged between BGP peers. BGP is commonly used to advertise aggregated prefixes of an AS's internal prefixes towards another ASs.

BGP supports two different neighbor relationships between peers. If two peers participate in the same AS number, they establish an *internal BGP (IBGP) peering*, if both peers are from a different AS they form an *external BGP (EBGP)* peering. By default, EBGP sessions can only be formed between directly connected neighbors as the time-to-live (TTL) of EBGP packets is 1. EBGP multihop should be used to establish session more than one hop away. IBGP sessions have no such limitation by default and usually are established between the loopback interfaces of the routers for stability purposes.

BGP follows strict rules when advertising routes / prefixes to peers.

- IBGP peers advertise routes received from EBGP peers to other IBGP peers.

© 2020 Juniper Networks, Inc. All rights reserved.

- EBGP peers advertise routes learned from IBGP or EBGP peers to other EBGP peers.
- IBGP peers do not advertise routes known from IBGP peers to other IBGP peers.

The last rule is needed to prevent routing loops inside an AS. IBGP peers participate in the same AS, so the AS path cannot be used for loop detection. Due to this rule the requirement is that IBGP neighbors must be fully meshed to ensure consistency in BGP routing information.

BGP sessions are manually defined. Peer IP addresses are explicitly configured, no neighbor discovery process is used. BGP session can be in one of the following states:

- *Idle state:* This is the initial state when all BGP connection request are rejected by the local BGP process.
- *Connect state:* In this state BGP is waiting for the transport connection to be established. If it is established, the local system sends an Open message and the state transitions to *OpenSent* state. If the connection fails, the local router starts listening for incoming connection request and transitions to the *Active* state.
- *Active state:* local peer is trying to establish connection, if it succeeds connection transition to *OpenSent* state. If not physical connectivity, configuration or the routing to the remote peer should be checked.
- *OpenSent state:* BGP waits for and Open message received from the peer. When it is received, BGP peer checks if it is valid and sends a Keepalive message. If not valid, the peer falls back to the Idle state.
- *OpenConfirm state:* In this state BGP waits for a Keepalive or Notification message. If a Notification massage is received, it means that the hold timer has expired and the peer transition to Idle state. If a Keepalive message is received the connection transitions to the Established state.
- *Established state:* In this state both neighbors can exchange routing information. They can exchange Update, Keepalive and Notification messages.

BGP uses the following messages types:

- Open message: Open message is sent immediately after the TCP connection is established. It is used to initiate the connection and to negotiate the options of the connection.
- Update message: This message is used to exchange routing information between peers. Update messages describe path information and a list of prefixes with their attributes that can be reached through the path.
- Keepalive message: Keepalive messages are sent to ensure that the hold timer doesn't expire.

© 2020 Juniper Networks, Inc. All rights reserved.

- Notification message: Notifications are used to signal an error. If some capability is not supported or for some reason the neighbor cannot send an update message. When an error is detected, the BGP session is terminated.
- Refresh message: Support for soft refresh of prefixes.

One of the main advantages of BGP is the capability of using different attributes to control the advertisement of the routes between ASs and influence path selection. Junos uses routing policies to control route selection by changing the BGP attributes.

BGP attributes are grouped in four main classes:

- Well-known mandatory: These attributes must be supported by all BGP neighbors and must be included in all BGP updates.
- Well-known discretionary: These attributes must be supported by all BGP neighbors, but do not have to be included in all BGP updates.
- Optional transitive: It is not required to be supported by all BGP neighbors, but if present it must be passed to other neighbors.
- Optional nontransitive: Again, not required to be supported by all BGP neighbors. If such an attribute is not recognized, it is ignored and not passed to other neighbors.

These attributes are very important to master for the JNCIE-SP lab exam.

- Next-hop: A well-known attribute. It has to be included for every prefix. It represents the IP address of the router advertising the route. The next-hop should be reachable by the local router to be installed in the routing table. EBGP peers change the next-hop when advertising routes to other EBGP peers. IBGP peer do not change Next-hop attribute by default.
- AS path: This attribute is also well-known mandatory and represents the path to the source of the prefix. Another purpose is loop prevention mechanism as it was stated already. Each EBGP peer prepends its AS to the AS path when advertising prefixes to EBGP peer. The AS path is not updated between IBGP peers.
- Origin: Third well-known mandatory attribute. Origin describes the way the route is injected into BGP. The router that advertises a prefix assigns an origin attribute to that route. IGP (I), EGP (E) and Incomplete (2) are three possible values. Junos by default injects all routes with origin of IGP.
- Local preference: This is well-known discretionary attribute that defines the preferred way out of the local AS. All traffic from the AS is flowing to the peer advertised routes with highest local preference. This is the only attribute where the highest value is preferred. Local preference is transmitted only across IBGP links. By default, the local preference is set to value 100.

© 2020 Juniper Networks, Inc. All rights reserved.

- Multi-Exit Discriminator: MED is an optional non-transitive attribute used to influence the preferred path into the AS when multiple links between two ASs exist. MED is not transferred across neighboring AS. The MED is often called metric and is used for neighboring AS to evaluate which path to use if several links are available between the local and neighboring AS.
- Community: This optional transitive attribute is used to tag (a group of prefixes) to be identified easily. It helps to simplify the policies applied to the routes for routing decisions. They can be used to signal other attributes between ASs.

With so many attributes, BGP uses a complicated algorithm to select the best path. For each prefix received BGP verifies if the next-hop is reachable and that no routing loop exist.

BGP uses the following algorithm if more than one non-hidden path to the same prefix is received.

1) The router first compares the prefixes with the highest local preference. As it was stated this is the only attribute where the highest value is preferred.
2) Next the AS-path is checked where the short AS-path is preferred.
3) Origin code is evaluated. As the lowest origin code is preferred. I < E <?
4) Next if several prefixes are advertised from the same neighboring AS, MED is compared, and lower value is chosen.
5) Routes learned by EBGP neighbors are preferred over routes learned from IBGP peers. If there are only EBGP routes to be considered, the algorithm skips step 6 and proceeds with step 7
6) If remaining routes are learned by IBGP, algorithm choses the route with lowest IGP cost to the IBGP peer.
7) Next BGP evaluates the lowest Router ID of the peers, usually this is the loopback IP address. If the routes are from external peers from two distinct ASs, BGP choses the current active route. This can be changed with a configuration knob so that the Router ID is always compared.
8) Next if Route Reflection is used in the network, the router evaluates the Cluster ID attribute. This is similar to the AS-path. The shortest Cluster list is preferred.
9) As a last resort BGP choses prefixes from peers with the lowest peer IP address.

The outcome of the algorithm is always a single path to a given prefix. If you use the `multipath` option, the algorithm stops before step 7 and all remaining routes are used for forwarding.

© 2020 Juniper Networks, Inc. All rights reserved.

The need for full-mesh IBGP neighbors has scalability problem. If you need to add a new peer, you have to reconfigure all other IBGP peers in the network. It was stated that neighbors in BGP are configured manually. Two solutions are available to limit the IBGP sessions in the network: Route Reflector and Confederations. Each of them will be described in detail later in this chapter.

The first task requires the configuration of confederations on top of the IS-IS network you created in Chapter Two.

Confederations break the global AS in multiple smaller ASs. Inside these smaller ASs (sub-AS) the IBGP peers should be full-meshed again, but the number of the sessions will be much lower. EBGP sessions are needed between each sub-AS. Except for the AS-path attribute, the BGP attributes are not changed on confederation EBGP (C-EBGP) sessions. As mentioned only the AS-path is changed between the sub-ASs, which is used as a loop prevention mechanism. When a prefix is advertised to an EBGP peer, the sub-AS numbers are removed from the AS-path.

AS numbers are assigned by the Regional Internet Registers (RIR). Initially AS numbers were 16bits, so called 2-byte ASs. AS numbers from 64512 to 65535 are for private use and the numbers from that range can be used for sub-AS numbers in confederation deployments.

As with IPv4 addresses, AS numbers were exhausted. 32bits or 4-byte ASs where introduced with backward compatibility to 2-byte AS numbers.

There are two formats of identifying the 4-byte AS numbers: ASPLAIN and ASDOT.

- ASPLAIN is the same as 2-byte AS numbers. The AS numbers are represented in 32bits integers.
- ASDOT format is using two 16bits integers with a single dot in between. For example, AS number 65546 will be displayed in ASDOT format as 1.10, where 1 from left is 65536 plus 10 from right is equal of 65546.

Note: Junos allows the configuration of 4-byte AS numbers in both formats. By default, they are shown in ASPLAIN format, which can be changed with a configuration knob.

A special reserved AS number 23456 is advertised to peers that do not understand new 4-bytes AS numbers. A new attribute (AS4_PATH) is added to preserve the 4-byte AS-path when traversing such peers. When the prefix is advertised to 4-byte capable peers then the original AS path is restored from AS4_PATH attribute.

1) Configure global confederation parameters.

© 2020 Juniper Networks, Inc. All rights reserved.

You have to configure confederations in the network. You have been given the limitation on the number and the type of the BGP sessions that you can establish. You cannot have more than 2 IBGP connections per router and no limit on the CBGP sessions between the routers. One possible solution is to split the routers to 4 sub-ASs with 2 IBGP peers inside. This way each router will have one IBGP session and one CBGP session.

The AS number is globally configured in Junos and the confederation configuration as well. As stated before, private AS numbers are used for sub-ASs. The number 54591 is used for the global AS number.

a) R1 and R2

```
[edit]
lab@R1# set routing-options autonomous-system 65000

[edit]
lab@R1# set routing-options confederation 54591 members [ 65000 65001 65002 65003 ]

[edit]
lab@R1# commit
commit complete
```

Sub-AS number 65000 is configured as the Autonomous System number on R1. The confederation configuration includes all sub-ASs under the global AS number.

Note: The configuration on R2 is the same, as it participates in the same sub-AS.

b) R3 and R4

```
[edit]
lab@R3# set routing-options autonomous-system 65001

[edit]
lab@R3# set routing-options confederation 54591 members [ 65000 65001 65002 65003 ]

[edit]
lab@R3# commit
commit complete
```

Note: Configuration on R4 is the same, as it participates in the same sub-AS.

c) R5 and R6

```
[edit]
lab@R5# set routing-options autonomous-system 65002
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set routing-options confederation 54591 members [ 65000 65001 65002 65003 ]

[edit]
lab@R5# commit
commit complete
```

Note: Configuration on R6 is the same, as it participates in the same sub-AS.

d) R7 and R8

```
[edit]
lab@R7# set routing-options autonomous-system 65003

[edit]
lab@R7# set routing-options confederation 54591 members [ 65000 65001 65002 65003 ]

[edit]
lab@R7# commit
commit complete
```

Note: Configuration on R8 is the same, as it participates in the same sub-AS.

2) Configure IBGP and CBGP sessions.

a) R1

R1 and every other router in the network will have one IBGP and one CBGP session. BGP neighbors are explicitly configured as shown below.

```
[edit]
lab@R1# set protocols bgp group ibgp type internal

[edit]
lab@R1# set protocols bgp group ibgp local-address 172.30.5.1

[edit]
lab@R1# set protocols bgp group ibgp neighbor 172.30.5.2

[edit]
lab@R1# set protocols bgp group ibgp authentication-key jncie

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

BGP neighbors are configured under groups. Here `group ibgp` is created and configured with `type internal`. By default, BGP groups are type external. Configure the group as internal ensures that all neighbors from that group do not need a peer AS configured under the group. As it was stated already usually IBGP neighbors are configured to peer with their loopback IP addresses, but by default BGP peers on directly connected IP addresses. The `local-address` knob instructs IBGP to source the TCP connection from loopback IP address.

BGP also supports session authentication. MD5, SHA1 and AES are the supported authentication algorithms. By default, BGP uses MD5 for authentication.

```
[edit]
lab@R1# set protocols bgp group cbgp type external

[edit]
lab@R1# set protocols bgp group cbgp local-address 172.30.5.1

[edit]
lab@R1# set protocols bgp group cbgp peer-as 65003

[edit]
lab@R1# set protocols bgp group cbgp neighbor 172.30.5.8

[edit]
lab@R1# set protocols bgp group cbgp authentication-key jncie

[edit]
lab@R1# set protocols bgp group cbgp multihop

[edit]
lab@R1# commit
commit complete
```

CBGP configurations is more or less similar to EBGP session config. You can omit `type external` as by default neighbors are external in Junos. The `peer-as` statement is mandatory for the EBGP and CBGP groups. R1 will peer with R8, which participates in sub-AS 65003. With CBGP the loopback IP addresses are used as the source of the BGP TCP connection. Unlike IBGP groups, which can establish sessions more than hop away, the TTL for EBGP sessions is set to 1. To change that behavior and use loopback IP addresses for CBGP sessions you also have to configure `multihop`.

The authentication configuration is the same as for the IBGP sessions.

b) R2

© 2020 Juniper Networks, Inc. All rights reserved.

R2 is in the same sub-AS as R1. For this reason, R2 will have IBGP peering with R1 and a CBGP peering with R3 from sub-AS 65001.

```
[edit]
lab@R2# set protocols bgp group ibgp type internal

[edit]
lab@R2# set protocols bgp group ibgp local-address 172.30.5.2

[edit]
lab@R2# set protocols bgp group ibgp neighbor 172.30.5.1

[edit]
lab@R2# set protocols bgp group ibgp authentication-key jncie

[edit]
lab@R2# set protocols bgp group cbgp type external

[edit]
lab@R2# set protocols bgp group cbgp local-address 172.30.5.2

[edit]
lab@R2# set protocols bgp group cbgp peer-as 65001

[edit]
lab@R2# set protocols bgp group cbgp neighbor 172.30.5.3

[edit]
lab@R2# set protocols bgp group cbgp authentication-key jncie

[edit]
lab@R2# set protocols bgp group cbgp multihop

[edit]
lab@R2# commit
commit complete
```

c) R3

R3 is in the same sub-AS as R4. For this reason, R3 will have IBGP peering with R4 and a CBGP peering with R2 from sub-AS 65000.

```
[edit]
lab@R3# set protocols bgp group ibgp type internal

[edit]
lab@R3# set protocols bgp group ibgp local-address 172.30.5.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols bgp group ibgp neighbor 172.30.5.4

[edit]
lab@R3# set protocols bgp group ibgp authentication-key jncie

[edit]
lab@R3# set protocols bgp group cbgp type external

[edit]
lab@R3# set protocols bgp group cbgp local-address 172.30.5.3

[edit]
lab@R3# set protocols bgp group cbgp peer-as 65000

[edit]
lab@R3# set protocols bgp group cbgp neighbor 172.30.5.2

[edit]
lab@R3# set protocols bgp group cbgp authentication-key jncie

[edit]
lab@R3# set protocols bgp group cbgp multihop

[edit]
lab@R3# commit
commit complete
```

d) R4

R4 is in the same sub-AS as R3. For this reason, R4 will have IBGP peering with R3 and a CBGP peering with R5 from sub-AS 65002.

```
[edit]
lab@R4# set protocols bgp group ibgp type internal

[edit]
lab@R4# set protocols bgp group ibgp local-address 172.30.5.4

[edit]
lab@R4# set protocols bgp group ibgp neighbor 172.30.5.3

[edit]
lab@R4# set protocols bgp group ibgp authentication-key jncie

[edit]
lab@R4# set protocols bgp group cbgp type external
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set protocols bgp group cbgp local-address 172.30.5.4

[edit]
lab@R4# set protocols bgp group cbgp peer-as 65002

[edit]
lab@R4# set protocols bgp group cbgp neighbor 172.30.5.5

[edit]
lab@R4# set protocols bgp group cbgp authentication-key jncie

[edit]
lab@R4# set protocols bgp group cbgp multihop

[edit]
lab@R4# commit
commit complete
```

e) R5

R5 is in the same sub-AS as R6. For this reason, R5 will have IBGP peering with R6 and a CBGP peering with R4 from sub-AS 65001.

```
[edit]
lab@R5# set protocols bgp group ibgp type internal

[edit]
lab@R5# set protocols bgp group ibgp local-address 172.30.5.5

[edit]
lab@R5# set protocols bgp group ibgp neighbor 172.30.5.6

[edit]
lab@R5# set protocols bgp group ibgp authentication-key jncie

[edit]
lab@R5# set protocols bgp group cbgp type external

[edit]
lab@R5# set protocols bgp group cbgp local-address 172.30.5.5

[edit]
lab@R5# set protocols bgp group cbgp peer-as 65001

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set protocols bgp group cbgp neighbor 172.30.5.4

[edit]
lab@R5# set protocols bgp group cbgp authentication-key jncie

[edit]
lab@R5# set protocols bgp group cbgp multihop

[edit]
lab@R5# commit
commit complete
```

f) R6

R6 is in the same sub-AS as R5. For this reason, R6 will have IBGP peering with R5 and a CBGP peering with R7 from sub-AS 65003.

```
[edit]
lab@R6# set protocols bgp group ibgp type internal

[edit]
lab@R6# set protocols bgp group ibgp local-address 172.30.5.6

[edit]
lab@R6# set protocols bgp group ibgp neighbor 172.30.5.5

[edit]
lab@R6# set protocols bgp group ibgp authentication-key jncie

[edit]
lab@R6# set protocols bgp group cbgp type external

[edit]
lab@R6# set protocols bgp group cbgp local-address 172.30.5.6

[edit]
lab@R6# set protocols bgp group cbgp peer-as 65003

[edit]
lab@R6# set protocols bgp group cbgp neighbor 172.30.5.7

[edit]
lab@R6# set protocols bgp group cbgp authentication-key jncie

[edit]
lab@R6# set protocols bgp group cbgp multihop
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# commit
commit complete
```

g) R7

R7 is in the same sub-AS as R8. For this reason, R7 will have IBGP peering with R8 and a CBGP peering with R6 from sub-AS 65002.

```
[edit]
lab@R7# set protocols bgp group ibgp type internal

[edit]
lab@R7# set protocols bgp group ibgp local-address 172.30.5.7

[edit]
lab@R7# set protocols bgp group ibgp neighbor 172.30.5.8

[edit]
lab@R7# set protocols bgp group ibgp authentication-key jncie

[edit]
lab@R7# set protocols bgp group cbgp type external

[edit]
lab@R7# set protocols bgp group cbgp local-address 172.30.5.7

[edit]
lab@R7# set protocols bgp group cbgp peer-as 65002

[edit]
lab@R7# set protocols bgp group cbgp neighbor 172.30.5.6

[edit]
lab@R7# set protocols bgp group cbgp authentication-key jncie

[edit]
lab@R7# set protocols bgp group cbgp multihop

[edit]
lab@R7# commit
commit complete
```

h) R8

© 2020 Juniper Networks, Inc. All rights reserved.

R8 is in sub-AS 65003 and has an IBGP neighborship with R7 and also an CBGP neighborship with R1 in sub-AS 65000.

```
[edit]
lab@R8# set protocols bgp group ibgp type internal

[edit]
lab@R8# set protocols bgp group ibgp local-address 172.30.5.8

[edit]
lab@R8# set protocols bgp group ibgp neighbor 172.30.5.7

[edit]
lab@R8# set protocols bgp group ibgp authentication-key jncie

[edit]
lab@R8# set protocols bgp group cbgp type external

[edit]
lab@R8# set protocols bgp group cbgp local-address 172.30.5.8

[edit]
lab@R8# set protocols bgp group cbgp neighbor 172.30.5.1

[edit]
lab@R8# set protocols bgp group cbgp peer-as 65000

[edit]
lab@R8# set protocols bgp group cbgp authentication-key jncie

[edit]
lab@R8# set protocols bgp group cbgp multihop

[edit]
lab@R8# commit
commit complete
```

Verify the IBGP and CBGP sessions on the selected routers.

```
[edit]
lab@R1# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                         0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps    Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.2           65000        388        390       0       0   2:54:30 0/0/0/0
0/0/0/0
172.30.5.8           65003          3          3       0       0         19 0/0/0/0
0/0/0/0
```

The output of the command above is used to receive information about BGP neighbors. You can see that both neighbors on R1 are established. When the sessions are established the number of the routes Active, Received, Accepted and Damped are shown. The word **`Established`** is not shown. On R1, 0 routes are received because nothing is advertised yet in BGP. Peers are shown with their IP addresses and AS numbers as they are configured.

In case there is a problem with the BGP session, the State is shown instead of the number of received routes.

```
[edit]
lab@R1# run show bgp summary
Groups: 2 Peers: 2 Down peers: 1
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                         0          0          0          0          0          0
Peer                     AS    InPkt    OutPkt  OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.2           65000        383        385     0  0     2:52:07 0/0/0/0
0/0/0/0
172.30.5.8           65003          0          0     0  0     2:49:15 Active
```

Neighbor R8 is in Active state and waiting for connection attempts.

Display the details for the BGP session with R2.

```
[edit]
lab@R1# run show bgp neighbor 172.30.5.2 | no-more
Peer: 172.30.5.2+60804 AS 65000 Local: 172.30.5.1+179 AS 65000
  Group: ibgp                   Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress AuthKey Refresh Confed>
  Authentication key is configured
  Local Address: 172.30.5.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 172.30.5.2      Local ID: 172.30.5.1        Active Holdtime: 90
  Keepalive Interval: 30         Group index: 0    Peer index: 0    SNMP index: 0
  I/O Session Thread: bgpio-0 State: Enabled
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
BFD: disabled, down
NLRI for restart configured on peer: inet-unicast
NLRI advertised by peer: inet-unicast
NLRI for this session: inet-unicast
Peer supports Refresh capability (2)
Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
Restart flag received from the peer: Notification
NLRI that restart is negotiated for: inet-unicast
NLRI of received end-of-rib markers: inet-unicast
NLRI of all end-of-rib markers sent: inet-unicast
Peer does not support LLGR Restarter functionality
Peer supports 4 byte AS extension (peer-as 65000)
Peer does not support Addpath
Table inet.0 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:              0
  Received prefixes:            0
  Accepted prefixes:            0
  Suppressed due to damping:    0
  Advertised prefixes:          0
Last traffic (seconds): Received 24   Sent 22   Checked 536
Input messages:  Total 22     Updates 1       Refreshes 0     Octets 466
Output messages: Total 20     Updates 0       Refreshes 0     Octets 384
Output Queue[1]: 0            (inet.0, inet-unicast)
```

This session is Internal, the local router services this BGP session and R2 is the initiator. The state is Established and the agreed keepalive and holdtime in seconds are shown. Family `inet-unicast` is the only negotiated NLRI for the session. This is the default address family and does not need to be manually configured. However, keep in mind that if you add additional families, you have to explicitly configure `family inet-unicast`.

Both peers on R2 are `Established` as well.

```
[edit]
lab@R2# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer             AS   InPkt   OutPkt  OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.1    65000     111      112     0  0       48:49 0/0/0/0
0/0/0/0
172.30.5.3    65001      48       48     0  0       20:18 0/0/0/0
0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                         0          0          0          0          0          0
Peer                 AS    InPkt    OutPkt  OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.1    65000       2        3     0   0        10 0/0/0/0
0/0/0/0
172.30.5.7    65003      19       21     0   0      8:03 0/0/0/0
0/0/0/0
```

R8 has two established BGP connections. One CBGP with R1 and one IBGP with R7.

Task 2: EBGP Configuration

The second task requires the configuration of EBGP sessions with the Internet Exchanges (IX), Peers (P) and Customer (C) BGP peers.

1) Configure EBGP on R1.

According to Figure 7, R1 and R2 have to build EBGP sessions with IX-1 and IX-2 in the same AS number 1620. Four routers share one broadcast segment between them.

You have to configure the interface on R1, which is part of that broadcast segment.

```
[edit]
lab@R1# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R1# set interfaces ge-0/0/5.300 vlan-id 300 family inet address 192.168.1.1/24

[edit]
lab@R1# commit
commit complete
```

Now you can configure the EBGP peer over the new interface.

```
[edit]
lab@R1# set protocols bgp group IX peer-as 1620

[edit]
lab@R1# set protocols bgp group IX neighbor 192.168.1.3

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set protocols bgp group IX neighbor 192.168.1.4

[edit]
lab@R1# set protocols bgp group IX log-updown

[edit]
lab@R1# commit
commit complete
```

Group IX includes both EBGP peers, IX-1 and IX-2 because they share the same parameters. The `log-update` configuration accomplishes the requirement that EBGP state changes must be logged as syslog message.

```
[edit]
lab@R1# run show bgp summary
Groups: 3 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    1128        564          0          0          0          0
Peer            AS  InPkt  OutPkt  OutQ  Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.2  65000    461     962     0      0     3:27:24 0/0/0/0
0/0/0/0
172.30.5.8  65003     77     575     0      0       33:13 0/0/0/0
0/0/0/0
192.168.1.3  1620     511       4     0      0          29 564/564/564/0
0/0/0/0
192.168.1.4  1620     504       2     0      0          26 0/564/564/0
0/0/0/0
```

After the sessions are established. R1 receives the same number of prefixes from both IX peers. The prefixes from IX-1 are chosen as the active ones.

2) Configure EBGP on R2.

Because R2 shares the same segment with R1 and both IX peers, the configuration for EBGP is practically the same as on R1.

Again, you have to configure the broadcast interface.

```
[edit]
lab@R2# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R2# set interfaces ge-0/0/5.300 vlan-id 300 family inet address 192.168.1.2/24

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# set protocols bgp group IX peer-as 1620

[edit]
lab@R2# set protocols bgp group IX neighbor 192.168.1.3

[edit]
lab@R2# set protocols bgp group IX neighbor 192.168.1.4


[edit]
lab@R2# set protocols bgp group IX log-updown

[edit]
lab@R2# commit
commit complete
```

Display the BGP summary in R2.

```
[edit]
lab@R2# run show bgp summary
Groups: 3 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    1692        564          0          0          0          0
Peer             AS   InPkt  OutPkt  OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.1   65000     965     937     0     0  3:29:06 0/564/564/0
0/0/0/0
172.30.5.3   65001     403    1373     0     0  3:00:35 0/0/0/0
0/0/0/0
192.168.1.3   1620     473       5     0     0     1:29 0/564/564/0
0/0/0/0
192.168.1.4   1620     474       6     0     0     1:32 564/564/564/0
0/0/0/0
```

R2 receives the same number prefixes from the IX peers in AS 1620. The BGP peer with IP address 192.168.1.4 is chosen as the gateway for the prefixes. R2 receives the prefixes from R1 also via the IBGP session.

Verify the received routes in the routing table. Display the extended output for one received prefix through BGP on R1 and R2.

```
[edit]
lab@R1# run show route protocol bgp 1.84.160.0/20 extensive

inet.0: 631 destinations, 1759 routes (631 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
1.84.160.0/20 (3 entries, 1 announced)
TSI:
KRT in-kernel 1.84.160.0/20 -> {192.168.1.3}
Page 0 idx 0 Type 1 val 1794134
    Nexthop: 192.168.1.3
    Localpref: 100
    AS path: [65000] 1620 33112 I
    Communities:
Page 0 idx 1 Type 1 val 1547e74
    Nexthop: 192.168.1.3
    Localpref: 100
    AS path: [65000] 1620 33112 I
    Communities:
Path 1.84.160.0 from 192.168.1.3 Vector len 4.  Val: 0 1
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 590
                Address: 0x1574c88
                Next-hop reference count: 1692
                Source: 192.168.1.3
                Next hop: 192.168.1.3 via ge-0/0/5.300, selected
                State: <Active Ext>
                Local AS: 65000 Peer AS:  1620
                Age: 7:26
                Task: BGP_1620.192.168.1.3+179
                Announcement bits (3): 0-KRT 2-BGP_RT_Background 3-Resolve tree 4
                AS path: 1620 33112 I
                Accepted
                Localpref: 100
                Router ID: 172.17.63.1
         BGP    Preference: 170/-101
                Next hop type: Router
                Address: 0x1574c3c
                Next-hop reference count: 1129
                Source: 192.168.1.4
                Next hop: 192.168.1.4 via ge-0/0/5.300, selected
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65000 Peer AS:  1620
                Age: 7:22
                Task: BGP_1620.192.168.1.4+60474
                AS path: 1620 33112 I
                Accepted
                Localpref: 100
                Router ID: 172.17.63.2
BGP             Preference: 170/-101
                Next hop type: Indirect
                Address: 0x15749dc
                Next-hop reference count: 564
                Source: 172.30.5.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                Next hop type: Router
                Next hop: 192.168.1.4 via ge-0/0/5.300, selected
                Protocol next hop: 192.168.1.4
                Indirect next hop: 1773830 -
                State: <NotBest Int Ext>
                Inactive reason: Not Best in its group - Interior > Exterior >
Exterior via Interior
                Local AS: 65000 Peer AS: 65000
                Age: 6:47       Metric2: 0
                Task: BGP_65000.172.30.5.2+179
                AS path: 1620 33112 I
                Accepted
                Localpref: 100
                Router ID: 172.30.5.2
                Indirect next hops: 1
                        Protocol next hop: 192.168.1.4
                        Indirect next hop: 1773830 -
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 192.168.1.4 via ge-0/0/5.300
                        192.168.1.0/24 Originating RIB: inet.0
                          Node path count: 1
                          Forwarding nexthops: 1
                                Next hop type: Interface
                                Nexthop: via ge-0/0/5.300
```

The chosen active route is marked with "*". R1 received the route from three peers. IX-1 with IP address 192.168.1.3, IX-2 with IP address 192.168.1.4 and from R2 with IP address 172.30.5.2. The prefix received from R2 is not preferred because it is received from an Internal peer. In fact, it is an External prefix received from an internal source. From the other two, R1 choses the prefix from IX-1 because it was first entered in the routing table.

Verify the same route on R2.

```
[edit]
lab@R2# run show route protocol bgp 1.84.160.0/20 extensive

inet.0: 631 destinations, 1759 routes (631 active, 0 holddown, 0 hidden)
1.84.160.0/20 (3 entries, 1 announced)
TSI:
KRT in-kernel 1.84.160.0/20 -> {192.168.1.4}
Page 0 idx 0 Type 1 val 17daae0
    Nexthop: 192.168.1.4
    Localpref: 100
    AS path: [65000] 1620 33112 I
    Communities:
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Page 0 idx 1 Type 1 val 1545ab0
    Nexthop: 192.168.1.4
    Localpref: 100
    AS path: [65000] 1620 33112 I
    Communities:
Path 1.84.160.0 from 192.168.1.4 Vector len 4.  Val: 0 1
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 590
                Address: 0x1574a74
                Next-hop reference count: 1692
                Source: 192.168.1.4
                Next hop: 192.168.1.4 via ge-0/0/5.300, selected
                State: <Active Ext>
                Local AS: 65000 Peer AS:  1620
                Age: 3:08
                Task: BGP_1620.192.168.1.4+54080
                Announcement bits (3): 0-KRT 2-BGP_RT_Background 3-Resolve tree 4
                AS path: 1620 33112 I
                Accepted
                Localpref: 100
                Router ID: 172.17.63.2
    BGP         Preference: 170/-101
                Next hop type: Router, Next hop index: 591
                Address: 0x1574e04
                Next-hop reference count: 1130
                Source: 192.168.1.3
                Next hop: 192.168.1.3 via ge-0/0/5.300, selected
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65000 Peer AS:  1620
                Age: 3:05
                Task: BGP_1620.192.168.1.3+179
                AS path: 1620 33112 I
                Accepted
                Localpref: 100
                Router ID: 172.17.63.1
            BGP Preference: 170/-101
                Next hop type: Indirect
                Address: 0x1574b0c
                Next-hop reference count: 564
                Source: 172.30.5.1
                Next hop type: Router, Next hop index: 591
                Next hop: 192.168.1.3 via ge-0/0/5.300, selected
                Protocol next hop: 192.168.1.3
                Indirect next hop: 176c000 262142
                State: <NotBest Int Ext>
                Inactive reason: Not Best in its group - Interior > Exterior >
Exterior via Interior
                Local AS: 65000 Peer AS: 65000
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Age: 3:47       Metric2: 0
Task: BGP_65000.172.30.5.1+56672
AS path: 1620 33112 I
Accepted
Localpref: 100
Router ID: 172.30.5.1
Indirect next hops: 1
        Protocol next hop: 192.168.1.3
        Indirect next hop: 176c000 262142
        Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 192.168.1.3 via ge-0/0/5.300
        192.168.1.0/24 Originating RIB: inet.0
          Node path count: 1
          Forwarding nexthops: 1
                Next hop type: Interface
                Nexthop: via ge-0/0/5.300
```

R2 also has the same prefix received from IX-1, IX-2 and from R1. The prefix from R1 is not chosen as it is received from an Internal BGP peer. R2 selects the prefix from IX-2 as it was first received.

3) Configuration of EBGP on R3.

Router R3 has two peers, P2-1 from AS 53732.2005 and P2-2 43208.365. You have to add two VLAN sub-interfaces on R3 for reachability towards those peers.

```
[edit]
lab@R3# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R3# set interfaces ge-0/0/5.301 vlan-id 301 family inet address 192.168.0.1/30

[edit]
lab@R3# set interfaces ge-0/0/5.302 vlan-id 302 family inet address 192.168.0.5/30

[edit]
lab@R3# commit
commit complete
```

Now the EBGP neighbors can be configured. Two separate BGP groups are used for both peers.

```
[edit]
lab@R3# set protocols bgp group P2-1 peer-as 53732.2005

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set protocols bgp group P2-1 neighbor 192.168.0.2

[edit]
lab@R3# set protocols bgp group P2-1 log-updown

[edit]
lab@R3# set protocols bgp group P3-1 peer-as 43208.365


[edit]
lab@R3# set protocols bgp group P3-1 neighbor 192.168.0.6

[edit]
lab@R3# set protocols bgp group P3-1 log-updown

[edit]
lab@R3# commit
commit complete
```

Verify the BGP session table for the new EBGP neighbors.

```
[edit]
lab@R3# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     770        211          0          0          0          0
Peer                 AS   InPkt  OutPkt  OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.2        65000    1561     600     0       0  3:15:48 0/557/557/0
0/0/0/0
172.30.5.4        65001     428     593     0       0  3:12:29 0/0/0/0
0/0/0/0
192.168.0.2  3521382357      81      98     0       0       22 96/97/97/0
0/0/0/0
192.168.0.6  2831679853      90      81     0       0       27 115/116/116/0
0/0/0/0
```

Both neighbors advertise a different number of prefixes. R3 also receives External routes from R2, but they are not active at this moment.

4) Configure EBGP sessions on R5.

Because R4 has no EBGP neighbors (according Figure 7) you proceed with R5. R5 should have two sessions with the Customer peers.

Again, the interfaces should be configured for connectivity to the BGP peers.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R5# set interfaces ge-0/0/5.303 vlan-id 303 family inet address 192.168.0.9/30

[edit]
lab@R5# set interfaces ge-0/0/5.304 vlan-id 304 family inet address 192.168.0.13/30

[edit]
lab@R5# commit
commit complete
```

Now you can configure both EBGP neighbors in single groups, as both Customer peers are from the same neighboring AS.

```
[edit]
lab@R5# set protocols bgp group C3 peer-as 64514

[edit]
lab@R5# set protocols bgp group C3 neighbor 192.168.0.10

[edit]
lab@R5# set protocols bgp group C3 neighbor 192.168.0.14

[edit]
lab@R5# set protocols bgp group C3 multipath

[edit]
lab@R5# set protocols bgp group C3 log-updown

[edit]
lab@R5# commit
commit complete
```

To accomplish the load balancing requirement, you need to use `multipath` for these two neighbors. This way the BGP best path algorithm will stop before step 7, where the Router ID is evaluated.

```
[edit]
lab@R5# run show bgp summary
Groups: 3 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     243        243          0          0          0          0
Peer             AS  InPkt  OutPkt  OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.4      65001     614      612      0      0    3:22:00 211/211/211/0
0/0/0/0
172.30.5.6      65002     432      595      0      0    3:14:12 0/0/0/0
0/0/0/0
192.168.0.10    64514       3      170      0      0          9 16/16/16/0
0/0/0/0
192.168.0.14    64514       3      161      0      0          5 16/16/16/0
0/0/0/0
```

16 routes are received from both External neighbors. All 16 routes are active in routing table because of the `multipath` configuration.

```
[edit]
lab@R5# run show route protocol bgp aspath-regex "64514 *"

inet.0: 300 destinations, 316 routes (300 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.32.0/24     *[BGP/170] 00:01:45, localpref 100, from 192.168.0.10
                      AS path: 64514 I
                      to 192.168.0.10 via ge-0/0/5.303
                    > to 192.168.0.14 via ge-0/0/5.304
                    [BGP/170] 00:01:41, localpref 100
                      AS path: 64514 I
                    > to 192.168.0.14 via ge-0/0/5.304
172.31.33.0/24     *[BGP/170] 00:01:45, localpref 100
                      AS path: 64514 I
                    > to 192.168.0.10 via ge-0/0/5.303
                      to 192.168.0.14 via ge-0/0/5.304
                    [BGP/170] 00:01:41, localpref 100
                      AS path: 64514 I
                    > to 192.168.0.14 via ge-0/0/5.304
...
```

Note: You can use `aspath-regex` to filter the output from the routing table. In the example above a regular expression is used to filter the prefixes, which have the AS path 64514 listed on the most left of the AS-PATH string. The AS numbers on the right of the 64514 number, can be any number and with any length. This is represented with “*”.

From the output above you can see that prefixes received from AS 64514, have both neighbors IP addresses selected as next-hop. Despite both next-hops in the routing table, Junos selects a single next-hop for the forwarding table. Let's display the forwarding-table entry for one of the routes with two next-hops.

```
[edit]
lab@R5# run show route forwarding-table destination 172.31.32.0/24
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
172.31.32.0/24   user      0 192.168.0.14  ucst   628      8 ge-0/0/5.304

Routing table: __master.anon__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   527     1
```

A single next-hop is selected, in this case 192.168.0.14 using the interface ge-0/0/5.304. You may actually see the other next hop selected. The key here is to recognize the same next hop is also in the forwarding table as a single entry.

You have to alter the default behavior. This is accomplished with a load-balancing policy with a special terminating action.

```
[edit]
lab@R5# set policy-options policy-statement load_balance term 1 then load-balance per-
packet

[edit]
lab@R5# set routing-options forwarding-table export load_balance

[edit]
lab@R5# commit
commit complete
```

Note: Despite the `per-packet` configuration in the `then` statement. Junos devices support only per-flow load-balancing.

The load-balance policy is applied in the export direction to the forwarding table.

Verify the result of the load-balance policy with the same route in the forwarding-table.

```
[edit]
lab@R5# run show route forwarding-table destination 172.31.32.0/24
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
172.31.32.0/24   user      0              ulst 262147    17
                           192.168.0.10   ucst   615     3 ge-0/0/5.303
                           192.168.0.14   ucst   628     4 ge-0/0/5.304

Routing table: __master.anon__.inet
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   527     1
...
```

As you can see, now both peers are selected as next-hops and traffic is effectively load-balanced in the forwarding table.

5) Configure EBGP sessions on R6.

The task requires R6 to establish an EBGP session to the loopback IP address of C2-1 neighbor. RIP should be used to receive the route of the loopback IP address. Because a failure of any of the two physical interfaces should not result in a neighbor failure, you have to configure RIP on both interfaces.

First you have to configure the interface to establish connectivity to both peers C2-1 and C1-1.

```
[edit]
lab@R6# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R6# set interfaces ge-0/0/5.305 vlan-id 305 family inet address 192.168.0.17/30

[edit]
lab@R6# set interfaces ge-0/0/5.306 vlan-id 306 family inet address 192.168.0.21/30

[edit]
lab@R6# set interfaces ge-0/0/5.307 vlan-id 307 family inet address 192.168.0.25/30

[edit]
lab@R6# commit
commit complete
```

Configure the RIP protocol to learn the loopback IP address.

```
[edit]
lab@R6# set protocols rip group peer neighbor ge-0/0/5.306

[edit]
lab@R6# set protocols rip group peer neighbor ge-0/0/5.307
```

As you already know the RIP configuration is simple. You just have to add both interface under the RIP group.

© 2020 Juniper Networks, Inc. All rights reserved.

If you remember from the previous chapter, Junos only accepts routes from RIP neighbors, but does not advertise anything by default. You have to configure and apply export policy to export R6 loopback IP address to both neighbors.

```
[edit]
lab@R6# set policy-options policy-statement loopback-to-rip term 1 from protocol direct

[edit]
lab@R6# set policy-options policy-statement loopback-to-rip term 1 from route-filter
172.30.5.6/32 exact

[edit]
lab@R6# set policy-options policy-statement loopback-to-rip term 1 then accept

[edit]
lab@R6# set protocols rip group peer export loopback-to-rip

[edit]
lab@R6# commit
commit complete
```

Verify if the loopback IP address of C2-1 is received by R6.

```
[edit]
lab@R6# run show rip neighbor
            Local  Source          Destination     Send   Receive   In
Neighbor    State  Address         Address         Mode   Mode     Met
--------    -----  -------         -----------     ----   -------  ---
ge-0/0/5.306   Up 192.168.0.21    224.0.0.9       mcast  both       1
ge-0/0/5.307   Up 192.168.0.25    224.0.0.9       mcast  both       1

[edit]
lab@R6# run show route protocol rip

inet.0: 300 destinations, 300 routes (73 active, 0 holddown, 227 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.31.1/32     *[RIP/100] 00:04:43, metric 2, tag 0
                    > to 192.168.0.22 via ge-0/0/5.306
                      to 192.168.0.26 via ge-0/0/5.307
224.0.0.9/32       *[RIP/100] 00:04:43, metric 1
                       MultiRecv

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

inet6.0: 31 destinations, 32 routes (31 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

C2-1 neighbor is reachable through both interfaces and the loopback IP address is present in the routing table.

```
[edit]
lab@R6# run show route advertising-protocol rip 192.168.0.21

inet.0: 300 destinations, 300 routes (73 active, 0 holddown, 227 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.6/32       *[Direct/0] 05:39:35
                     > via lo0.0

[edit]
lab@R6# run show route advertising-protocol rip 192.168.0.25

inet.0: 300 destinations, 300 routes (73 active, 0 holddown, 227 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.6/32       *[Direct/0] 05:39:42
                     > via lo0.0
```

R6 loopback IP address is advertised by RIP protocol on R6 to C2-1.

```
[edit]
lab@R6# set protocols bgp group C2-1 local-address 172.30.5.6

[edit]
lab@R6# set protocols bgp group C2-1 neighbor 172.31.31.1

[edit]
lab@R6# set protocols bgp group C2-1 peer-as 64513

[edit]
lab@R6# set protocols bgp group C2-1 multihop

[edit]
lab@R6# set protocols bgp group C2-1 log-updown

[edit]
lab@R6# commit
commit complete
```

You have to use `multihop` knob to specify that this EBGP neighbor is not using the direct link IP addresses (TTL) and the `local-address` to change the source IP address to loopback IP address.

Verify the BGP session table.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       243         16          0          0          0          0
Peer             AS  InPkt  OutPkt  OutQ  Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.5   65002    649     489     0      0  3:39:34 0/227/227/0
0/0/0/0
172.30.5.7   65003    480     482     0      0  3:36:25 0/0/0/0
0/0/0/0
172.31.31.1 64513      3       3     0      0       26 16/16/16/0
0/0/0/0
```

The EBGP multihop session with C2-1 is established and 16 routes are received.

The EBGP neighbor on R6 should be established with C1-1, here you should configure a limit on the prefixes that are received from this neighbor.

By default, Junos accepts all prefixes from BGP peers. You can limit the accepted prefixes for each BGP family independently. If the remote neighbor sends more than the configured prefixes you have the option to tear down the session for a specified time in minutes or forever. You can also specify that after reaching a certain percentage of the maximum prefixes configured, that Junos generates a syslog message.

```
[edit]
lab@R6# set protocols bgp group C1-1 neighbor 192.168.0.18

[edit]
lab@R6# set protocols bgp group C1-1 peer-as 64512

[edit]
lab@R6# set protocols bgp group C1-1 family inet unicast prefix-limit maximum 20

[edit]
lab@R6# set protocols bgp group C1-1 family inet unicast prefix-limit teardown idle-
timeout 3

[edit]
lab@R6# set protocols bgp group C1-1 log-updown

[edit]
lab@R6# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

EBGP is limited for 20 prefixes for family `inet unicast`. If the maximum number of prefixes is reached the session goes down for 3 minutes.

```
[edit]
lab@R6# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     259         32          0          0          0          0
Peer             AS   InPkt  OutPkt  OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.5   65002     659     502     0       0  3:44:06 0/227/227/0
0/0/0/0
172.30.5.7   65003     491     495     0       0  3:40:57 0/0/0/0
0/0/0/0
172.31.31.1  64513      13      16     0       0     4:58 16/16/16/0
0/0/0/0
192.168.0.18 64512       7       6     0       0     1:09 16/16/16/0
0/0/0/0
```

Neighbor with C1-1 is established and 16 routes are received.

6) Configure EBGP on R7.

R7 should have two EBGP neighbors, one with C1-1 and one with P1-2. As with R6, session with C1-1 should be with prefix limit configured.

First the interfaces should be configured for IP connectivity.

```
[edit]
lab@R7# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R7# set interfaces ge-0/0/5.308 vlan-id 308 family inet address 192.168.0.29/30

[edit]
lab@R7# set interfaces ge-0/0/5.309 vlan-id 309 family inet address 192.168.0.33/30

[edit]
lab@R7# commit
commit complete
```

Configure the EBGP neighbor with C1-1.

```
[edit]
lab@R7# set protocols bgp group C1-1 peer-as 64512
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols bgp group C1-1 neighbor 192.168.0.34

[edit]
lab@R7# set protocols bgp group C1-1 family inet unicast prefix-limit maximum 20


[edit]
lab@R7# set protocols bgp group C1-1 family inet unicast prefix-limit teardown idle-
timeout 3

[edit]
lab@R7# set protocols bgp group C1-1 log-updown
```

Configure the EBGP neighbor with P1-2.

```
[edit]
lab@R7# set protocols bgp group P1-2 peer-as 1679.12483

[edit]
lab@R7# set protocols bgp group P1-2 neighbor 192.168.0.30

[edit]
lab@R7# set protocols bgp group P1-2 log-updown

[edit]
lab@R7# commit
commit complete
```

Verify the BGP session table for the EBGP neighbors.

```
[edit]
lab@R7# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     612        595          0          0          0          0
Peer                 AS   InPkt OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.6       65002   504   995       0       0 3:44:49 15/32/32/0
0/0/0/0
172.30.5.8       65003   237   732       0       0 1:45:20 0/0/0/0
0/0/0/0
192.168.0.30 110047427   492     7       0       0       4 564/564/564/0
0/0/0/0
192.168.0.34     64512     5   496       0       0       6 16/16/16/0
0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Both EBGP sessions are established. As on R7, R7 is receiving 16 prefixes from C1-1. More than 500 prefixes are received from P1-2.

7) Configure EBGP session on R8.

R8 should have just one EBGP neighbor with P1-1 form the same 4-byte AS as the neighbor P1-2 on R7.

Configure the interface for reaching P1-1 on R8.

```
[edit]
lab@R8# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R8# set interfaces ge-0/0/5.310 vlan-id 310 family inet address 192.168.0.37/30
```

Adding the EBGP neighbor configuration.

```
[edit]
lab@R8# set protocols bgp group P1-1 peer-as 1679.12483

[edit]
lab@R8# set protocols bgp group P1-1 neighbor 192.168.0.38

[edit]
lab@R8# set protocols bgp group P1-1 log-updown

[edit]
lab@R8# commit
commit complete
```

Verify the BGP session table.

```
[edit]
lab@R8# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    1357        791          0          0          0          0
Peer               AS  InPkt  OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.1     65000  974      1180       0       0     1:43:28 652/654/654/0
0/0/0/0
172.30.5.7     65003  770      1240       0       0     1:51:21 31/139/139/0
0/0/0/0
192.168.0.38 110047427 492      570        0       0          11 108/564/564/0
0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

R8 receives the same number of prefixes as R7 from neighbor P1-1. But only part of them are active.

8) Configure IPv6 EBGP sessions.

a) R3

For the IPv6 sessions to P2-1 on R3, you have to use link-local IPv6 addresses.

For this purpose, you have to enable `family inet6` on the interfaces towards P2-1.

Note: Depending on the kit you are working on you may need to get the link local address form ge-0/0/5.301 and apply it as the neighbor address in the routing-instances P2-1's BGP neighborship.

```
[edit]
lab@R3# set interfaces ge-0/0/5.301 family inet6

[edit]
lab@R3# commit
commit complete
```

You have to use `traceoptions` to find the IPv6 link-local address of neighbor P2-1. Activate `traceoptions` for BGP packets with flag detail.

```
[edit]
lab@R3# set protocols bgp traceoptions file bgp.log

[edit]
lab@R3# set protocols bgp traceoptions flag packets detail

[edit]
lab@R3# commit
commit complete
```

After committing the configuration, you can check the log file for IPv6 session attempts from P2-1. Because there is a lot of noise from the other BGP connection on the "wire" you can use **`match`** to search through the log file.

```
[edit]
lab@R3# run show log bgp.log | match bgp_listen
Aug 29 20:15:28.738997 bgp_listen_accept: Connection attempt from unconfigured
neighbor: fe80::250:5601:2da9:5bf3+52343
```


© 2020 Juniper Networks, Inc. All rights reserved.

You can see that in the log there are connection attempts from an un-configured neighbor for the local BGP process. The message includes IP address and the source port of the connection attempt.

Configure IPv6 EBGP neighbor on R3.

```
[edit]
lab@R3# set protocols bgp group P2-1-ipv6 peer-as 53732.2005

[edit]
lab@R3# set protocols bgp group P2-1-ipv6 local-interface ge-0/0/5.301

[edit]
lab@R3# set protocols bgp group P2-1-ipv6 neighbor fe80::250:5601:2da9:5bf3

[edit]
lab@R3# set protocols bgp group P2-1-ipv6 log-updown

[edit]
lab@R3# commit
commit complete
```

Because you are using IPv6 link-local addresses for this session, you must specify which interface link-local IPv6 address to be used. Link-local IPv6 addresses on each interface share the same subnet – fe80::/10, so by specifying the `local-interface` the BGP process knows on which interface to establish the connection.

Verify the IPv6 EBGP session.

```
[edit]
lab@R3# run show bgp summary
Groups: 5 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     832        211          0          0          0          0
inet6.0
                      16         16          0          0          0          0
Peer                     AS   InPkt  OutPkt  OutQ  Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.2       65000   1320      572     0      0  2:26:50 0/588/588/0
0/0/0/0
172.30.5.4       65001   334       568     0      0  2:25:22 0/31/31/0
0/0/0/0
192.168.0.2 3521382357   132       150     0      1    23:46 96/97/97/0
0/0/0/0
192.168.0.6 2831679853   401       486     0      0   2:24:51 115/116/116/0
0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
fe80::223:9c01:2d8b:6c81   3521382357    10     11     0        0          3:34 Establ
  inet6.0: 16/16/16/0
```

The IPv6 EBGP session is established. The output of the BGP summary command displays the IPV6 neighbor in a different way. Now you can see the word `Establ` as the state and the routes received for table `inet6.0` one line below. As with IPv4 when you configure IPv6 addresses peer addresses it automatically enables family inet6 unicast for that specific BGP group.

```
[edit]
lab@R3# run show route receive-protocol bgp fe80::250:5601:2da9:5bf3

inet.0: 878 destinations, 919 routes (282 active, 0 holddown, 635 hidden)

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

inet6.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  fd01:aaaa:bbbb:1::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:1::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:2::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:3::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:4::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:5::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:6::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:7::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:8::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:9::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:a::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:b::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:c::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:d::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
  fd01:aaaa:bbbb:1:e::/80
*                         fe80::250:5601:2da9:5bf3               3521382357 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  fd01:aaaa:bbbb:1:f::/80
*                         fe80::250:5601:2da9:5bf3                    3521382357 I
```

16 IPv6 prefixes are received from R3.

Before moving on, delete or deactivate traceoptions on BGP:

```
[edit]
lab@R3# delete protocols bgp traceoptions

[edit]
lab@R3# commit
commit complete
```

b) R5.

R5 should use the existing IPv4 EBGP connection to C3-1 and C3-2. To resolve the IPv6 routes through the IPv4 connection, IPv4-compatible IPv6 prefixes are configured. The format of the addresses is in the form of ::w.x.y.z, where the w.x.y.z is the IPv4 address of the interface. The subnet mask of the address must be longer than 96 bits.

Configure IPv4-compatible IPv6 addresses on interfaces towards C3-1 and C3-2.

```
[edit]
lab@R5# set interfaces ge-0/0/5.303 family inet6 address ::192.168.0.9/126

[edit]
lab@R5# set interfaces ge-0/0/5.304 family inet6 address ::192.168.0.13/126

[edit]
lab@R5# commit
commit complete
```

Now you must enable the appropriate family for the existing IPv4 session.

```
[edit]
lab@R5# set protocols bgp group C3 family inet6 unicast

[edit]
lab@R5# commit
commit complete
```

Verify the BGP session table.

```
[edit]
lab@R5# run show bgp summary
Groups: 3 Peers: 4 Down peers: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       243        242          0          0          0          0
inet6.0
                        32         32          0          0          0          0
Peer              AS  InPkt  OutPkt  OutQ  Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.4    65001    536     359     0      0  2:04:24 211/211/211/0
0/0/0/0
172.30.5.6    65002    278     357     0      0  2:03:42 31/32/32/0
0/0/0/0
192.168.0.10  64514     11      11     0      0     3:59 Establ
  inet6.0: 16/16/16/0
192.168.0.14  64514     11      10     0      0     3:55 Establ
  inet6.0: 16/16/16/0
```

IPv6 routes are received from both neighbors in AS 64514. But looking carefully you can see that IPv4 routes are missing. List the routes received from those neighbors:

```
[edit]
lab@R5# run show route receive-protocol bgp 192.168.0.10

inet.0: 316 destinations, 316 routes (315 active, 0 holddown, 1 hidden)

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

inet6.0: 55 destinations, 75 routes (55 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  fd18:cccc:dddd:2::/80
*                         ::192.168.0.10                         64514 I
  fd18:cccc:dddd:2:1::/80
*                         ::192.168.0.10                         64514 I
  fd18:cccc:dddd:2:2::/80
*                         ::192.168.0.10                         64514 I
  fd18:cccc:dddd:2:3::/80
*                         ::192.168.0.10                         64514 I
  fd18:cccc:dddd:2:4::/80
*                         ::192.168.0.10                         64514 I
  fd18:cccc:dddd:2:5::/80
*                         ::192.168.0.10                         64514 I
  fd18:cccc:dddd:2:6::/80
*                         ::192.168.0.10                         64514 I
...
```

Indeed, only IPv6 routes are received, the IPv4 prefixes are gone. Verify the detailed information for the connection between R5 and C3-1

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# run show bgp neighbor 192.168.0.10
Peer: 192.168.0.10+179 AS 64514 Local: 192.168.0.9+61244 AS 54591
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LogUpDown AddressFamily PeerAS Multipath Refresh>
  Address families configured: inet6-unicast
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 172.31.47.1     Local ID: 172.30.5.5        Active Holdtime: 90
  Keepalive Interval: 30         Peer index: 0
  BFD: disabled, down
  Local Interface: ge-0/0/5.303
  NLRI for restart configured on peer: inet6-unicast
  NLRI advertised by peer: inet-unicast inet6-unicast
  NLRI for this session: inet6-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet6-unicast
  NLRI of received end-of-rib markers: inet6-unicast
  NLRI of all end-of-rib markers sent: inet6-unicast
  Peer supports 4 byte AS extension (peer-as 64514)
  Peer does not support Addpath
  Table inet6.0 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              16
    Received prefixes:            16
    Accepted prefixes:            16
    Suppressed due to damping:    0
    Advertised prefixes:          0
  Last traffic (seconds): Received 1    Sent 23   Checked 18
  Input messages:  Total 7      Updates 2       Refreshes 0     Octets 378
  Output messages: Total 6      Updates 0       Refreshes 0     Octets 184
  Output Queue[1]: 0
```

From the output you can see the local peer R5 is configured only for family `inet6 unicast`. The remote peer C3-1 advertises both families `inet unicast` and `inet6 unicast`. Enable `family inet6 unicast` on R5, the default `family inet-unicast` is disabled. You must explicitly enable it again to have both families active over a single BGP session.

```
[edit]
lab@R5# set protocols bgp group C3 family inet unicast

[edit]
lab@R5# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

Verify if the IPv4 routes are now received together with IPv6 prefixes.

```
[edit]
lab@R5# run show bgp summary
Groups: 3 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     275        274          0          0          0          0
inet6.0
                      32         32          0          0          0          0
Peer                AS  InPkt  OutPkt  OutQ  Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.4      65001    661     483     0      0   2:59:42 211/211/211/0
0/0/0/0
172.30.5.6      65002    400     482     0      0   2:59:00 31/32/32/0
0/0/0/0
192.168.0.10    64514     82     252     0      0     35:54 Establ
  inet.0: 16/16/16/0
  inet6.0: 16/16/16/0
192.168.0.14    64514     82     243     0      0     35:50 Establ
  inet.0: 16/16/16/0
  inet6.0: 16/16/16/0
```

c) R7.

R7 should establish an IPv6 EBGP peer with P1-2.

The IPv6 address is configured on the interface towards P1-2.

```
[edit]
lab@R7# set interfaces ge-0/0/5.308 family inet6 address fc09:c0:ffee::1/126

[edit]
lab@R7# set protocols bgp group P1-2-ipv6 peer-as 1679.12483

[edit]
lab@R7# set protocols bgp group P1-2-ipv6 neighbor fc09:c0:ffee::2

[edit]
lab@R7# set protocols bgp group P1-2-ipv6 log-updown

[edit]
lab@R7# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

After committing the configuration, verify the BGP summary table.

```
[edit]
lab@R7# run show bgp summary
Groups: 5 Peers: 5 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      1372        791          0          0          0          0
inet6.0
                        16         16          0          0          0          0
Peer                     AS  InPkt  OutPkt  OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.6           65002    340    1742     0       0   2:32:01 15/32/32/0
0/0/0/0
172.30.5.8           65003   2051     841     0       1   2:25:38 652/760/760/0
0/0/0/0
192.168.0.30    110047427    824     979     0       0   2:32:23 108/564/564/0
0/0/0/0
192.168.0.34         64512    333    1364     0       0   2:32:26 16/16/16/0
0/0/0/0
fc09:c0:ffee::2  110047427    6       6       0       0     1:32 Establ
  inet6.0: 16/16/16/0
```

R7 receives 16 IPv6 routes from P1-2.

d) R8.

As with R7, R8 should establish an IPv6 BGP session with P1-1.

Configure the IPv6 address and IPv6 EBGP peer.

```
[edit]
lab@R8# set interfaces ge-0/0/5.310 family inet6 address fc09:c0:ffee::5/126

[edit]
lab@R8# set protocols bgp group P1-1-ipv6 peer-as 1679.12483

[edit]
lab@R8# set protocols bgp group P1-1-ipv6 neighbor fc09:c0:ffee::6

[edit]
lab@R8# set protocols bgp group P1-1-ipv6 log-updown

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify the BGP session table.

```
[edit]
lab@R8# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      1357        791          0          0          0          0
inet6.0
                        16         16          0          0          0          0
Peer                     AS  InPkt  OutPkt  OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.1            65000   1508     787     0       0   2:26:27 652/654/654/0
0/0/0/0
172.30.5.7            65003    835     787     0       0   2:26:14 31/139/139/0
0/0/0/0
192.168.0.38  110047427    808     404     0       0   2:26:38 108/564/564/0
0/0/0/0
fc09:c0:ffee::6   110047427      3       3     0     0          25 Establ
  inet6.0: 16/16/16/0
```

R8 receives 16 IPv6 prefixes from the IPv6 EBGP session to P1-1.

9) Enable route flap damping on R5, R6 and R7

You have to configure damping in case of flapping prefixes received from C1 and C3 customers. The prefixes from C1 should be damped more aggressive.

Damping is measure used by BGP to counteract the effect of flapping prefixes/routes. Route damping is used only on EBGP session. Damping means that BGP ignores prefixes that experienced flaps. BGP stops using these prefixes for best path selection and stops advertising them.

With damping enabled, BGP builds a database of instable prefixes received from neighbors. This means that BGP sets a number to each prefix and when this number is reached, the prefix is damped for some time period. This number is called figure of merit and by default is 0. Each time the prefix is withdrawn, the figure is increased with 1000, and so on.

Figure of merit depends on four parameters:

- ***Suppress*** – this is the threshold where the BGP prefix is considered unstable and is not used any more. By default, it is 3000.
- ***Reuse*** – this is the threshold where the BGP prefix is used again for best path calculation and eligible for advertisement. By default, it is 750.

© 2020 Juniper Networks, Inc. All rights reserved.

- ***Half-life*** – this is a value in minutes and represents the rate at which the figure is decreased by half if the value is greater than 0. Default value is 15 minutes.
- ***Max-suppress*** – this value is maximum time that a prefix can stay damped. By default, it is 60 minutes.

a) R5.

On R5, for both C3 customer peers you must enable default damping parameters.

```
[edit]
lab@R5# set protocols bgp group C3 damping

[edit]
lab@R5# commit
commit complete
```

Adding `damping` under the BGP group enables damping for all routes received from the neighbors in that group.

b) R6.

```
[edit]
lab@R6# set protocols bgp group C1-1 damping
```

The `damping` configuration should be enabled globally but as per requirement here you have to make damping parameters more aggressive. This is done with damping policy.

```
[edit]
lab@R6# set policy-options damping aggressive half-life 20

[edit]
lab@R6# set policy-options damping aggressive reuse 500

[edit]
lab@R6# set policy-options damping aggressive suppress 2500

[edit]
lab@R6# set policy-options policy-statement damp-aggressive then damping aggressive
```

Increasing the half-life to 20 minutes means that the rate of lowering the figure will decrease. The `Suppress` parameter is lowered to 2500 so that the prefixes will be in the damped state sooner. Lowering the `reuse` parameter to 500 will increase the time for the prefix to be considered again by BGP.

© 2020 Juniper Networks, Inc. All rights reserved.

After creating damping policy, you then have to use that policy in routing policy. Routing policy is applied in import direction to the peer for which prefixes should be damped more aggressively.

```
[edit]
lab@R6# set protocols bgp group C1-1 import damp-aggressive

[edit]
lab@R6# commit
commit complete
```

Policy is applied to the C1-1 peer only.

c) R7.

R7 has session with C1-1 as well. The same damping and routing policies must be configured.

```
[edit]
lab@R7# set protocols bgp group C1-1 damping

[edit]
lab@R7# set policy-options damping aggressive half-life 20

[edit]
lab@R7# set policy-options damping aggressive reuse 500

[edit]
lab@R7# set policy-options damping aggressive suppress 2500

[edit]
lab@R7# set policy-options policy-statement damp-aggressive term 1 then damping
aggressive

[edit]
lab@R7# set protocols bgp group C1-1 import damp-aggressive

[edit]
lab@R7# commit
commit complete
```

Note: With `show route damping` you can see the damped route in the routing table.

10) Ensuring that all routes from all EBGP peers are reachable on all router's routing table.

© 2020 Juniper Networks, Inc. All rights reserved.

EBGP prefixes are re-advertised to IBGP neighbors. However, IBGP sessions do not change the next-hop for the received EBGP prefixes, this is also true for CBGP sessions, those prefixes are hidden as the next-hop cannot be resolved.

The following solutions can be used to resolve next-hop issue.

- ***Next-hop self*** – using policy you can change the next-hop, before advertising the route to IBGP, to IP address of the advertising router. This IP address is known to the other internal routers (IGP usually). This is the preferred way and is the most used.
- ***Export direct routes to IGP*** – you can export the direct link IP addresses to IGP making them reachable.
- ***IGP passive interface*** – this method is almost the same as previous one. Difference here is that the routes are added to IGP as Internal not like External.
- ***Static routes*** – this method does not scale well but could solve the problem. You can create static route on each IBGP router to the EBGP next-hop.
- ***IGP adjacencies between AS network*** – this method is not preferred because of several reasons. But as the previous one it could solve the problem. Forming IGP adjacencies between both ASs means to stretch the IGP to the other AS. Choosing this option in a live network is an absolute guarantee to get fired!

a) R1

R1 and R2 should not use next-hop self for resolving the problem as per requirement. Here you are forced to use the IGP passive method. You have to add the broadcast interface towards IX to IS-IS Level 1.

Currently R1 has no hidden routes.

Note: If your device prefers the route learned from R8 then you will have a single hidden route (192.0.0.0/4) which will be outlined after the next few outputs.

```
[edit]
lab@R1# run show route hidden

inet.0: 875 destinations, 2004 routes (875 active, 0 holddown, 0 hidden)

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

inet6.0: 31 destinations, 32 routes (31 active, 0 holddown, 0 hidden)
```

R1 receives routes from IX routers, R8 and R2.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show bgp summary
Groups: 3 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    1921        791          0          0          0          0
Peer             AS  InPkt  OutPkt  OutQ  Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.2   65000   1409    1325     0      0    1:07:45 203/654/654/0
0/0/0/0
172.30.5.8   65003   1805    1607     0      1      56:27 137/139/139/0
0/0/0/0
192.168.1.3   1620   1113     498     0      1    1:04:37 451/564/564/0
0/0/0/0
192.168.1.4   1620   1113     471     0      1    1:04:34 0/564/564/0
0/0/0/0
```

Only the routes from R8 can have unresolved next-hop, because the next-hop for the prefixes from other the neighbors are known. The prefixes received from R8 have a next-hop of 192.168.0.38, which is resolved by 192.0.0.0/4, received from IX routers.

```
[edit]
lab@R1# run show route 192.168.0.2

inet.0: 875 destinations, 1963 routes (827 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.0.0/4         *[BGP/170] 00:23:11, localpref 100
                      AS path: 1620 33003 62186 23280 38961 3686 I
                     > to 192.168.1.3 via ge-0/0/5.300
                     [BGP/170] 00:23:08, localpref 100
                      AS path: 1620 33003 62186 23280 38961 3686 I
                     > to 192.168.1.4 via ge-0/0/5.300
                     [BGP/170] 00:23:50, localpref 100, from 172.30.5.2
                      AS path: 1620 33003 62186 23280 38961 3686 I
                     > to 192.168.1.4 via ge-0/0/5.300
```

Adding the broadcast interface in IS-IS as passive will resolve the next-hop issue on other routers in the network.

```
[edit]
lab@R1# set protocols isis interface ge-0/0/5.300 level 1 passive

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify the IS-IS interface table.

```
[edit]
lab@R1# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR              Level 2 DR              L1/L2 Metric
ae0.0                 1   0x1 Point to Point          Disabled                       5/5
ge-0/0/4.114          1   0x1 Point to Point          Disabled                     10/10
ge-0/0/4.118          1   0x1 Point to Point          Disabled                     10/10
ge-0/0/5.300          1   0x1 Passive                 Disabled                     10/10
lo0.0                 1   0x1 Passive                 Passive                       0/0
```

Verify the Level 1 LSP interface route is advertised to IS-IS.

```
[edit]
lab@R1# run show isis database level 1 R1.00-00 extensive
IS-IS level 1 link-state database:

R1.00-00 Sequence: 0x2f, Checksum: 0xf4d7, Lifetime: 1154 secs
   IPV4 Unicast IS neighbor: R2.00        Metric:         5
     Two-way fragment: R2.00-00, Two-way first fragment: R2.00-00
   IPV4 Unicast IS neighbor: R4.00      Metric:        10
     Two-way fragment: R4.00-00, Two-way first fragment: R4.00-00
   IPV4 Unicast IS neighbor: R8.00       Metric:        10
     Two-way fragment: R8.00-00, Two-way first fragment: R8.00-00
   IPV6 Unicast IS neighbor: R2.00        Metric:         5
     Two-way fragment: R2.00-00, Two-way first fragment: R2.00-00
   IPV6 Unicast IS neighbor: R8.00       Metric:        10
     Two-way fragment: R8.00-00, Two-way first fragment: R8.00-00
   IP IPV4 Unicast prefix: 172.30.0.0/30  Metric:         5 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.4/30  Metric:        10 Internal Up
   IP IPV4 Unicast prefix: 172.30.0.8/30  Metric:        10 Internal Up
   IP IPV4 Unicast prefix: 172.30.5.1/32  Metric:         0 Internal Up
   IP IPV4 Unicast prefix:  192.168.1.0/24 Metric:        10 Internal Up
   V6 IPV6 Unicast prefix: fd17:f0f4:f691:5::1/128 Metric:         0 Internal Up
...
```

b) R2

R2 configuration should match the configuration on R1. The interface route should be added as passive to IS-IS.

```
[edit]
lab@R2# set protocols isis interface ge-0/0/5.300 level 1 passive

[edit]
lab@R2# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

c) R3

R3 has two EBGP sessions with P2-1 and P3-1 and receives combined more than 200 prefixes.

```
[edit]
lab@R3# run show bgp summary
Groups: 5 Peers: 5 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       848        227          0          0          0          0
inet6.0
                        16         16          0          0          0          0
Peer                     AS  InPkt  OutPkt  OutQ  Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.2            65000   1647     901     0      0  3:31:51 13/601/601/0
0/0/0/0
172.30.5.4            65001    484    1430     0      0  3:30:23 3/34/34/0
0/0/0/0
192.168.0.2  3521382357    272     298     0      1  1:28:47 96/97/97/0
0/0/0/0
192.168.0.6  2831679853    540     634     0      0  3:29:52 115/116/116/0
0/0/0/0
fe80::223:9c01:2d8b:6c81  3521382357   150    153   0    0  1:08:35 Establ
  inet6.0: 16/16/16/0
```

Verify how R3 is advertising the prefixes to R4.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.4

inet.0: 879 destinations, 889 routes (879 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/4               192.168.1.3                  100        (65000) 1620
61671 I
* 0.0.0.0/5               192.168.1.3                  100        (65000) 1620
61671 27075 I
* 1.64.0.0/10             192.168.1.3                  100        (65000) 1620
61671 I
* 1.84.160.0/20           192.168.1.3                  100        (65000) 1620
33112 I
* 1.96.0.0/11             192.168.1.3                  100        (65000) 1620
33112 63164 40776 51777 I
* 1.161.192.0/21          192.168.1.3                  100        (65000) 1620
33112 30404 32138 45045 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 1.166.89.0/24             192.168.0.38                   100        (65000 65003)
110047427 56422 47123 34908 57085 3090 19128 21300 I
* 1.176.0.0/12              192.168.1.3                    100        (65000) 1620
33112 49129 16320 52954 I
* 2.0.0.0/7                 192.168.1.3                    100        (65000) 1620
61671 60177 3091 5721 1770 35283 6098 41407 ?
* 3.49.128.0/19             192.168.0.38                   100        (65000 65003)
110047427 56422 7923 42335 17359 9293 38097 38522 61520 I
* 3.189.0.0/16              192.168.1.3                    100        (65000) 1620
33112 41164 7136 6370 35389 39823 16965 I
* 3.214.25.0/24             192.168.1.3                    100        (65000) 1620
61671 62186 11282 7796 I
* 4.0.0.0/6                 192.168.0.38                   100        (65000 65003)
110047427 56422 235 4628 20905 9158 9476 23720 14923 ?
* 4.16.128.0/17             192.168.1.3                    100        (65000) 1620
33003 49129 50466 22491 57331 43943 48281 I
* 4.53.208.0/21             192.168.0.38                   100        (65000 65003)
110047427 36024 50292 3128 3431 18053 I
* 4.84.128.0/18             192.168.1.3                    100        (65000) 1620
33112 63164 28164 56111 39848 I
* 4.173.41.0/24             192.168.1.3                    100        (65000) 1620
33003 30404 22363 54666 40669 53365 I
* 5.127.0.0/17              192.168.0.6                    100        2831679853 9726
36659 30705 25538 37414 49276 ?
* 6.83.176.0/21             192.168.1.3                    100        (65000) 1620
33003 63164 14519 59474 20051 I
* 6.98.154.0/25             192.168.0.6                    100        2831679853
26697 36659 39905 I
* 7.0.0.0/8                 192.168.1.3                    100        (65000) 1620
33112 52411 37841 5705 22477 16750 ?
```

You can see that R3 is advertising the routes to R4 with their original next-hop.

```
[edit]
lab@R3# set policy-options policy-statement nhs term 1 from protocol bgp

[edit]
lab@R3# set policy-options policy-statement nhs term 1 from route-type external

[edit]
lab@R3# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R3# set protocols bgp group ibgp export nhs

[edit]
lab@R3# set protocols bgp group cbgp export nhs
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# commit
commit complete
```

The routing-policy matches external BGP routes and has single then statement, which is changing the next-hop to self. Then the policy is applied in the export direction to the IBGP sessions.

Verify the next hop advertised by R3 after the policy.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.4

inet.0: 879 destinations, 889 routes (879 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/4               192.168.1.3                  100        (65000) 1620
61671 I
* 0.0.0.0/5               192.168.1.3                  100        (65000) 1620
61671 27075 I
* 1.64.0.0/10             192.168.1.3                  100        (65000) 1620
61671 I
* 1.84.160.0/20           192.168.1.3                  100        (65000) 1620
33112 I
* 1.96.0.0/11             192.168.1.3                  100        (65000) 1620
33112 63164 40776 51777 I
* 1.161.192.0/21          192.168.1.3                  100        (65000) 1620
33112 30404 32138 45045 I
* 1.166.89.0/24           192.168.0.38                 100        (65000 65003)
110047427 56422 47123 34908 57085 3090 19128 21300 I
* 1.176.0.0/12            192.168.1.3                  100        (65000) 1620
33112 49129 16320 52954 I
* 2.0.0.0/7               192.168.1.3                  100        (65000) 1620
61671 60177 3091 5721 1770 35283 6098 41407 ?
* 3.49.128.0/19           192.168.0.38                 100        (65000 65003)
110047427 56422 7923 42335 17359 9293 38097 38522 61520 I
* 3.189.0.0/16            192.168.1.3                  100        (65000) 1620
33112 41164 7136 6370 35389 39823 16965 I
* 3.214.25.0/24           192.168.1.3                  100        (65000) 1620
61671 62186 11282 7796 I
* 4.0.0.0/6               192.168.0.38                 100        (65000 65003)
110047427 56422 235 4628 20905 9158 9476 23720 14923 ?
* 4.16.128.0/17           192.168.1.3                  100        (65000) 1620
33003 49129 50466 22491 57331 43943 48281 I
* 4.53.208.0/21           192.168.0.38                 100        (65000 65003)
110047427 36024 50292 3128 3431 18053 I
* 4.84.128.0/18           192.168.1.3                  100        (65000) 1620
33112 63164 28164 56111 39848 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 4.173.41.0/24              192.168.1.3                    100        (65000) 1620
33003 30404 22363 54666 40669 53365 I
* 5.127.0.0/17               Self                           100        2831679853 9726
36659 30705 25538 37414 49276 ?
* 6.83.176.0/21              192.168.1.3                    100        (65000) 1620
33003 63164 14519 59474 20051 I
* 6.98.154.0/25              Self                           100        2831679853
26697 36659 39905 I
```

d) R4.

R4 has no EBGP sessions, but it is a good candidate to explore the next-hop issue.

After the next-hop self-policy is applied on R3, the prefixes received from R3 are having a reachable next hop, which is the IP address of the link between R3 and R4.

```
[edit]
lab@R4# run show route terse

inet.0: 879 destinations, 894 routes (878 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination          P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 0.0.0.0/4            B 170        100                           (65000) 1620
61671 I
  unverified                                        >172.30.0.5
* ? 0.0.0.0/5            B 170        100                           (65000) 1620
61671 27075 I
  unverified                                        >172.30.0.5
* ? 1.64.0.0/10          B 170        100                           (65000) 1620
61671 I
  unverified                                        >172.30.0.5
* ? 1.84.160.0/20        B 170        100                           (65000) 1620
33112 I
  unverified                                        >172.30.0.5
* ? 1.96.0.0/11          B 170        100                           (65000) 1620
33112 63164 40776 51777 I
  unverified                                        >172.30.0.5
* ? 1.161.192.0/21       B 170        100                           (65000) 1620
33112 30404 32138 45045 I
  unverified                                        >172.30.0.5
* ? 1.166.89.0/24        B 170        100                           (65000 65003)
110047427 56422 47123 34908 57085 3090 19128 21300 I
  unverified                                        >172.30.0.5
* ? 1.176.0.0/12         B 170        100                           (65000) 1620
33112 49129 16320 52954 I
  unverified                                        >172.30.0.5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* ? 2.0.0.0/7          B 170        100                              (65000) 1620
61671 60177 3091 5721 1770 35283 6098 41407 ?
  unverified                                        >172.30.0.5
* ? 3.49.128.0/19      B 170        100                              (65000 65003)
110047427 56422 7923 42335 17359 9293 38097 38522 61520 I
  unverified                                        >172.30.0.5
* ? 3.189.0.0/16       B 170        100                              (65000) 1620
33112 41164 7136 6370 35389 39823 16965 I
  unverified                                        >172.30.0.5
* ? 3.214.25.0/24      B 170        100                              (65000) 1620
61671 62186 11282 7796 I
  unverified                                        >172.30.0.5
* ? 4.0.0.0/6          B 170        100                              (65000 65003)
110047427 56422 235 4628 20905 9158 9476 23720 14923 ?
  unverified                                        >172.30.0.5
* ? 4.16.128.0/17      B 170        100                              (65000) 1620
33003 49129 50466 22491 57331 43943 48281 I
  unverified                                        >172.30.0.5
* ? 4.53.208.0/21      B 170        100                              (65000 65003)
110047427 36024 50292 3128 3431 18053 I
  unverified                                        >172.30.0.5
* ? 4.84.128.0/18      B 170        100                              (65000) 1620
33112 63164 28164 56111 39848 I
  unverified                                        >172.30.0.5
* ? 4.173.41.0/24      B 170        100                              (65000) 1620
33003 30404 22363 54666 40669 53365 I
  unverified                                        >172.30.0.5
* ? 5.127.0.0/17       B 170        100                              2831679853
9726 36659 30705 25538 37414 49276 ?
  unverified                                        >172.30.0.21
...
```

Review the hidden route that appears on R4.

```
[edit]
lab@R4# run show route hidden

inet.0: 879 destinations, 894 routes (878 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0              [Aggregate] 00:06:36
                         Reject

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

inet6.0: 33 destinations, 33 routes (33 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show route hidden extensive

inet.0: 879 destinations, 894 routes (878 active, 0 holddown, 1 hidden)
0.0.0.0/0 (1 entry, 0 announced)
         Aggregate
                Next hop type: Reject, Next hop index: 0
                Address: 0xbe73d90
                Next-hop reference count: 1
                State: <Hidden Int Ext Changed>
                Local AS: 65001
                Age: 6:43
                Validation State: unverified
                Task: Aggregate
                AS path: ({65000 65002 65003}) {73 235 294 339 391 522 583 783
794 1039 1386 1435 1540 1620 1790 1900 2249 2442 2455 2621 2732 2840 2911 2964
3073 3079 3093 3205 3858 3907 4098 4126 4199 4322 4341 4364 4377 4432 4440 4473
4610 4637 4669 4724 4757 5144 5215 5572 5646 5697 6221 6363 6452 6621 6736 6803
6918 7109 7211 7291 7318 7408 7434 7478 7553 7796 7809 7810 7864 7923 8149 8250
8283 8392 8471 8532 8975 9166 9290 9573 9726 10129 10213 10512 10588 10670 10677
11015 11030 11189 11282 11735 11767 11824 11895 11964 11967 11989 11996 12070
12096 12254 12277 12546 12558 12603 12628 12630 12631 12833 12925 13064 13379
13859 13924 14052 14519 14529 14769 14967 15134 15270 15311 15324 15417 15595
15641 15698 15729 15751 15757 15804 16003 16076 16083 16214 16320 16367 16415
16474 16486 16644 16731 16804 16881 17027 17311 17359 17485 17519 17596 17752
17862 17952 18088 18133 18261 18421 18424 18566 18632 18695 18823 18856 18938
19022 19045 19115 19412 19579 19604 19626 19797 20051 20058 20201 20308 20922
21086 21422 21432 21847 21995 21997 22156 22315 22447 22508 22770 22896 22927
23011 23071 23369 23590 23618 23834 23946 23973 23980 24130 24293 24462 24527
24614 24783 24807 24825 24934 24984 25129 25291 25429 25505 25561 25728 25733
25782 26184 26290 26374 26697 27023 27075 27233 27260 27338 27622 27725 27765
27823 27841 27917 28069 28074 28095 28147 28164 28274 28324 28389 28392 28421
28432 28440 28515 28632 28850 28964 29201 29318 29414 29537 29886 30006} {30198
30307 30394 30404 30690 30709 30903 30990 30993 31395 31460 31825 31903 32041
32138 32218 32308 32310 32499 32692 32814 32856 32895 33003 33112 33163 33170
33806 33879 34167 34264 34427 34454 34502 34511 34519 34708 34877 34880 35154
35412 35580 35713 35738 35761 35823 35848 35851 36024 36263 36570 36621 36773
37067 37159 37678 37730 37739 37810 38202 38235 38353 38445 38568 38772 38966
39156 39235 39260 39310 39399 39595 39596 39848 39993 40234 40487 40522 40612
40776 40827 40946 40964 41116 41164 41240 41479 41551 41558 41729 41830 41865
41968 42490 42693 43174 43266 43790 44008 44051 44308 44354 44372 44473 44541
44977 45020 45045 45456 45461 45561 45579 45580 45703 45755 45948 46182 46288
46848 46859 46924 47109 47123 47191 47547 47650 47734 47832 47957 48013 48413
48485 48723 48750 48948 49129 49182 49255 49429 49911 50097 50102 50192 50234
50292 50767 50874 50996 51066 51176 51216 51234 51297 51492 51521 51777 52184
52411 52527 52702 52765 52954 53145 53146 53304 53324 53454 53536 53646 53716
53943 53963 54085 54246 54283 54870 54904 55090 55406 55461 55548 55612 56068
56075 56111 56153 56173 56251 56342 56355 56422 56647 56813 57017 57359 57375
57545 57763 57816 57983 58079 58105 58230 58524 58642 58669 58735 58800 58919
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
59056 59261 59474 59544 60050 60096 60177 60197 60224 60236 60239 60467 60637
60715 60930 61080 61219 61279 61320 61361 61370 61444 61671 61718 61968 62006
62164 62186 62239 62283 62295 62341 62703 62843 63075 63155 63160 63164 63861
64005 64234} ?  (LocalAgg OverflowASPathSize)
...
```

The reason why this route is hidden has to do with all the AS path entries that get included from all the contributing BGP routes. In this case, there are too many and it is hitting the max size.

```
[edit]
lab@R4# set routing-options aggregate defaults brief

[edit]
lab@R4# commit
commit complete
```

Verify the changes have fixed the hidden route on R4.

```
[edit]
lab@R4# run show route hidden

inet.0: 879 destinations, 894 routes (879 active, 0 holddown, 0 hidden)

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

inet6.0: 33 destinations, 33 routes (33 active, 0 holddown, 0 hidden)...
```

The aggregate route is no longer hidden. You will need to fix this route on R5 after verifying all the other information.

e) R5

R5 has two external BGP neighbors, both from AS 64514. Verify BGP session table on R5.

```
[edit]
lab@R5# run show bgp summary
Groups: 3 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     275        274          0          0          0          0
inet6.0
                      32         32          0          0          0          0
Peer             AS  InPkt  OutPkt  OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.4      65001     661      483      0        0    2:59:42 211/211/211/0
0/0/0/0
172.30.5.6      65002     400      482      0        0    2:59:00 31/32/32/0
0/0/0/0
192.168.0.10    64514      82      252      0        0      35:54 Establ
  inet.0: 16/16/16/0
  inet6.0: 16/16/16/0
192.168.0.14    64514      82      243      0        0      35:50 Establ
  inet.0: 16/16/16/0
  inet6.0: 16/16/16/0
```

Apply the next-hop self-policy on R5 IBGP and CBGP session.

```
[edit]
lab@R5# set policy-options policy-statement nhs term 1 from protocol bgp

[edit]
lab@R5# set policy-options policy-statement nhs term 1 from route-type external

[edit]
lab@R5# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R5# set protocols bgp group ibgp export nhs

[edit]
lab@R5# set protocols bgp group cbgp export nhs

[edit]
lab@R5# commit
commit complete
```

Review the 0/0 hidden route that appears on R5.

```
[edit]
lab@R5# run show route hidden 0/0 exact extensive

inet.0: 882 destinations, 914 routes (880 active, 0 holddown, 2 hidden)
0.0.0.0/0 (1 entry, 0 announced)
         Aggregate
                Next hop type: Reject, Next hop index: 0
                Address: 0xbe73d90
                Next-hop reference count: 1
                State: <Hidden Int Ext Changed>
                Local AS: 65002
                Age: 18:39
                Validation State: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                Task: Aggregate
                AS path: ({65000 65001 65003}) {73 82 105 221 235 362 391 448 647
1039 1044 1158 1435 1449 1498 1527 1556 1620 1810 1900 1938 2175 2455 2692 2772
2840 2911 2942 2964 2972 3002 3676 3913 3988 4098 4322 4341 4377 4450 4610 4637
4669 4685 4701 4850 5040 5144 5283 5376 5405 5463 5577 5697 5934 6047 6122 6221
6292 6363 6417 6500 6621 6738 6749 6803 6854 7090 7247 7408 7418 7458 7478 7546
7553 7923 8124 8142 8149 8162 8199 8471 8532 8739 8773 8858 8897 8975 9282 9348
9583 9726 9806 9830 10038 10275 10277 10512 10564 10573 10591 10856 11015 11030
11041 11114 11147 11163 11189 11211 11243 11678 11842 11931 11967 11993 11996
12180 12432 12631 12799 12943 12981 13600 13839 13859 13908 13924 13970 14053
14321 14421 14529 14615 14669 14769 14822 14964 15085 15134 15270 15287 15311
15370 15406 15417 15698 15751 15773 15816 15988 16104 16320 16353 16415 16490
16572 16881 16950 17104 17112 17205 17363 17650 17752 17944 17954 18566 18632
18695 18720 18725 18938 19045 19086 19115 19271 19379 19404 19412 19447 19516
19579 19593 19787 19797 19856 20058 20308 20438 20457 20738 20756 21121 21137
21151 21356 21422 21432 21997 22085 22156 22499 22508 22622 22685 22695 22770
22793 22825 22856 22889 22896 22971 23268 23369 23980 24074 24130 24148 24249
24563 24573 24686 24807 24878 25129 25214 25366 25505 25709 25728 25782 25972
26116 26184 26192 26223 26472 26586 26593 26697 26793 26964 27013 27075 27127
27134 27233 27365 27452 27550 27725 27841 28074 28213} {28513 28515 28552 28807
29191 29197 29410 29537 29611 29982 30212 30394 30404 30962 31133 31155 31395
31460 31532 31606 31785 31825 32138 32218 32308 32523 32814 33003 33112 33207
33320 33332 33751 33786 33806 33909 34040 34070 34427 34440 34454 34476 34511
34519 34562 34880 35069 35248 35488 35560 35580 35621 35713 35728 35761 36024
36176 36570 36621 36659 36805 36832 36854 36956 37067 37080 37352 37730 37739
38425 38445 38523 38845 39136 39310 39339 39587 39596 39819 40234 40494 40612
40715 40751 40776 40827 40946 40964 40969 40979 41116 41164 41533 41551 41726
41830 41968 42039 42110 42216 42693 43067 43096 43174 43481 43545 43727 43790
43964 44008 44298 44354 44544 44678 44771 44929 45045 45273 45303 45383 45561
45730 45796 45948 46305 46312 46848 46859 46923 47078 47082 47123 47209 47734
47736 47780 48413 48485 48591 48750 48898 48948 49044 49060 49129 49179 49182
49383 49542 49557 49911 49930 50097 50192 50234 50274 50292 50293 50419 50476
50563 50666 50874 50887 50889 51020 51068 51216 51297 51417 51434 51521 51614
51777 51805 52229 52411 52461 52954 52989 53239 53355 53598 53943 54232 54253
54284 54491 54617 54928 55493 55548 55567 55803 56041 56173 56251 56321 56342
56422 56512 56813 56894 57359 57816 57819 57862 57966 57968 57983 58079 58384
58955 59056 59071 59521 59544 59878 59959 60071 60177 60236 60637 60715 60886
60930 61080 61361 61405 61444 61649 61671 62064 62186 62504 62582 62703 62822
63075 63155 63164 63202 63280 63284 63300 63669 63776 63827 63887 64005} ?
(LocalAgg OverflowASPathSize)
...
```

The reason why this route is hidden is the same reason the route was hidden on R4.

```
[edit]
lab@R5# set routing-options aggregate defaults brief

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# commit
commit complete
```

Verify the changes have fixed the hidden route on R4.

```
[edit]
lab@R5# run show route hidden extensive

inet.0: 882 destinations, 914 routes (881 active, 0 holddown, 1 hidden)
172.31.31.0/24 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Next hop type: Unusable, Next hop index: 0
                Address: 0xbe73e50
                Next-hop reference count: 1
                State: <Hidden Int Ext>
                Local AS: 65002 Peer AS: 65002
                Age: 22
                Validation State: unverified
                Task: BGP_65002.172.30.5.6
                AS path: 64513 I
                Accepted
                Localpref: 100
                Router ID: 172.30.5.6
                Indirect next hops: 1
                        Protocol next hop: 172.31.31.1
                        Indirect next hop: 0x0 - INH Session ID: 0x0

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

inet6.0: 55 destinations, 71 routes (55 active, 0 holddown, 0 hidden)
```

The aggregate 0/0 route is no longer hidden. But there is still a hidden route. The reason why this route is hidden is because the next-hop is part of the same subnet. The prefix cannot be resolved by a next-hop, which is part of that same subnet. The prefix is resolved by the C2-1 loopback IP address. The problem will be fixed when we apply the next-hop self-policy on R6.

f) R6

R6 has two EBGP neighbors. It has a lot of hidden prefixes.

```
[edit]
lab@R6# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    930        808          0          0          0          0
Peer                     AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.5            65002       1727       1605       0       0       19:37
669/775/775/0         0/0/0/0
172.30.5.7            65003       1605       1652       0       0       19:46
107/123/123/0         0/0/0/0
172.31.31.1           64513         46        675       0       0       19:46
16/16/16/0            0/0/0/0
192.168.0.18          64512         47        672       0       0       19:59
16/16/16/0            0/0/0/0
```

Adding the next-hop self-policy on R6.

```
[edit]
lab@R6# set policy-options policy-statement nhs term 1 from protocol bgp

[edit]
lab@R6# set policy-options policy-statement nhs term 1 from route-type external

[edit]
lab@R6# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R6# set protocols bgp group ibgp export nhs

[edit]
lab@R6# set protocols bgp group cbgp export nhs

[edit]
lab@R6# commit
commit complete
```

g) R7

R7 has an EBGP session to one customer and one peer router in different ASs.

```
[edit]
lab@R7# run show bgp summary
Groups: 5 Peers: 5 Down peers: 1
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      1995        808          0          0          0          0
inet6.0
                        16         16          0          0          0          0
Peer                     AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.6              65002       1655       1607        0       1       21:01
685/701/701/0           0/0/0/0
172.30.5.8              65003       2256       2656        0       0       22:40
0/714/714/0             0/0/0/0
192.168.0.30         110047427        507        607        0       0       23:27
107/564/564/0           0/0/0/0
192.168.0.34            64512         55       1067        0       0       23:29
16/16/16/0              0/0/0/0
fc09:c0:ffee::2      110047427         52         52        0       0       23:20
Establ
  inet6.0: 16/16/16/0
```

Apply the next-hop self-policy on R7 IBGP and CBGP session.

```
[edit]
lab@R7# set policy-options policy-statement nhs term 1 from protocol bgp

[edit]
lab@R7# set policy-options policy-statement nhs term 1 from route-type external

[edit]
lab@R7# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R7# set protocols bgp group ibgp export nhs

[edit]
lab@R7# set protocols bgp group cbgp export nhs

[edit]
lab@R7# commit
commit complete
```

h) R8

R8 has single EBGP session with P1-1 router.

```
[edit]
lab@R8# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      1980        808          0          0          0          0
inet6.0
                        16         16          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.1              65000       1856       1404        0          0          24:20
607/608/608/0           0/0/0/0
172.30.5.7              65003       2660       2258        0          0          24:16
94/808/808/0            0/0/0/0
192.168.0.38         110047427        511        599        0          0          25:03
107/564/564/0           0/0/0/0
fc09:c0:ffee::6      110047427         56         56        0          0          24:59
Establ
  inet6.0: 16/16/16/0
```

Apply the next-hop self-policy on R8 IBGP and CBGP sessions.

```
[edit]
lab@R8# set policy-options policy-statement nhs from protocol bgp

[edit]
lab@R8# set policy-options policy-statement nhs from route-type external

[edit]
lab@R8# set policy-options policy-statement nhs then next-hop self

[edit]
lab@R8# set protocols bgp group ibgp export nhs

[edit]
lab@R8# set protocols bgp group cbgp export nhs

[edit]
lab@R8# commit
commit complete
```

11) Verify the routing table for hidden routes.

a) R1

```
[edit]
lab@R1# run show route hidden terse

inet.0: 875 destinations, 2004 routes (875 active, 0 holddown, 0 hidden)

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

inet6.0: 31 destinations, 32 routes (31 active, 0 holddown, 0 hidden)
```

No hidden routes are present in R1 routing table.

b) R2

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show route hidden

inet.0: 874 destinations, 2022 routes (874 active, 0 holddown, 0 hidden)

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

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
```

No hidden routes on R2 as well.

c) R3

```
[edit]
lab@R3# run show route hidden

inet.0: 880 destinations, 906 routes (880 active, 0 holddown, 0 hidden)

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

inet6.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

d) R4

```
[edit]
lab@R4# run show route hidden terse

inet.0: 879 destinations, 1003 routes (878 active, 0 holddown, 0 hidden)

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

inet6.0: 32 destinations, 32 routes (32 active, 0 holddown, 0 hidden)
```

e) R5

```
[edit]
lab@R5# run show route hidden

inet.0: 882 destinations, 1005 routes (882 active, 0 holddown, 0 hidden)

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

inet6.0: 55 destinations, 75 routes (55 active, 0 holddown, 0 hidden)
```

f) R6

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R6# run show route hidden

inet.0: 882 destinations, 1004 routes (882 active, 0 holddown, 0 hidden)

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

inet6.0: 31 destinations, 31 routes (31 active, 0 holddown, 0 hidden)
```

g) R7

```
[edit]
lab@R7# run show route hidden

inet.0: 878 destinations, 2105 routes (878 active, 0 holddown, 0 hidden)

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

inet6.0: 50 destinations, 50 routes (50 active, 0 holddown, 0 hidden)
```

h) R8

```
[edit]
lab@R8# run show route hidden

inet.0: 876 destinations, 2088 routes (876 active, 0 holddown, 0 hidden)

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

inet6.0: 50 destinations, 50 routes (50 active, 0 holddown, 0 hidden)
```

All routers do not have hidden routes anymore after applying the next-hop self-policies and adding the external interface addresses as IS-IS passive on R1 and R2.

Task 3: Routing Policies

This task involves controlling BGP routes with route policy. BGP is a very policy-oriented protocol, export and import policies are used for influencing path selection.

BGP stores the routes in three tables: RIB-IN, RIB-LOCAL and RIB-OUT.

The RIB-IN table is created for every established peer to store all routes received from that peer. BGP import policies are applied after the routes are in RIB-IN.

RIB-LOCAL stores routes that the router uses for forwarding traffic.

© 2020 Juniper Networks, Inc. All rights reserved.

RIB-OUT table is again created for every peer and stores routes that are advertised to that peer. Routes stored in RIB-OUT already passed BGP export policies.

This task includes different requirements that you should accomplish with BGP routing policies on the routes received and advertised to the BGP neighbors.

1) Configure the policies.

Repeat this on all eight devices.

Some of the requirements demand to configure the same policy and communities on all routers. One such policy is to accomplish the requirement of not accepting routes 0.0.0.0 with any subnet mask length from all EBGP peers. You have to configure communities on each router because all IPv4 routes should be identified with standard communities.

Standard communities can identify the received IPv4 routes from any of the neighboring ASs.

A community is an attribute used to group routes that share common parameters.

The community format is formed by two 16bits values, combined 32bits divided by ":". The higher 16bits represent the AS number whereas the lower 16bits the number is chosen by the administrator. Using the AS number makes the community unique in Internet.

```
[edit]
lab@R1# set policy-options community C1 members 54591:64512

[edit]
lab@R1# set policy-options community C2 members 54591:64513

[edit]
lab@R1# set policy-options community C3 members 54591:64514

[edit]
lab@R1# set policy-options community IX members 54591:1620

[edit]
lab@R1# set policy-options community P1 members 54591:1679

[edit]
lab@R1# set policy-options community P2 members 54591:53732

[edit]
lab@R1# set policy-options community P3 members 54591:43208

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set policy-options community rtbh members 6451.:666


[edit]
lab@R1# commit
commit complete
```

Note: The last community is using regular expressions. Because you have to match routes from any Customer AS, where the second part of the community value equals 666, we used the "." instead of the last digit. This way `community rtbh` will match the members 64511:666, 64512:666 and etc.

Besides a list of the communities you must also configure policy called default-filter, to filter all 0.0.0.0 routes with any mask length from any EBGP peer.

```
[edit]
lab@R1# set policy-options policy-statement default-filter term 1 from route-filter 0/0
through 0/32

[edit]
lab@R1# set policy-options policy-statement default-filter term 1 then reject
```

The policy matches 0.0.0.0 route with any mask, including 0,1,2 through 32. It matches exactly the routes 0/0, 0/1, 0/2 and etc. but not the prefixes in between. For example, 0.0.0.1/32 is not matched.

You also have to black hole routes advertised from customers with community "<customer AS>:666". A community list matching all those communities was created. Now you have to create a policy, which will be implemented in the import direction on all IBGP an CBGP sessions on each router.

```
[edit]
lab@R1# set policy-options policy-statement rtbh term 1 from community rtbh

[edit]
lab@R1# set policy-options policy-statement rtbh term 1 then next-hop discard
```

The policy "rtbh" matches community list `rtbh` and changes the next-hop of the routes to discard. This way traffic towards these destinations is dropped by the forwarding engine.

To advertising an aggregate summary route of all internal IPv4 routes including RIP and OSPF you have to create aggregate route aggregating all that routes. This route should be configured on all devices.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set routing-options aggregate route 172.30.0.0/16

[edit]
lab@R1# commit
commit complete
```

You must export the aggregate route to the IX peers, but first reject the routes originated from P1 AS. Create an IX export policy where the first term matches all routes with a community of P1 and rejects them. A second term will match the aggregate route from the routing table and accept it for export.

a) R1

```
[edit]
lab@R1# set policy-options policy-statement IX-export term 1 from protocol bgp

[edit]
lab@R1# set policy-options policy-statement IX-export term 1 from community P1

[edit]
lab@R1# set policy-options policy-statement IX-export term 1 then reject

[edit]
lab@R1# set policy-options policy-statement IX-export term 2 from protocol aggregate

[edit]
lab@R1# set policy-options policy-statement IX-export term 2 from route-filter
172.30.0.0/16 exact

[edit]
lab@R1# set policy-options policy-statement IX-export term 2 then accept
```

An import policy must be applied to allow only routes with a mask shorter than /8 and longer than /24. It also sets the community IX to the accepted routes. Everything else should be filtered, the second term accomplishes this.

```
[edit]
lab@R1# set policy-options policy-statement IX-filter term 1 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R1# set policy-options policy-statement IX-filter term 1 then community set IX

[edit]
lab@R1# set policy-options policy-statement IX-filter term 1 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set policy-options policy-statement IX-filter term 2 then reject
```

The first term is using the `prefix-length-range` match statement with prefix 0.0.0.0/0 and a mask between /8 and /24. This will ensure that all prefixes with a mask between /8 and /24 including is matched. Then community `IX` is set, and the prefixes are accepted. The second term is mandatory because the default BGP policy is to accept all routes.

Next, the policies should be applied in the right direction.

```
[edit]
lab@R1# set protocols bgp group IX import default-filter

[edit]
lab@R1# set protocols bgp group IX import IX-filter
```

Policies default-filter and IX-filter are applied in import direction.

```
[edit]
lab@R1# set protocols bgp group IX export IX-export
```

IX-export is the only policy in export direction.

Policy rtbh is applied in import direction to all IBGP and CBGP sessions.

```
[edit]
lab@R1# set protocols bgp group ibgp import rtbh

[edit]
lab@R1# set protocols bgp group cbgp import rtbh

[edit]
lab@R1# commit
commit complete
```

The result of the policies will be checked after all routers are configured with policies.

b) R2

R2 policy configuration is very similar with the policies of R1 because R2 has EBGP sessions to the external peers.

```
[edit]
lab@R2# set policy-options policy-statement IX-export term 1 from protocol bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set policy-options policy-statement IX-export term 1 from community P1

[edit]
lab@R2# set policy-options policy-statement IX-export term 1 then reject

[edit]
lab@R2# set policy-options policy-statement IX-export term 2 from protocol aggregate

[edit]
lab@R2# set policy-options policy-statement IX-export term 2 from route-filter
172.30.0.0/16 exact

[edit]
lab@R2# set policy-options policy-statement IX-export term 2 then as-path-prepend
"54591 54591 54591"

[edit]
lab@R2# set policy-options policy-statement IX-export term 2 then accept

[edit]
lab@R2# set policy-options policy-statement IX-export term 3 from protocol bgp

[edit]
lab@R2# set policy-options policy-statement IX-export  term 3 then as-path-prepend
"54591 54591 54591"

[edit]
lab@R2# set policy-options policy-statement IX-export term 3 then accept
```

The first term is the same as on R1. Its purpose is to filter routes with community P1. The next two terms are different because you have to make sure that the IX peers prefer prefixes advertised by R1. You can accomplish that by artificially make the AS path longer. Prepending your own AS number three times is enough for the R1 routes to be preferred. Two terms are needed because you need to add one term more for the BGP routes which are matched by the default policy on R1, but on R2 you have add prepend to them as well.

Note: You can use any AS number in the prepend policy, but its best to use your own. This will ensure that the routes will not be rejected in the Internet by the BGP loop prevention mechanism.

Import policy is basically the same as on R1.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# set policy-options policy-statement IX-filter term 1 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R2# set policy-options policy-statement IX-filter term 1 then community set IX

[edit]
lab@R2# set policy-options policy-statement IX-filter term 1 then accept

[edit]
lab@R2# set policy-options policy-statement IX-filter term 2 then reject
```

Apply the policies on R2.

```
[edit]
lab@R2# set protocols bgp group IX import default-filter

[edit]
lab@R2# set protocols bgp group IX import IX-filter

[edit]
lab@R2# set protocols bgp group IX export IX-export

[edit]
lab@R2# set protocols bgp group ibgp import rtbh

[edit]
lab@R2# set protocols bgp group cbgp import rtbh

[edit]
lab@R2# commit
commit complete
```

c) R3

R3 has two external sessions with P2 and P3.

Create policies for session with P2.

```
[edit]
lab@R3# set policy-options policy-statement P2-export term 1 from protocol aggregate

[edit]
lab@R3# set policy-options policy-statement P2-export term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R3# set policy-options policy-statement P2-export term 1 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

The export policy to P2 is advertising the aggregate route of all internal routes and all other BGP routes selected as best path, because of the default BGP policy.

The import policy should accept only routes with mask length between /8 and /24.

```
[edit]
lab@R3# set policy-options policy-statement P2-filter term 1 from route-filter
0.0.0.0/0 prefix-length-range /8-/24

[edit]
lab@R3# set policy-options policy-statement P2-filter term 1 then local-preference 200

[edit]
lab@R3# set policy-options policy-statement P2-filter term 1 then community set P2

[edit]
lab@R3# set policy-options policy-statement P2-filter term 1 then accept

[edit]
lab@R3# set policy-options policy-statement P2-filter term 2 then reject
```

Here the local-preference BGP attribute is changed to 200, this means that the routes from P2 are made more preferred compared to the same routes received from IX.

Create a policy for P3's session.

```
[edit]
lab@R3# set policy-options policy-statement P3-export term 1 from protocol aggregate

[edit]
lab@R3# set policy-options policy-statement P3-export term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R3# set policy-options policy-statement P3-export term 1 then accept
```

Export policy is advertising the aggregate route from the routing table.

You have a requirement to accept /32 routes only for AS 43208.365. Peer P3 is part of that AS. Adding a term which accepts routes originated from that AS before the term accepting routes with a mask between /8 and /24 will accomplish this. For this purpose, you have to create an `as-path` under `policy-options` and use that `as-path` in the policy to match routes originated from AS 43208.365.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set policy-options as-path P3-local-routes 2831679853
```

Note: Here ASPLAIN format is used for describing the AS number. You can see the ASPLAIN format from the output of the command `show bgp summary`.

The configuration of the as-path accepts regular expressions as well. In the current example only routes that have that AS number in AS path are matched. Routes coming from another AS and passing through AS 43208.365 are not matched.

Create the import policy for session with P3.

```
[edit]
lab@R3# set policy-option policy-statement P3-filter term 1 from as-path P3-local-
routes

[edit]
lab@R3# set policy-option policy-statement P3-filter term 1 from route-filter 0/0
prefix-length-range /32-/32

[edit]
lab@R3# set policy-option policy-statement P3-filter term 1 then local-preference 200

[edit]
lab@R3# set policy-option policy-statement P3-filter term 1 then community set P3

[edit]
lab@R3# set policy-option policy-statement P3-filter term 1 then accept

[edit]
lab@R3# set policy-option policy-statement P3-filter term 2 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R3# set policy-option policy-statement P3-filter term 2 then local-preference 200

[edit]
lab@R3# set policy-option policy-statement P3-filter term 2 then community set P3

[edit]
lab@R3# set policy-option policy-statement P3-filter term 2 then accept

[edit]
lab@R3# set policy-option policy-statement P3-filter term 3 then reject
```

Apply the policies to P2 and P3 sessions.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set protocols bgp group P2-1 export P2-export

[edit]
lab@R3# set protocols bgp group P2-1 import default-filter

[edit]
lab@R3# set protocols bgp group P2-1 import P2-filter

[edit]
lab@R3# set protocols bgp group P3-1 export P3-export

[edit]
lab@R3# set protocols bgp group P3-1 import default-filter

[edit]
lab@R3# set protocols bgp group P3-1 import P3-filter

[edit]
lab@R3# set protocols bgp group ibgp import rtbh

[edit]
lab@R3# set protocols bgp group cbgp import rtbh

[edit]
lab@R3# commit
commit complete
```

d) R4

R4 has no external BGP sessions. Here only the policies for Remote triggered black hole should be implemented.

```
[edit]
lab@R4# set protocols bgp group ibgp import rtbh

[edit]
lab@R4# set protocols bgp group cbgp import rtbh

[edit]
lab@R4# commit
commit complete
```

e) R5

R5 has two sessions with customers from AS 64514. Here you must be careful because these BGP sessions carry IPv6 and IPv4 prefixes. Routing policies should deal with IPv4 traffic also.

© 2020 Juniper Networks, Inc. All rights reserved.

For this reason, you must match the whole family inet6 and accept all prefixes.

Create an export policy `as-internal` which will match the aggregate route and exports it to both EBGP sessions.

```
[edit]
lab@R5# set policy-options policy-statement as-internal term 1 from protocol aggregate

[edit]
lab@R5# set policy-options policy-statement as-internal term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R5# set policy-options policy-statement as-internal term 1 then accept
```

There is a requirement that customer routes should be preferred over the same routes learned from any other peer. Only if the routes are marked with specific community, then these routes should be less preferred. For this purpose, you have to create and apply one additional community to the routes having EBGP sessions with customer neighbors.

```
[edit]
lab@R5# set policy-options community C3-low-pref members 64514:90

[edit]
lab@R5# set policy-options policy-statement C3-filter term 1 from family inet6

[edit]
lab@R5# set policy-options policy-statement C3-filter term 1 then accept

[edit]
lab@R5# set policy-options policy-statement C3-filter term 2 from community C3-low-pref

[edit]
lab@R5# set policy-options policy-statement C3-filter term 2 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R5# set policy-options policy-statement C3-filter term 2 then local-preference 90

[edit]
lab@R5# set policy-options policy-statement C3-filter term 2 then community add C3

[edit]
lab@R5# set policy-options policy-statement C3-filter term 2 then accept

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set policy-options policy-statement C3-filter term 3 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R5# set policy-options policy-statement C3-filter term 3 then local-preference 300

[edit]
lab@R5# set policy-options policy-statement C3-filter term 3 then community add C3

[edit]
lab@R5# set policy-options policy-statement C3-filter term 3 then accept

[edit]
lab@R5# set policy-options policy-statement C3-filter term 4 then reject
```

Here term 2 matches routes with a community value of `64514:90` and sets a local-preference of 90. Other routes with a mask between /8 and /24 are matched by `term 3`. The term sets the local-preference to 300, which makes the routes more preferred. In term 3 you can add match criteria based on the as-path to be sure that only the routes originated from that AS are matched. However usually a customer AS are not considered a transit AS.

Note: Here the community is applied with "`add`" not with "`set`". The difference is that with add you are adding the community to the list of communities already carried with the route, while with set you override the community list and set the new one. With that policy you need to use `add` because you are expecting routes with community `rtbh` to come from the customers.

Apply the policies on the EBGP sessions.

```
[edit]
lab@R5# set protocols bgp group C3 export as-internal

[edit]
lab@R5# set protocols bgp group C3 import default-filter

[edit]
lab@R5# set protocols bgp group C3 import C3-filter

[edit]
lab@R5# set protocols bgp group ibgp import rtbh

[edit]
lab@R5# set protocols bgp group cbgp import rtbh

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# commit
commit complete
```

f) R6

R6 has EBGP sessions with two customers from different ASs.

Here you must create two new communities for matching less preferred routes, one for C2 and one C1.

```
[edit]
lab@R6# set policy-options community C1-low-pref members 64512:90

[edit]
lab@R6# set policy-options community C2-low-pref members 64513:90
```

Because you have to advertise only a default route to C2 you have to create an aggregate route. Make sure you also apply the fix for the AS Path length restriction.

```
[edit]
lab@R6# set routing-options aggregate route 0/0

[edit]
lab@R6# set routing-options aggregate defaults brief
```

Create export policies for the sessions with C1-1 neighbor.

```
[edit]
lab@R6# set policy-options policy-statement as-internal term 1 from protocol aggregate

[edit]
lab@R6# set policy-options policy-statement as-internal term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R6# set policy-options policy-statement as-internal term 1 then metric 10

[edit]
lab@R6# set policy-options policy-statement as-internal term 1 then accept
```

The policy `as-internal` is advertising all internal routes. The difference with the policy on R5 is the setting of the BGP MED attribute on the aggregate route send to C1-1. This is needed because R6 should be the preferred exit and entry point for customer C1. You are

© 2020 Juniper Networks, Inc. All rights reserved.

setting a MED of 10 on R6 and MED of 20 on R7. This will ensure that C1 will choose the routes with a lower MED coming from R6.

The policy above changed the MED on the aggregate route advertised to C1. To set the MED on other BGP routes advertised by the default BGP policy you have to create a policy, which matches these routes and set the MED to 10.

```
[edit]
lab@R6# set policy-options policy-statement med-10 term 1 from protocol bgp

[edit]
lab@R6# set policy-options policy-statement med-10 term 1 then metric 10

[edit]
lab@R6# set policy-options policy-statement med-10 term 1 then accept
```

Configure import policy for C1 session.

```
[edit]
lab@R6# set policy-option policy-statement C1-filter term 1 from community C1-low-pref

[edit]
lab@R6# set policy-option policy-statement C1-filter term 1 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R6# set policy-option policy-statement C1-filter term 1 then local-preference 90

[edit]
lab@R6# set policy-option policy-statement C1-filter term 1 then community add C1

[edit]
lab@R6# set policy-option policy-statement C1-filter term 1 then accept

[edit]
lab@R6# set policy-option policy-statement C1-filter term 2 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R6# set policy-option policy-statement C1-filter term 2 then local-preference 400

[edit]
lab@R6# set policy-option policy-statement C1-filter term 2 then community add C1

[edit]
lab@R6# set policy-option policy-statement C1-filter term 2 then accept

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R6# set policy-option policy-statement C1-filter term 3 then reject
```

Term 2 is setting the local-preference to 400 to make R6 the preferred exit point to neighbor C1.

Configure an export policy for the session with C2-1. You should advertise only a default route to C2.

```
[edit]
lab@R6# set policy-options policy-statement default-only term 1 from protocol aggregate

[edit]
lab@R6# set policy-options policy-statement default-only term 1 from route-filter 0/0
exact

[edit]
lab@R6# set policy-options policy-statement default-only term 1 then accept

[edit]
lab@R6# set policy-options policy-statement default-only term 2 then reject
```

Unlike the export policies on the other routers here you must configure `term 2` to `reject` the readvertisement of other BGP routes, because you have to advertise only a default route.

Configure import policy for C2-1.

```
[edit]
lab@R6# set policy-options policy-statement C2-filter term 1 from community C2-low-pref

[edit]
lab@R6# set policy-options policy-statement C2-filter term 1 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R6# set policy-options policy-statement C2-filter term 1 then local-preference 90

[edit]
lab@R6# set policy-options policy-statement C2-filter term 1 then community add C2

[edit]
lab@R6# set policy-options policy-statement C2-filter term 1 then accept

[edit]
lab@R6# set policy-options policy-statement C2-filter term 2 from route-filter 0/0
prefix-length-range /8-/24
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set policy-options policy-statement C2-filter term 2 then local-preference 300

[edit]
lab@R6# set policy-options policy-statement C2-filter term 2 then community add C2

[edit]
lab@R6# set policy-options policy-statement C2-filter term 2 then accept

[edit]
lab@R6# set policy-options policy-statement C2-filter term 3 then reject
```

Term 1 sets the local-preference to less preferred for routes received with a community value of 64513:90. Whereas other routes are configured with a local-preference of 300 to be more preferred compared to the same routes received from other neighbors.

Apply the policies to the BGP sessions.

```
[edit]
lab@R6# set protocols bgp group C1-1 export as-internal

[edit]
lab@R6# set protocols bgp group C1-1 export med-10

[edit]
lab@R6# set protocols bgp group C1-1 import default-filter

[edit]
lab@R6# set protocols bgp group C1-1 import C1-filter

[edit]
lab@R6# set protocols bgp group C2-1 export default-only

[edit]
lab@R6# set protocols bgp group C2-1 import default-filter

[edit]
lab@R6# set protocols bgp group C2-1 import C2-filter

[edit]
lab@R6# set protocols bgp group ibgp import rtbh

[edit]
lab@R6# set protocols bgp group cbgp import rtbh

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R6# commit
commit complete
```

g) R7.

R7 has an EBGP session with C1 and P1-2. R6 should be preferred by C1 for both inbound and outbound traffic. R7 should advertise BGP routes to C1 with higher MED. Again, you need to configure a community for recognizing the less preferred traffic from C1.

```
[edit]
lab@R7# set policy-options community C1-low-pref members 64512:90
```

Configure the export policies for the session with C1.

```
[edit]
lab@R7# set policy-options policy-statement as-internal term 1 from protocol aggregate

[edit]
lab@R7# set policy-options policy-statement as-internal term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R7# set policy-options policy-statement as-internal term 1 then metric 20

[edit]
lab@R7# set policy-options policy-statement as-internal term 1 then accept
```

Configure a policy to set the MED to 20 for other re-advertised BGP routes.

```
[edit]
lab@R7# set policy-options policy-statement med-20 term 1 from protocol bgp

[edit]
lab@R7# set policy-options policy-statement med-20 term 1 then metric 20

[edit]
lab@R7# set policy-options policy-statement med-20 term 1 then accept
```

Configure import policy for C1 neighbor.

```
[edit]
lab@R7# set policy-options policy-statement C1-filter term 1 from community C1-low-pref
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set policy-options policy-statement C1-filter term 1 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R7# set policy-options policy-statement C1-filter term 1 then local-preference 90

[edit]
lab@R7# set policy-options policy-statement C1-filter term 1 then community add C1

[edit]
lab@R7# set policy-options policy-statement C1-filter term 1 then accept

[edit]
lab@R7# set policy-options policy-statement C1-filter term 2 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R7# set policy-options policy-statement C1-filter term 2 then local-preference 300

[edit]
lab@R7# set policy-options policy-statement C1-filter term 2 then community add C1

[edit]
lab@R7# set policy-options policy-statement C1-filter term 2 then accept

[edit]
lab@R7# set policy-options policy-statement C1-filter term 3 then reject
```

The routes received from C1 are set with a local-preference of 300, which makes them more preferred from the same routes received from IX, but less preferred then the same routes received via R6.

The export policy for P1s session should advertise part of your internal prefixes to ensure equal load-balancing for traffic entering your AS from P1 AS. For that purpose, you have to create a new aggregate route and advertise it to P1 with the well-known `no-export` community. This means that only routes from P1 AS will have those communities and will not be advertised to other ASs. To ensure that your AS routes will be visible by the neighboring ASs of P1, you should advertise the summary route for your local AS routes to P1 as well.

```
[edit]
lab@R7# set routing-options aggregate route 172.30.128.0/17

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# set policy-options community no-export members no-export

[edit]
lab@R7# set policy-options policy-statement P1-export term 1 from community IX

[edit]
lab@R7# set policy-options policy-statement P1-export term 1 then reject

[edit]
lab@R7# set policy-options policy-statement P1-export term 2 from protocol aggregate

[edit]
lab@R7# set policy-options policy-statement P1-export term 2 from route-filter
172.30.128.0/17 exact

[edit]
lab@R7# set policy-options policy-statement P1-export term 2 then community set no-
export

[edit]
lab@R7# set policy-options policy-statement P1-export term 2 then accept

[edit]
lab@R7# set policy-options policy-statement P1-export term 3 from protocol aggregate

[edit]
lab@R7# set policy-options policy-statement P1-export term 3 from route-filter
172.30.0.0/16 exact

[edit]
lab@R7# set policy-options policy-statement P1-export term 3 then accept
```

The first term rejects all routes received from IX AS, this as per requirement.

Note: Although the `no-export` community is a well-known community it should be explicitly configured as any other custom community.

Configure import policy for P1 session.

You should accept only routes that originated in P1 AS, no transit routes should be accepted. R8 should be the preferred exit point for those routes.

```
[edit]
lab@R7# set policy-options as-path P1 110047427
```

Creating AS path that matches routes originated by P1 AS.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set policy-options policy-statement P1-filter term 1 from as-path P1

[edit]
lab@R7# set policy-options policy-statement P1-filter term 1 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R7# set policy-options policy-statement P1-filter term 1 then local-preference 200

[edit]
lab@R7# set policy-options policy-statement P1-filter term 1 then community set P1

[edit]
lab@R7# set policy-options policy-statement P1-filter term 1 then accept

[edit]
lab@R7# set policy-options policy-statement P1-filter term 2 then reject
```

Applying the policies to the BGP sessions.

```
[edit]
lab@R7# set protocols bgp group C1-1 import default-filter

[edit]
lab@R7# set protocols bgp group C1-1 import C1-filter

[edit]
lab@R7# set protocols bgp group C1-1 export as-internal

[edit]
lab@R7# set protocols bgp group C1-1 export med-20

[edit]
lab@R7# set protocols bgp group P1-2 import default-filter

[edit]
lab@R7# set protocols bgp group P1-2 import P1-filter

[edit]
lab@R7# set protocols bgp group P1-2 export P1-export

[edit]
lab@R7# set protocols bgp group ibgp import rtbh

[edit]
lab@R7# set protocols bgp group cbgp import rtbh

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# commit
commit complete
```

h) R8

R8 has just one EBGP session with P1-1.

Create an export policy with the first term advertising the other half of your local routes with the no-export community.

```
[edit]
lab@R8# set routing-options aggregate route 172.30.0.0/17

[edit]
lab@R8# set policy-options community no-export members no-export

[edit]
lab@R8# set policy-options policy-statement P1-export term 1 from protocol bgp

[edit]
lab@R8# set policy-options policy-statement P1-export term 1 from community IX

[edit]
lab@R8# set policy-options policy-statement P1-export term 1 then reject

[edit]
lab@R8# set policy-options policy-statement P1-export term 2 from protocol aggregate

[edit]
lab@R8# set policy-options policy-statement P1-export term 2 from route-filter
172.30.0.0/17 exact

[edit]
lab@R8# set policy-options policy-statement P1-export term 2 then community set no-
export

[edit]
lab@R8# set policy-options policy-statement P1-export term 2 then accept

[edit]
lab@R8# set policy-options policy-statement P1-export term 3 from protocol aggregate

[edit]
lab@R8# set policy-options policy-statement P1-export term 3 from route-filter
172.30.0.0/16 exact

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# set policy-options policy-statement P1-export term 3 then accept
```

Configure an import policy for P1 EBGP session. The routes are assigned a local-preference of 210, which makes them more preferred then the same routes received from R7.

```
[edit]
lab@R8# set policy-options as-path P1 110047427

[edit]
lab@R8# set policy-options policy-statement P1-filter term 1 from as-path P1

[edit]
lab@R8# set policy-options policy-statement P1-filter term 1 from route-filter 0/0
prefix-length-range /8-/24

[edit]
lab@R8# set policy-options policy-statement P1-filter term 1 then local-preference 210

[edit]
lab@R8# set policy-options policy-statement P1-filter term 1 then community set P1

[edit]
lab@R8# set policy-options policy-statement P1-filter term 1 then accept

[edit]
lab@R8# set policy-options policy-statement P1-filter term 2 then reject

[edit]
lab@R8# commit
commit complete
```

Applying the policies to the BGP sessions.

```
[edit]
lab@R8# set protocols bgp group P1-1 import default-filter

[edit]
lab@R8# set protocols bgp group P1-1 import P1-filter

[edit]
lab@R8# set protocols bgp group P1-1 export P1-export

[edit]
lab@R8# set protocols bgp group ibgp import rtbh

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# set protocols bgp group cbgp import rtbh

[edit]
lab@R8# commit
commit complete
```

2) Verify the policy configuration.

Now you have to verify if all requirements are accomplished with the policies implemented.

a) R1

Verify if unresolved routes are present in the routing table.

```
[edit]
lab@R1# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

No unresolved routes are present.

Verify if routes with a mask shorter than /8 and longer than /24 are present in routing table.

```
[edit]
lab@R1# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* ? 172.17.47.2/32     B 170        200                    (65001) 2831679853 I
* ? 172.17.47.3/32     B 170        200                    (65001) 2831679853 I
```

Only both routes received from P3-1 are present with mask longer than /32. They are allowed by the task.

Verify the routing table for 0.0.0.0 with any mask.

```
[edit]
lab@R1# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and advertised to other internal BGP peers.

```
[edit]
lab@R1# run show route advertising-protocol bgp 172.30.5.8 aspath-regex "1620 .*"
community-name IX

inet.0: 830 destinations, 1796 routes (668 active, 0 holddown, 324 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Prefix          Nexthop          MED   Lclpref    AS path
* 1.64.0.0/10     192.168.1.3            100        1620 61671 I
* 1.84.160.0/20   192.168.1.3            100        1620 33112 I
* 1.96.0.0/11     192.168.1.3            100        1620 33112 63164 40776 51777 I
* 1.161.192.0/21  192.168.1.3            100        1620 33112 30404 32138 45045 I
* 1.166.89.0/24   192.168.1.3            100        1620 110047427 56422 47123
34908 57085 3090 19128 21300 I
* 1.176.0.0/12    192.168.1.3            100        1620 33112 49129 16320 52954 I
* 3.49.128.0/19   192.168.1.3            100        1620 110047427 56422 7923
42335 17359 9293 38097 38522 61520 I
* 3.189.0.0/16    192.168.1.3            100        1620 33112 41164 7136 6370
35389 39823 16965 I
...
```

R1 is advertising routes to R8 with community IX.

Verify if the customer routes are advertised to IX peers.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.1.3 aspath-regex
"64512|64513|64514"

inet.0: 830 destinations, 1796 routes (668 active, 0 holddown, 324 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24      Self                                    64512 I
* 172.31.1.0/24      Self                                    64512 I
* 172.31.2.0/24      Self                                    64512 I
* 172.31.3.0/24      Self                                    64512 I
* 172.31.4.0/24      Self                                    64512 I
* 172.31.5.0/24      Self                                    64512 I
* 172.31.6.0/24      Self                                    64512 I
* 172.31.7.0/24      Self                                    64512 I
* 172.31.8.0/24      Self                                    64512 I
* 172.31.9.0/24      Self                                    64512 I
* 172.31.10.0/24     Self                                    64512 I
* 172.31.11.0/24     Self                                    64512 I
* 172.31.12.0/24     Self                                    64512 I
* 172.31.13.0/24     Self                                    64512 I
* 172.31.14.0/24     Self                                    64512 I
* 172.31.15.0/24     Self                                    64512 I
* 172.31.16.0/24     Self                                    64513 I
* 172.31.17.0/24     Self                                    64513 I
* 172.31.18.0/24     Self                                    64513 I
* 172.31.19.0/24     Self                                    64513 I
* 172.31.20.0/24     Self                                    64513 I
* 172.31.21.0/24     Self                                    64513 I
* 172.31.22.0/24     Self                                    64513 I
* 172.31.23.0/24     Self                                    64513 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.24.0/24    Self                                  64513 I
* 172.31.25.0/24    Self                                  64513 I
* 172.31.26.0/24    Self                                  64513 I
* 172.31.27.0/24    Self                                  64513 I
* 172.31.28.0/24    Self                                  64513 I
* 172.31.29.0/24    Self                                  64513 I
* 172.31.30.0/24    Self                                  64513 I
* 172.31.31.0/24    Self                                  64513 I
* 172.31.32.0/24    Self                                  64514 I
* 172.31.33.0/24    Self                                  64514 I
* 172.31.34.0/24    Self                                  64514 I
* 172.31.35.0/24    Self                                  64514 I
* 172.31.36.0/24    Self                                  64514 I
* 172.31.37.0/24    Self                                  64514 I
* 172.31.38.0/24    Self                                  64514 I
* 172.31.39.0/24    Self                                  64514 I
* 172.31.40.0/24    Self                                  64514 I
* 172.31.41.0/24    Self                                  64514 I
* 172.31.42.0/24    Self                                  64514 I
* 172.31.43.0/24    Self                                  64514 I
* 172.31.44.0/24    Self                                  64514 I
* 172.31.45.0/24    Self                                  64514 I
* 172.31.46.0/24    Self                                  64514 I
* 172.31.47.0/24    Self                                  64514 I
...
```

Verify if the local summary prefix is sent to EBGP peers.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.1.3 172.30/16

inet.0: 830 destinations, 1796 routes (668 active, 0 holddown, 324 hidden)
  Prefix              Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16       Self                                    I
```

Verify the local-preference value of the customer routes received by R1.

```
[edit]
lab@R1# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 830 destinations, 1796 routes (668 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24       *[BGP/170] 00:41:57, localpref 90, from 172.30.5.8
                       AS path: (65003) 64512 I
                     > to 172.30.0.2 via ae0.0
172.31.1.0/24       *[BGP/170] 01:13:47, localpref 400, from 172.30.5.2
                       AS path: (65001 65002) 64512 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.2 via ae0.0
                      to 172.30.0.6 via ge-0/0/4.114
                      to 172.30.0.10 via ge-0/0/4.118
172.31.2.0/24       *[BGP/170] 01:13:47, localpref 400, from 172.30.5.2
                      AS path: (65001 65002) 64512 I
                    > to 172.30.0.2 via ae0.0
                      to 172.30.0.6 via ge-0/0/4.114
                      to 172.30.0.10 via ge-0/0/4.118
172.31.3.0/24       *[BGP/170] 01:13:47, localpref 400, from 172.30.5.2
                      AS path: (65001 65002) 64512 I
                    > to 172.30.0.2 via ae0.0
                      to 172.30.0.6 via ge-0/0/4.114
                      to 172.30.0.10 via ge-0/0/4.118
```

Verify the remote triggered black hole routes with community rtbh.

```
[edit]
lab@R1# run show route protocol bgp community-name rtbh

inet.0: 830 destinations, 1796 routes (668 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24       *[BGP/170] 00:00:48, localpref 400, from 172.30.5.8
                      AS path: (65003 65002) 64512 I
                      Discard
...
```

Verify if the routes received from P1, P2 and P3 peers are preferred to routes received by IX peers.

```
[edit]
lab@R1# run show route 172.17.0.0/24

inet.0: 830 destinations, 1796 routes (668 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.0.0/24       *[BGP/170] 00:53:49, localpref 210, from 172.30.5.8
                      AS path: (65003) 110047427 I
                    > to 172.30.0.10 via ge-0/0/4.118
                    [BGP/170] 01:55:17, localpref 100
                      AS path: 1620 110047427 I
                    > to 192.168.1.3 via ge-0/0/5.300
                    [BGP/170] 01:55:13, localpref 100
                      AS path: 1620 110047427 I
                    > to 192.168.1.4 via ge-0/0/5.300
```


© 2020 Juniper Networks, Inc. All rights reserved.

The routes received from R8 are preferred even though R1 has direct connection to IX peers.

Verify if the routes with P1 are not send to IX peers.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.1.3 community-name P1

inet.0: 830 destinations, 2021 routes (668 active, 0 holddown, 549 hidden)

[edit]
lab@R1# run show route advertising-protocol bgp 192.168.1.4 community-name P1

inet.0: 830 destinations, 2021 routes (668 active, 0 holddown, 549 hidden)
```

b) R2.

Verify if unresolved routes are present in the routing table.

```
[edit]
lab@R2# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify if the routes with a mask shorter than /8 and longer than /24 are present in routing table.

```
[edit]
lab@R2# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32    B 170    200  >172.30.0.14       (65001) 2831679853 I
* 172.17.47.3/32    B 170    200  >172.30.0.14       (65001) 2831679853 I
```

Here again only both the /32 routes from P3-1 are present in the routing table with a mask shorter than /8 and longer than /24.

Verify the routing table for route 0.0.0.0 with any mask.

```
[edit]
lab@R2# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and advertised to other internal BGP peers.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show route  advertising-protocol bgp 172.30.5.3 aspath-regex "1620
.*" community-name IX

inet.0: 830 destinations, 1852 routes (668 active, 0 holddown, 324 hidden)
  Prefix          Nexthop         MED   Lclpref    AS path
* 1.64.0.0/10     192.168.1.4             100        1620 61671 I
* 1.84.160.0/20   192.168.1.4             100        1620 33112 I
* 1.96.0.0/11     192.168.1.4             100        1620 33112 63164 40776 51777 I
* 1.161.192.0/21  192.168.1.4             100        1620 33112 30404 32138 45045 I
* 1.166.89.0/24   192.168.1.4             100        1620 110047427 56422 47123
34908 57085 3090 19128 21300 I
* 1.176.0.0/12    192.168.1.4             100        1620 33112 49129 16320 52954 I
* 3.49.128.0/19   192.168.1.4             100        1620 110047427 56422 7923 42335
17359 9293 38097 38522 61520 I
* 3.189.0.0/16    192.168.1.4             100        1620 33112 41164 7136 6370
35389 39823 16965 I
...
```

Verify if customer routes are advertised to IX peers.

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.1.4 aspath-regex ".*
[64512-64514]"

inet.0: 830 destinations, 1960 routes (830 active, 0 holddown, 324 hidden)
  Prefix            Nexthop    MED     Lclpref    AS path
* 172.31.0.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.1.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.2.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.3.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.4.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.5.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.6.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.7.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.8.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.9.0/24     Self                       54591 54591 54591 [65000] 64512 I
* 172.31.10.0/24    Self                       54591 54591 54591 [65000] 64512 I
* 172.31.11.0/24    Self                       54591 54591 54591 [65000] 64512 I
* 172.31.12.0/24    Self                       54591 54591 54591 [65000] 64512 I
* 172.31.13.0/24    Self                       54591 54591 54591 [65000] 64512 I
* 172.31.14.0/24    Self                       54591 54591 54591 [65000] 64512 I
* 172.31.15.0/24    Self                       54591 54591 54591 [65000] 64512 I
* 172.31.16.0/24    Self                       54591 54591 54591 [65000] 64513 I
* 172.31.17.0/24    Self                       54591 54591 54591 [65000] 64513 I
* 172.31.18.0/24    Self                       54591 54591 54591 [65000] 64513 I
* 172.31.19.0/24    Self                       54591 54591 54591 [65000] 64513 I
* 172.31.20.0/24    Self                       54591 54591 54591 [65000] 64513 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.21.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.22.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.23.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.24.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.25.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.26.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.27.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.28.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.29.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.30.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.31.0/24    Self                    54591 54591 54591 [65000] 64513 I
* 172.31.32.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.33.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.34.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.35.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.36.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.37.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.38.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.39.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.40.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.41.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.42.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.43.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.44.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.45.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.46.0/24    Self                    54591 54591 54591 [65000] 64514 I
* 172.31.47.0/24    Self                    54591 54591 54591 [65000] 64514 I
```

You can see that the AS is prepended three times to the advertised routes.

Note: A different regular expression is used to display the customer routes compared to the output of R1. The result of the output is the same.

Verify if the local summary prefix is sent to EBGP peers.

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.1.4 172.30/16

inet.0: 830 destinations, 1797 routes (668 active, 0 holddown, 324 hidden)
  Prefix          Nexthop   MED   Lclpref  AS path
* 172.30.0.0/16      Self                  54591 54591 54591 [65000] I
```

Verify the local-preference of the customer routes received by R2.

```
[edit]
lab@R2# run show route protocol bgp aspath-regex ".* [64512-64514]"
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 830 destinations, 1960 routes (830 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[BGP/170] 01:28:19, localpref 90, from 172.30.5.1
                      AS path: (65003) 64512 I
                    > to 172.30.0.18 via ge-0/0/4.127
                    [BGP/170] 01:28:18, localpref 90, from 172.30.5.3
                      AS path: (65001 65002) 64512 I
                    > to 172.30.0.14 via ge-0/0/4.123
                      to 172.30.0.18 via ge-0/0/4.127
172.31.1.0/24      *[BGP/170] 00:33:53, localpref 400, from 172.30.5.1
                      AS path: (65003 65002) 64512 I
                      Discard
172.31.2.0/24      *[BGP/170] 02:00:09, localpref 400, from 172.30.5.3
                      AS path: (65001 65002) 64512 I
                    > to 172.30.0.14 via ge-0/0/4.123
                      to 172.30.0.18 via ge-0/0/4.127
...
```

Verify the remote triggered black hole routes with community rtbh.

```
[edit]
lab@R2# run show route protocol bgp community-name rtbh

inet.0: 830 destinations, 1797 routes (668 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 00:34:39, localpref 400, from 172.30.5.1
                      AS path: (65003 65002) 64512 I
                      Discard
```

Verify if the routes received from P1, P2 and P3 peers are preferred to the routes received by IX peers.

```
[edit]
lab@R2# run show route 172.17.0.0/24

inet.0: 830 destinations, 1797 routes (668 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.0.0/24      *[BGP/170] 01:24:36, localpref 210, from 172.30.5.1
                      AS path: (65003) 110047427 I
                    > to 172.30.0.1 via ae0.0
                    [BGP/170] 02:10:08, localpref 100
                      AS path: 1620 110047427 I
                    > to 192.168.1.3 via ge-0/0/5.300
                    [BGP/170] 02:10:13, localpref 100
                      AS path: 1620 110047427 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 192.168.1.4 via ge-0/0/5.300
```

Verify if routes from P1 are not send to IX peers.

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.1.3 community-name P1

inet.0: 830 destinations, 1960 routes (830 active, 0 holddown, 324 hidden)

[edit]
lab@R2# run show route advertising-protocol bgp 192.168.1.4 community-name P1

inet.0: 830 destinations, 1960 routes (830 active, 0 holddown, 324 hidden)
```

c) R3.

Verify if unresolved routes are present in the routing table.

```
[edit]
lab@R3# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify if routes with a mask shorter than /8 and longer than /24 are present in routing table.

```
[edit]
lab@R3# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32  B 170    200          >192.168.0.6      2831679853 I
* 172.17.47.3/32  B 170    200          >192.168.0.6      2831679853 I
```

Verify the routing table for route/ prefix 0.0.0.0 with any mask.

```
[edit]
lab@R3# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and advertised to other internal BGP peers.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.4 community-name P2

inet.0: 881 destinations, 896 routes (835 active, 0 holddown, 60 hidden)
  Prefix           Nexthop    MED     Lclpref    AS path
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 12.173.96.0/19    Self                  200        3521382357 4637 7478 16572 33786
?
* 15.36.0.0/15      Self                  200        3521382357 4637 57359 16881
43174 64323 I
* 17.96.0.0/13      Self                  200        3521382357 4637 22896 I
* 26.0.0.0/8        Self                  200        3521382357 46848 I
* 30.0.0.0/9        Self                  200        3521382357 4637 24130 34708
32310 4432 54904 ?
...

[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.4 community-name P3

inet.0: 881 destinations, 896 routes (835 active, 0 holddown, 60 hidden)
  Prefix             Nexthop      MED     Lclpref    AS path
* 5.127.0.0/17       Self                  200        2831679853 9726 36659 30705
25538 37414 49276 ?
* 10.128.0.0/11      Self                  200        2831679853 26697 4341 43012
28104 39181 51157 ?
* 17.205.88.0/21     Self                  200        2831679853 9726 19412 16490 6122
I
* 18.100.0.0/17      Self                  200        2831679853 9726 19412 40487
63999 13005 28592 5369 9061 I
* 21.0.0.0/8         Self                  200        2831679853 26697 56251 10277
21151 I
...
```

Both outputs show that the correct communities are set.

Verify if the customer routes are advertised to P2 and P3 peers.

```
[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.2 aspath-regex
"64512|64513|64514"

inet.0: 881 destinations, 896 routes (835 active, 0 holddown, 60 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24      Self                                    64512 I
* 172.31.1.0/24      Self                                    64512 I
* 172.31.2.0/24      Self                                    64512 I
* 172.31.3.0/24      Self                                    64512 I
* 172.31.4.0/24      Self                                    64512 I
* 172.31.5.0/24      Self                                    64512 I
* 172.31.6.0/24      Self                                    64512 I
* 172.31.7.0/24      Self                                    64512 I
* 172.31.8.0/24      Self                                    64512 I
* 172.31.9.0/24      Self                                    64512 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.10.0/24 Self                                  64512 I
* 172.31.11.0/24 Self                                  64512 I
* 172.31.12.0/24 Self                                  64512 I
* 172.31.13.0/24 Self                                  64512 I
* 172.31.14.0/24 Self                                  64512 I
* 172.31.15.0/24 Self                                  64512 I
* 172.31.16.0/24 Self                                  64513 I
* 172.31.17.0/24 Self                                  64513 I
* 172.31.18.0/24 Self                                  64513 I
* 172.31.19.0/24 Self                                  64513 I
* 172.31.20.0/24 Self                                  64513 I
* 172.31.21.0/24 Self                                  64513 I
* 172.31.22.0/24 Self                                  64513 I
* 172.31.23.0/24 Self                                  64513 I
* 172.31.24.0/24 Self                                  64513 I
* 172.31.25.0/24 Self                                  64513 I
* 172.31.26.0/24 Self                                  64513 I
* 172.31.27.0/24 Self                                  64513 I
* 172.31.28.0/24 Self                                  64513 I
* 172.31.29.0/24 Self                                  64513 I
* 172.31.30.0/24 Self                                  64513 I
* 172.31.31.0/24 Self                                  64513 I
* 172.31.32.0/24 Self                                  64514 I
* 172.31.33.0/24 Self                                  64514 I
...
```

```
[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.6 aspath-regex
"64512|64513|64514"

inet.0: 881 destinations, 896 routes (835 active, 0 holddown, 60 hidden)
  Prefix         Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24  Self                                   64512 I
* 172.31.1.0/24  Self                                   64512 I
* 172.31.2.0/24  Self                                   64512 I
* 172.31.3.0/24  Self                                   64512 I
* 172.31.4.0/24  Self                                   64512 I
* 172.31.5.0/24  Self                                   64512 I
* 172.31.6.0/24  Self                                   64512 I
* 172.31.7.0/24  Self                                   64512 I
* 172.31.8.0/24  Self                                   64512 I
* 172.31.9.0/24  Self                                   64512 I
* 172.31.10.0/24 Self                                   64512 I
* 172.31.11.0/24 Self                                   64512 I
* 172.31.12.0/24 Self                                   64512 I
* 172.31.13.0/24 Self                                   64512 I
* 172.31.14.0/24 Self                                   64512 I
* 172.31.15.0/24 Self                                   64512 I
* 172.31.16.0/24 Self                                   64513 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
...
```

Verify if the local summary prefix is sent to EBGP peers.

```
[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.2 172.30/16

inet.0: 881 destinations, 896 routes (835 active, 0 holddown, 60 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                                    I

[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.6 172.30/16

inet.0: 881 destinations, 896 routes (835 active, 0 holddown, 60 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                                    I
```

Verify the local-preference of the customer routes received by R3.

```
[edit]
lab@R3# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 881 destinations, 896 routes (835 active, 0 holddown, 60 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24           *[BGP/170] 01:48:31, localpref 90, from 172.30.5.4
                          AS path: (65002) 64512 I
                        > to 172.30.0.26 via ge-0/0/4.136
                        [BGP/170] 01:48:32, localpref 90, from 172.30.5.2
                          AS path: (65000 65003) 64512 I
                        > to 172.30.0.13 via ge-0/0/4.123
                          to 172.30.0.26 via ge-0/0/4.136
172.31.1.0/24           *[BGP/170] 00:54:06, localpref 400, from 172.30.5.2
                          AS path: (65000 65003 65002) 64512 I
                          Discard
172.31.2.0/24           *[BGP/170] 02:20:39, localpref 400, from 172.30.5.4
                          AS path: (65002) 64512 I
                        > to 172.30.0.26 via ge-0/0/4.136
...
```

Verify the remote triggered black hole routes with community rtbh.

```
[edit]
lab@R3# run show route protocol bgp community-name rtbh

inet.0: 881 destinations, 896 routes (835 active, 0 holddown, 60 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

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

172.31.1.0/24          *[BGP/170] 00:54:24, localpref 400, from 172.30.5.2
                          AS path: (65000 65003 65002) 64512 I
                          Discard
```

d) R4

R4 has no external BGP sessions but it receives external routes from R3 and R5.

Verify if unresolved routes are present in the routing table.

```
[edit]
lab@R4# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify if routes with a mask shorter than /8 and longer than /24 are present in routing table.

```
[edit]
lab@R4# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32  B 170   200          >172.30.0.21      2831679853 I
* 172.17.47.3/32  B 170   200          >172.30.0.21      2831679853 I
```

Verify the routing table for 0.0.0.0 with any mask.

```
[edit]
lab@R4# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if customer routes are received by R4.

```
[edit]
lab@R4# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 835 destinations, 866 routes (834 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24          *[BGP/170] 01:51:51, localpref 90, from 172.30.5.5
                          AS path: (65002) 64512 I
                        > to 172.30.0.30 via ge-0/0/4.145
172.31.1.0/24          *[BGP/170] 00:57:26, localpref 400, from 172.30.5.3
                          AS path: (65000 65003 65002) 64512 I
                        > to 172.30.0.21 via ge-0/0/4.134
                        [BGP/170] 00:57:26, localpref 400, from 172.30.5.5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                      AS path: (65002) 64512 I
                    > to 172.30.0.30 via ge-0/0/4.145
172.31.2.0/24       *[BGP/170] 02:33:16, localpref 400, from 172.30.5.5
                      AS path: (65002) 64512 I
                    > to 172.30.0.30 via ge-0/0/4.145
...
```

Verify the remote triggered black hole routes with community rtbh.

```
[edit]
lab@R4# run show route protocol bgp community-name rtbh

inet.0: 673 destinations, 704 routes (673 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24       *[BGP/170] 00:11:24, localpref 400, from 172.30.5.3
                      AS path: (65000 65003 65002) 64512 I, validation-state:
unverified
                      Discard
                    [BGP/170] 00:11:29, localpref 400, from 172.30.5.5
                      AS path: (65002) 64512 I, validation-state: unverified
                      Discard

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

inet6.0: 33 destinations, 33 routes (33 active, 0 holddown, 0 hidden)
```

e) R5

Verify if unresolved routes are present in the routing table.

```
[edit]
lab@R5# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify if routes with mask shorter than /8 and longer than /24 are present in routing table.

```
[edit]
lab@R5# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32   B 170   200  >172.30.0.34      (65001) 2831679853 I
* 172.17.47.3/32   B 170   200  >172.30.0.34      (65001) 2831679853 I
```

Verify the routing table for 0.0.0.0 with any mask.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if customer routes are advertised to the EBGP peers.

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.10 aspath-regex
"64512|64513|64514"

inet.0: 837 destinations, 853 routes (836 active, 0 holddown, 1 hidden)
  Prefix              Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24       Self                                    64512 I
* 172.31.1.0/24       Self                                    64512 I
* 172.31.2.0/24       Self                                    64512 I
* 172.31.3.0/24       Self                                    64512 I
* 172.31.4.0/24       Self                                    64512 I
* 172.31.5.0/24       Self                                    64512 I
* 172.31.6.0/24       Self                                    64512 I
* 172.31.7.0/24       Self                                    64512 I
* 172.31.8.0/24       Self                                    64512 I
* 172.31.9.0/24       Self                                    64512 I
* 172.31.10.0/24      Self                                    64512 I
* 172.31.11.0/24      Self                                    64512 I
* 172.31.12.0/24      Self                                    64512 I
* 172.31.13.0/24      Self                                    64512 I
* 172.31.14.0/24      Self                                    64512 I
* 172.31.15.0/24      Self                                    64512 I
* 172.31.16.0/24      Self                                    64513 I
* 172.31.17.0/24      Self                                    64513 I
* 172.31.18.0/24      Self                                    64513 I
* 172.31.19.0/24      Self                                    64513 I
* 172.31.20.0/24      Self                                    64513 I
* 172.31.21.0/24      Self                                    64513 I
* 172.31.22.0/24      Self                                    64513 I
* 172.31.23.0/24      Self                                    64513 I
* 172.31.24.0/24      Self                                    64513 I
* 172.31.25.0/24      Self                                    64513 I
* 172.31.26.0/24      Self                                    64513 I
* 172.31.27.0/24      Self                                    64513 I
* 172.31.28.0/24      Self                                    64513 I
* 172.31.29.0/24      Self                                    64513 I
* 172.31.30.0/24      Self                                    64513 I
* 172.31.31.0/24      Self                                    64513 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify if the local summary prefix is sent to EBGP peers.

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.10 172.30/16

inet.0: 837 destinations, 853 routes (836 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                                    I
```

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.14 172.30/16

inet.0: 837 destinations, 853 routes (836 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                                    I
```

Verify the local-preference of the customer routes received by R5.

```
[edit]
lab@R5# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 837 destinations, 853 routes (836 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24          *[BGP/170] 02:00:43, localpref 90, from 172.30.5.6
                          AS path: 64512 I
                        > to 172.30.0.34 via ae0.0
172.31.1.0/24          *[BGP/170] 01:06:18, localpref 400, from 172.30.5.6
                          AS path: 64512 I
                          Discard
172.31.2.0/24          *[BGP/170] 02:42:58, localpref 400, from 172.30.5.6
                          AS path: 64512 I
                        > to 172.30.0.34 via ae0.0
172.31.3.0/24          *[BGP/170] 02:42:58, localpref 400, from 172.30.5.6
                          AS path: 64512 I
                        > to 172.30.0.34 via ae0.0
172.31.4.0/24          *[BGP/170] 02:42:58, localpref 400, from 172.30.5.6
                          AS path: 64512 I
                        > to 172.30.0.34 via ae0.0
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify the remote triggered black hole routes with community rtbh.

```
[edit]
lab@R5# run show route protocol bgp community-name rtbh

inet.0: 837 destinations, 853 routes (836 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24          *[BGP/170] 01:06:29, localpref 400, from 172.30.5.6
                          AS path: 64512 I
                        Discard
```

Check the multipath load balancing.

```
[edit]
lab@R5# run show route aspath-regex 64514

inet.0: 837 destinations, 853 routes (836 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.32.0/24     *[BGP/170] 02:46:20, localpref 300, from 192.168.0.14
                      AS path: 64514 I
                    > to 192.168.0.10 via ge-0/0/5.303
                      to 192.168.0.14 via ge-0/0/5.304
                    [BGP/170] 02:46:12, localpref 300
                      AS path: 64514 I
                    > to 192.168.0.10 via ge-0/0/5.303
172.31.33.0/24     *[BGP/170] 02:46:20, localpref 300, from 192.168.0.14
                      AS path: 64514 I
                    > to 192.168.0.10 via ge-0/0/5.303
                      to 192.168.0.14 via ge-0/0/5.304
                    [BGP/170] 02:46:12, localpref 300
                      AS path: 64514 I
                    > to 192.168.0.10 via ge-0/0/5.303
...

[edit]
lab@R5# run show route forwarding-table destination 172.31.32.0/24
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
172.31.32.0/24     user     0                    ulst 262142    17
                                192.168.0.10     ucst   618     3 ge-0/0/5.303
                                192.168.0.14     ucst   616     3 ge-0/0/5.304

...
```


© 2020 Juniper Networks, Inc. All rights reserved.

f) R6

Verify if unresolved routes are present in the routing table.

```
[edit]
lab@R6# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify if routes with a mask shorter than /8 and longer than /24 are present in routing table.

```
[edit]
lab@R6# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32  B 170    200     >172.30.0.25      (65001) 2831679853 I
* 172.17.47.3/32  B 170    200     >172.30.0.25      (65001) 2831679853 I
```

Verify the routing table for a route 0.0.0.0 with any mask.

```
[edit]
lab@R6# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if customer routes are advertised to the EBGP peers.

```
[edit]
lab@R6# run show route advertising-protocol bgp 172.31.31.1 aspath-regex
"64512|64513|64514"

[edit]
lab@R6#
```

```
[edit]
lab@R6# run show route advertising-protocol bgp 192.168.0.18 aspath-regex
"64512|64513|64514"

inet.0: 838 destinations, 1080 routes (837 active, 0 holddown, 0 hidden)
  Prefix              Nexthop              MED     Lclpref    AS path
* 172.31.16.0/24      Self                 10                 64513 I
* 172.31.17.0/24      Self                 10                 64513 I
* 172.31.18.0/24      Self                 10                 64513 I
* 172.31.19.0/24      Self                 10                 64513 I
* 172.31.20.0/24      Self                 10                 64513 I
* 172.31.21.0/24      Self                 10                 64513 I
* 172.31.22.0/24      Self                 10                 64513 I
* 172.31.23.0/24      Self                 10                 64513 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.24.0/24     Self                10         64513 I
* 172.31.25.0/24     Self                10         64513 I
* 172.31.26.0/24     Self                10         64513 I
* 172.31.27.0/24     Self                10         64513 I
* 172.31.28.0/24     Self                10         64513 I
* 172.31.29.0/24     Self                10         64513 I
* 172.31.30.0/24     Self                10         64513 I
* 172.31.31.0/24     Self                10         64513 I
* 172.31.32.0/24     Self                10         64514 I
* 172.31.33.0/24     Self                10         64514 I
* 172.31.34.0/24     Self                10         64514 I
* 172.31.35.0/24     Self                10         64514 I
* 172.31.36.0/24     Self                10         64514 I
* 172.31.37.0/24     Self                10         64514 I
* 172.31.38.0/24     Self                10         64514 I
* 172.31.39.0/24     Self                10         64514 I
* 172.31.40.0/24     Self                10         64514 I
* 172.31.41.0/24     Self                10         64514 I
* 172.31.42.0/24     Self                10         64514 I
* 172.31.43.0/24     Self                10         64514 I
* 172.31.44.0/24     Self                10         64514 I
* 172.31.45.0/24     Self                10         64514 I
* 172.31.46.0/24     Self                10         64514 I
* 172.31.47.0/24     Self                10         64514 I
```

Here you can see that the routes advertised to C1-1 are assigned a MED of 10, because R6 should be preferred over R7 for inbound traffic from C1-1.

Verify if the local summary prefix is sent to EBGP peers.

```
[edit]
lab@R6# run show route advertising-protocol bgp 192.168.0.18 172.30/16

inet.0: 838 destinations, 1080 routes (837 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                 10                 I
```

Check if only the default route is advertised to C2.

```
[edit]
lab@R6# run show route advertising-protocol bgp 172.31.31.1


inet.0: 676 destinations, 693 routes (676 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    ?
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: AS path holds all AS numbers of the contributing routes.

Verify the local-preference of the customer routes received by R6.

```
[edit]
lab@R6# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 838 destinations, 1080 routes (837 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[BGP/170] 03:08:34, localpref 90
                      AS path: 64512 I
                    > to 192.168.0.18 via ge-0/0/5.305
                    [BGP/170] 03:07:10, localpref 90, from 172.30.5.7
                      AS path: (65003) 64512 I
                    > to 172.30.0.42 via ge-0/0/4.167
172.31.1.0/24      *[BGP/170] 03:08:34, localpref 400
                      AS path: 64512 I
                    > to 192.168.0.18 via ge-0/0/5.305
172.31.2.0/24      *[BGP/170] 03:08:34, localpref 400
                      AS path: 64512 I
                    > to 192.168.0.18 via ge-0/0/5.305
...
```

Verify the remote triggered black hole routes with community rtbh.

```
[edit]
lab@R6# run show route protocol bgp community-name rtbh

inet.0: 838 destinations, 1080 routes (837 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 03:08:47, localpref 400
                      AS path: 64512 I
                    > to 192.168.0.18 via ge-0/0/5.305
```

Verify multihop load balancing.

```
[edit]
lab@R6# run show route aspath-regex 64513

inet.0: 838 destinations, 1080 routes (837 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.16.0/24   *[BGP/170] 03:10:17, localpref 300, from 172.31.31.1
                      AS path: 64513 I
                    > to 192.168.0.22 via ge-0/0/5.306
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                      to 192.168.0.26 via ge-0/0/5.307
172.31.17.0/24    *[BGP/170] 03:10:17, localpref 300, from 172.31.31.1
                      AS path: 64513 I
                    > to 192.168.0.22 via ge-0/0/5.306
                      to 192.168.0.26 via ge-0/0/5.307
...
```

Verify that R6 is preferred for C1 outbound traffic.

```
[edit]
lab@R6# run show route 172.31.1/24

inet.0: 838 destinations, 1080 routes (837 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24          *[BGP/170] 03:16:03, localpref 400
                          AS path: 64512 I
                        > to 192.168.0.18 via ge-0/0/5.305
```

g) R7

Verify if unresolved routes are present in the routing table.

```
[edit]
lab@R7# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify if routes with mask shorter than /8 and longer than /24 are present in routing table.

```
[edit]
lab@R7# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32    B 170    200      >172.30.0.17      (65002 65001) 2831679853 I
* 172.17.47.3/32    B 170    200      >172.30.0.17      (65002 65001) 2831679853 I
```

Verify the routing table for a route 0.0.0.0 with any mask.

```
[edit]
lab@R7# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if customer routes are advertised to the EBGP peers.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.30 aspath-regex
"64512|64513|64514"

inet.0: 834 destinations, 1887 routes (834 active, 0 holddown, 548 hidden)
  Prefix          Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24   Self                                   64512 I
* 172.31.1.0/24   Self                                   64512 I
* 172.31.2.0/24   Self                                   64512 I
* 172.31.3.0/24   Self                                   64512 I
* 172.31.4.0/24   Self                                   64512 I
* 172.31.5.0/24   Self                                   64512 I
* 172.31.6.0/24   Self                                   64512 I
* 172.31.7.0/24   Self                                   64512 I
* 172.31.8.0/24   Self                                   64512 I
* 172.31.9.0/24   Self                                   64512 I
* 172.31.10.0/24  Self                                   64512 I
* 172.31.11.0/24  Self                                   64512 I
* 172.31.12.0/24  Self                                   64512 I
* 172.31.13.0/24  Self                                   64512 I
* 172.31.14.0/24  Self                                   64512 I
* 172.31.15.0/24  Self                                   64512 I
* 172.31.16.0/24  Self                                   64513 I
* 172.31.17.0/24  Self                                   64513 I
* 172.31.18.0/24  Self                                   64513 I
* 172.31.19.0/24  Self                                   64513 I
...

[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.34 aspath-regex
"64512|64513|64514"

inet.0: 834 destinations, 1887 routes (834 active, 0 holddown, 548 hidden)
  Prefix          Nexthop              MED     Lclpref    AS path
* 172.31.16.0/24  Self                 20                 64513 I
* 172.31.17.0/24  Self                 20                 64513 I
* 172.31.18.0/24  Self                 20                 64513 I
* 172.31.19.0/24  Self                 20                 64513 I
* 172.31.20.0/24  Self                 20                 64513 I
* 172.31.21.0/24  Self                 20                 64513 I
* 172.31.22.0/24  Self                 20                 64513 I
* 172.31.23.0/24  Self                 20                 64513 I
* 172.31.24.0/24  Self                 20                 64513 I
* 172.31.25.0/24  Self                 20                 64513 I
* 172.31.26.0/24  Self                 20                 64513 I
* 172.31.27.0/24  Self                 20                 64513 I
* 172.31.28.0/24  Self                 20                 64513 I
* 172.31.29.0/24  Self                 20                 64513 I
* 172.31.30.0/24  Self                 20                 64513 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.31.0/24  Self                  20                 64513 I
* 172.31.32.0/24  Self                  20                 64514 I
* 172.31.33.0/24  Self                  20                 64514 I
* 172.31.34.0/24  Self                  20                 64514 I
* 172.31.35.0/24  Self                  20                 64514 I
* 172.31.36.0/24  Self                  20                 64514 I
* 172.31.37.0/24  Self                  20                 64514 I
* 172.31.38.0/24  Self                  20                 64514 I
* 172.31.39.0/24  Self                  20                 64514 I
* 172.31.40.0/24  Self                  20                 64514 I
* 172.31.41.0/24  Self                  20                 64514 I
* 172.31.42.0/24  Self                  20                 64514 I
* 172.31.43.0/24  Self                  20                 64514 I
* 172.31.44.0/24  Self                  20                 64514 I
* 172.31.45.0/24  Self                  20                 64514 I
* 172.31.46.0/24  Self                  20                 64514 I
* 172.31.47.0/24  Self                  20                 64514 I
```

Verify if the local summary prefix is sent to the EBGP peers.

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.30 172.30/16

inet.0: 834 destinations, 1887 routes (834 active, 0 holddown, 548 hidden)
  Prefix                 Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16          Self                                    I
* 172.30.128.0/17        Self                                    I
```

Note: R7 is advertising the second half of the internal routes as is for equal load balancing purposes.

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.34 172.30/16

inet.0: 834 destinations, 1887 routes (834 active, 0 holddown, 548 hidden)
  Prefix                 Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16          Self                 20                 I
```

Here you can see that the route to C1-1 is advertised with a MED of 20, making R6 preferred.

Verify that the local-preference of the customer routes received by R7.

```
[edit]
lab@R7# run show route protocol bgp aspath-regex ".* [64512-64514]"
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 834 destinations, 1887 routes (834 active, 0 holddown, 548 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[BGP/170] 03:32:17, localpref 90
                      AS path: 64512 I
                    > to 192.168.0.34 via ge-0/0/5.309
                    [BGP/170] 02:48:36, localpref 90, from 172.30.5.6
                      AS path: (65002) 64512 I
                    > to 172.30.0.41 via ge-0/0/4.167
172.31.1.0/24      *[BGP/170] 01:54:10, localpref 400, from 172.30.5.6
                      AS path: (65002) 64512 I
                      Discard
                    [BGP/170] 03:32:17, localpref 300
                      AS path: 64512 I
                    > to 192.168.0.34 via ge-0/0/5.309
172.31.2.0/24      *[BGP/170] 03:30:46, localpref 400, from 172.30.5.6
                      AS path: (65002) 64512 I
                    > to 172.30.0.41 via ge-0/0/4.167
                    [BGP/170] 02:48:36, localpref 400, from 172.30.5.8
                      AS path: (65000 65001 65002) 64512 I
                    > to 172.30.0.41 via ge-0/0/4.167
                    [BGP/170] 03:32:17, localpref 300
...
```

Verify the remote triggered black hole routes with community rtbh.

```
[edit]
lab@R7# run show route protocol bgp community-name rtbh

inet.0: 834 destinations, 1887 routes (834 active, 0 holddown, 548 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 01:54:26, localpref 400, from 172.30.5.6
                      AS path: (65002) 64512 I
                      Discard
                    [BGP/170] 03:32:33, localpref 300
                      AS path: 64512 I
                    > to 192.168.0.34 via ge-0/0/5.309
```

Verify that R6 is preferred for C1 outbound traffic.

```
[edit]
lab@R7# run show route 172.31.1/24

inet.0: 834 destinations, 1887 routes (834 active, 0 holddown, 548 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 01:58:44, localpref 400, from 172.30.5.6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  AS path: (65002) 64512 I
  Discard
[BGP/170] 03:36:51, localpref 300
  AS path: 64512 I
> to 192.168.0.34 via ge-0/0/5.309
```

Verify if routes not native to P1 are accepted.

```
[edit]
lab@R7# run show route receive-protocol bgp 192.168.0.30 aspath-regex "110047427
.+"

inet.0: 834 destinations, 1887 routes (834 active, 0 holddown, 548 hidden)

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

inet6.0: 50 destinations, 50 routes (50 active, 0 holddown, 0 hidden)
```

Verify if the routes to P1 are advertised with community no-export.

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.30 172.30/16 extensive

inet.0: 834 destinations, 1887 routes (834 active, 0 holddown, 548 hidden)
* 172.30.0.0/16 (1 entry, 1 announced)
 BGP group P1-2 type External
     Nexthop: Self
     AS path: [65003] I (LocalAgg)

* 172.30.128.0/17 (1 entry, 1 announced)
 BGP group P1-2 type External
     Nexthop: Self
     AS path: [65003] I (LocalAgg)
     Communities: no-export
```

h) R8

Verify if unresolved routes are present in the routing table.

```
[edit]
lab@R8# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify if routes with mask shorter than /8 and longer than /24 are present in routing table.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32  B 170   200     >172.30.0.9        (65000 65001) 2831679853 I
* 172.17.47.3/32  B 170   200     >172.30.0.9        (65000 65001) 2831679853 I
```

Verify the routing table for route 0.0.0.0 with any mask.

```
[edit]
lab@R8# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if customer routes are advertised to the EBGP peers.

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.38 aspath-regex
"64512|64513|64514"

inet.0: 832 destinations, 1660 routes (832 active, 0 holddown, 0 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24      Self                                    64512 I
* 172.31.1.0/24      Self                                    64512 I
* 172.31.2.0/24      Self                                    64512 I
* 172.31.3.0/24      Self                                    64512 I
* 172.31.4.0/24      Self                                    64512 I
* 172.31.5.0/24      Self                                    64512 I
* 172.31.6.0/24      Self                                    64512 I
* 172.31.7.0/24      Self                                    64512 I
* 172.31.8.0/24      Self                                    64512 I
* 172.31.9.0/24      Self                                    64512 I
* 172.31.10.0/24     Self                                    64512 I
* 172.31.11.0/24     Self                                    64512 I
* 172.31.12.0/24     Self                                    64512 I
* 172.31.13.0/24     Self                                    64512 I
* 172.31.14.0/24     Self                                    64512 I
* 172.31.15.0/24     Self                                    64512 I
* 172.31.16.0/24     Self                                    64513 I
* 172.31.17.0/24     Self                                    64513 I
* 172.31.18.0/24     Self                                    64513 I
* 172.31.19.0/24     Self                                    64513 I
* 172.31.20.0/24     Self                                    64513 I
* 172.31.21.0/24     Self                                    64513 I
* 172.31.22.0/24     Self                                    64513 I
* 172.31.23.0/24     Self                                    64513 I
* 172.31.24.0/24     Self                                    64513 I
* 172.31.25.0/24     Self                                    64513 I
* 172.31.26.0/24     Self                                    64513 I
* 172.31.27.0/24     Self                                    64513 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.28.0/24    Self                                  64513 I
* 172.31.29.0/24    Self                                  64513 I
* 172.31.30.0/24    Self                                  64513 I
* 172.31.31.0/24    Self                                  64513 I
* 172.31.32.0/24    Self                                  64514 I
* 172.31.33.0/24    Self                                  64514 I
* 172.31.34.0/24    Self                                  64514 I
* 172.31.35.0/24    Self                                  64514 I
* 172.31.36.0/24    Self                                  64514 I
* 172.31.37.0/24    Self                                  64514 I
* 172.31.38.0/24    Self                                  64514 I
* 172.31.39.0/24    Self                                  64514 I
* 172.31.40.0/24    Self                                  64514 I
* 172.31.41.0/24    Self                                  64514 I
* 172.31.42.0/24    Self                                  64514 I
* 172.31.43.0/24    Self                                  64514 I
* 172.31.44.0/24    Self                                  64514 I
* 172.31.45.0/24    Self                                  64514 I
* 172.31.46.0/24    Self                                  64514 I
* 172.31.47.0/24    Self                                  64514 I
```

Verify if the local summary prefix is sent to EBGP peers.

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.38 172.30/16

inet.0: 832 destinations, 1660 routes (832 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                                    I
* 172.30.0.0/17           Self                                    I
```

Note: R8 is advertising the first half of the internal routes as is requested for equal load balancing.

Verify the local-preference of the customer routes received by R8.

```
[edit]
lab@R8# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 832 destinations, 1660 routes (832 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24       *[BGP/170] 03:18:07, localpref 90, from 172.30.5.7
                       AS path: 64512 I
                     > to 172.30.0.45 via ge-0/0/4.178
172.31.1.0/24       *[BGP/170] 02:02:37, localpref 400, from 172.30.5.7
                       AS path: (65002) 64512 I
                       Discard
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.31.2.0/24      *[BGP/170] 03:17:16, localpref 400, from 172.30.5.1
                      AS path: (65000 65001 65002) 64512 I
                    > to 172.30.0.37 via ge-0/0/4.158
                    [BGP/170] 03:18:07, localpref 400, from 172.30.5.7
                      AS path: (65002) 64512 I
                    > to 172.30.0.37 via ge-0/0/4.158
...
```

Verify the remote triggered black hole routes with community rtbh.

```
[edit]
lab@R8# run show route protocol bgp community-name rtbh

inet.0: 832 destinations, 1660 routes (832 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 02:02:50, localpref 400, from 172.30.5.7
                      AS path: (65002) 64512 I
                      Discard
```

Verify if the P1, P2 and P3 routes are preferred to IX routes.

```
[edit]
lab@R8# run show route 172.17.0.0/24

inet.0: 832 destinations, 1660 routes (832 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.0.0/24      *[BGP/170] 03:18:56, localpref 210
                      AS path: 110047427 I
                    > to 192.168.0.38 via ge-0/0/5.310
```

R8 receives the route only from P1. It has a higher local-preference. Therefore, R1 does not advertise the route learned from IX peers to R8.

Verify if routes not native to P1 are accepted.

```
[edit]
lab@R8# run show route receive-protocol bgp 192.168.0.38 aspath-regex "110047427
.+"

inet.0: 832 destinations, 1660 routes (832 active, 0 holddown, 0 hidden)

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

inet6.0: 50 destinations, 52 routes (50 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify if the routes to P1 are advertised with community no-export.

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.38 172.30/16 extensive

inet.0: 832 destinations, 1660 routes (832 active, 0 holddown, 0 hidden)
* 172.30.0.0/16 (1 entry, 1 announced)
 BGP group P1-1 type External
     Nexthop: Self
     AS path: [65003] I (LocalAgg)

* 172.30.0.0/17 (1 entry, 1 announced)
 BGP group P1-1 type External
     Nexthop: Self
     AS path: [65003] I (LocalAgg)
     Communities: no-export
```

Task 4: IBGP and Route Reflection

The last Task of Chapter Four is focused on another method to reduce IBGP full mesh connectivity. You must redesign the network by introducing Route Reflection instead of confederations. You must delete all IBGP configuration including the confederations and configure the sessions using a route reflector.

Route Reflector overview

A Router Reflector (RR) is basically a dedicated router which can re-advertise IBGP routes from one IBGP neighbor to another. Other routers are clients to the route reflector, they do not require any specific configuration. On the route reflector under the IBGP group, a cluster ID is configured to identify route reflector functionality. All neighbors configured under that group are treated are clients.

A route reflector does not change any of the BGP attributes. They are solely used to reflect the routing information to all clients. Because the AS is not changed as with Confederations, new attributes are introduced to prevent loops within AS.

- Cluster list ID – this attribute is analogous to the AS path attribute and is used to prevent loops. If a Route reflector receives route with its own cluster ID in the list, the route is discarded. Additionally, the routers in the network can use this attribute for BGP best path selection algorithm. BGP prefers a route with the shortest Cluster list. Cluster ID should be unique in the network. It is added when a RR re-advertises the routes to the clients and non-clients.

© 2020 Juniper Networks, Inc. All rights reserved.

- Originator ID – this attribute contains the Router ID of the first router which advertises the route to the AS. It is also used as a loop prevention mechanism in case the Cluster list does not prevent a loop.

Because a route reflector is an IBGP router it can change any BGP attribute with routing policy. You have to be careful with changing attributes on the route reflector though. For example, if you apply next-hop self on route reflector all the traffic will be attract by that route reflector. Usually the route reflector is capable of carrying a lot of prefixes in the RIB and having less power in the forwarding engine.

1) Delete IBGP configuration on all routers including the confederation.

Repeat this on all eight devices.

```
[edit]
lab@R1# delete routing-options confederation

[edit]
lab@R1# delete protocols bgp group ibgp

[edit]
lab@R1# delete protocols bgp group cbgp

[edit]
lab@R1# commit
commit complete
```

2) Configure additional interfaces on R1 and R2.

For the route reflector we will be using a dedicated vMX router. It is connected to R1 and R2 as shown in the beginning of the task.

a) R1.

```
[edit]
lab@R1# set interfaces ge-0/0/4.206 vlan-id 206

[edit]
lab@R1# set interfaces ge-0/0/4.206 family inet address 172.30.0.65/30

[edit]
lab@R1# set interfaces ge-0/0/4.206 family iso

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Interface is configured with `family iso`. Because you previously configured "interface all" under `protocols isis level 1`, nothing should be changed to add this interface for IS-IS routing.

b) R2

```
[edit]
lab@R2# set interfaces ge-0/0/4.207 vlan-id 207

[edit]
lab@R2# set interfaces ge-0/0/4.207 family inet address 172.30.0.69/30

[edit]
lab@R2# set interfaces ge-0/0/4.207 family iso

[edit]
lab@R2# commit
commit complete
```

c) Verify the connectivity.

```
[edit]
lab@R1# run ping 172.30.0.66 count 3
PING 172.30.0.66 (172.30.0.66): 56 data bytes
64 bytes from 172.30.0.66: icmp_seq=0 ttl=64 time=3.431 ms
64 bytes from 172.30.0.66: icmp_seq=1 ttl=64 time=2.805 ms
64 bytes from 172.30.0.66: icmp_seq=2 ttl=64 time=3.397 ms

--- 172.30.0.66 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 2.805/3.211/3.431/0.287 ms
```

```
[edit]
lab@R2# run ping 172.30.0.70 count 3
PING 172.30.0.70 (172.30.0.70): 56 data bytes
64 bytes from 172.30.0.70: icmp_seq=0 ttl=64 time=24.673 ms
64 bytes from 172.30.0.70: icmp_seq=1 ttl=64 time=1.138 ms
64 bytes from 172.30.0.70: icmp_seq=2 ttl=64 time=1.229 ms

--- 172.30.0.70 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.138/9.013/24.673/11.073 ms
```

3) Configure the autonomous system.

© 2020 Juniper Networks, Inc. All rights reserved.

All routers should be configured with a global AS instead of the one configured for the confederation.

Repeat this on all eight devices.

```
[edit]
lab@R1# set routing-options autonomous-system 54591
```

4) Configure IBGP configuration.

Only one IBGP neighbor, the route reflector, should be configured on every router in the network. As stated earlier, the configuration on the clients is the same as with any other IBGP neighbor. You should add additional configuration for authentication and the requirement for BFD. RTBH policy is applied in the import direction.

Repeat this on all eight devices. Change local-address to match respective device.

```
[edit]
lab@R1# set protocols bgp group ibgp type internal

[edit]
lab@R1# set protocols bgp group ibgp local-address 172.30.5.1

[edit]
lab@R1# set protocols bgp group ibgp import rtbh

[edit]
lab@R1# set protocols bgp group ibgp authentication-key jncie

[edit]
lab@R1# set protocols bgp group ibgp bfd-liveness-detection minimum-interval 3000

[edit]
lab@R1# set protocols bgp group ibgp neighbor 172.30.5.41
```

Make a small alteration to the protect RE firewall filter to allow multi-hop BFD session's UDP messages to be accepted on port 4784.

```
[edit]
lab@R1# set firewall family inet filter protect-re term bfd from port 4784

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

The configuration on all routers are is the same except the local-address. You can use any application for text editing to change the local-address for each router and then load merge terminal to apply the configuration.

5) Applying next-hop self on all routers except R1, R2, and R4.

The policy configuration is available from the previous confederation exercise. You only have to re-apply.

```
[edit]
lab@R3# set protocols bgp group ibgp export nhs

[edit]
lab@R3# commit
commit complete
```

6) Configure the router reflector.

a) Enabling family iso and Configure NET address.

```
[edit]
lab@route-reflector# set interfaces ge-0/0/1.206 family iso

[edit]
lab@route-reflector# set interfaces ge-0/0/1.207 family iso

[edit]
lab@route-reflector# set interface lo0.0 family iso address 49.0001.1720.3000.5041.00
```

b) Configure IS-IS protocol.

```
[edit]
lab@route-reflector# set protocols isis level 2 disable

[edit]
lab@route-reflector# set protocols isis level 1 authentication-type md5

[edit]
lab@route-reflector# set protocols isis level 1 authentication-key jncie

[edit]
lab@route-reflector# set protocols isis interface all point-to-point

[edit]
lab@route-reflector# set protocols isis interface all family inet bfd-liveness-
detection multiplier 3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# set protocols isis interface all family inet bfd-liveness-
detection minimum-interval 2000
```

c) Configure the autonomous system.

```
[edit]
lab@route-reflector# set routing-options autonomous-system 54591
```

d) Configure IBGP.

You have requirement to create two clusters and any client should be a member of only one cluster. Statement `cluster 0.0.0.1` configures the neighbors in the group to be in cluster.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 type internal

[edit]
lab@route-reflector# set protocols bgp group cluster-1 local-address 172.30.5.41

[edit]
lab@route-reflector# set protocols bgp group cluster-1  authentication-key jncie

[edit]
lab@route-reflector# set protocols bgp group cluster-1 cluster 0.0.0.1

[edit]
lab@route-reflector# set protocols bgp group cluster-1 bfd-liveness-detection minimum-
interval 3000

[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 172.30.5.1

[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 172.30.5.6

[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 172.30.5.7

[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 172.30.5.8
```

Configure the second cluster group with `cluster 0.0.0.2`

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 type internal

[edit]
lab@route-reflector# set protocols bgp group cluster-2 local-address 172.30.5.41

[edit]
lab@route-reflector# set protocols bgp group cluster-2 authentication-key jncie

[edit]
lab@route-reflector# set protocols bgp group cluster-2 cluster 0.0.0.2

[edit]
lab@route-reflector# set protocols bgp group cluster-2 bfd-liveness-detection minimum-
interval 3000

[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 172.30.5.2

[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 172.30.5.3

[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 172.30.5.4

[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 172.30.5.5
```

To simulate that the RR does not support a 4-byte AS number disable this through configuration.

Note: The `disable-4byte-as` option is a hidden command and must be typed out completely.

```
[edit]
lab@route-reflector# set protocols bgp disable-4byte-as

[edit]
lab@route-reflector# commit
commit complete
```

e) Verify the operation of both RR clusters.

```
[edit]
lab@route-reflector# run show bgp summary
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Groups: 2 Peers: 8 Down peers: 0
Table            Tot Paths  Act Paths Suppressed  History Damp State    Pending
inet.0                 984        600          0        0          0          0
Peer                    AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
172.30.5.1          54591        396        272       0       0        8:29 383/383/0
0/0/0
172.30.5.2          54591        379        492       0       0        2:52 0/383/0
0/0/0
172.30.5.3          54591        124        352       0       0        2:48 153/153/0
0/0/0
172.30.5.4          54591          1        458       0       0          14 0/0/0
0/0/0
172.30.5.5          54591          8        462       0       0        2:40 16/16/0
0/0/0
172.30.5.6          54591         24        606       0       0        8:21 31/32/0
0/0/0
172.30.5.7          54591         25        610       0       0        8:21 1/1/0
0/0/0
172.30.5.8          54591         21        612       0       0        8:29 16/16/0
0/0/0
```

All bfd sessions are established.

```
[edit]
lab@route-reflector# run show bfd session
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.30.0.65              Up        ge-0/0/1.206   9.000     3.000        3
172.30.0.69              Up        ge-0/0/1.207   9.000     3.000        3
172.30.5.1               Up                       9.000     3.000        3
172.30.5.2               Up                       9.000     3.000        3
172.30.5.3               Up                       9.000     3.000        3
172.30.5.4               Up                       9.000     3.000        3
172.30.5.5               Up                       9.000     3.000        3
172.30.5.6               Up                       9.000     3.000        3
172.30.5.7               Up                       9.000     3.000        3
172.30.5.8               Up                       9.000     3.000        3

10 sessions, 10 clients
Cumulative transmit rate 3.3 pps, cumulative receive rate 3.3 pps
```

All 8 client routers sessions are established.

```
[edit]
lab@route-reflector# run show bgp group brief
Group Type: Internal    AS: 54591                    Local AS: 54591
  Name: cluster-1       Index: 0                     Flags: <>
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Options: <Cluster>
  Holdtime: 0
  Total peers: 4        Established: 4
  172.30.5.1+179
  172.30.5.6+179
  172.30.5.7+179
  172.30.5.8+60097
  inet.0: 431/432/0

Group Type: Internal    AS: 54591                    Local AS: 54591
  Name: cluster-2       Index: 2                     Flags: <>
  Options: <Cluster>
  Holdtime: 0
  Total peers: 4        Established: 4
  172.30.5.2+179
  172.30.5.3+179
  172.30.5.4+179
  172.30.5.5+179
  inet.0: 169/552/0

Groups: 2  Peers: 8    External: 0    Internal: 8    Down peers: 0   Flaps: 0
Table    Tot Paths  Act Paths Suppressed  History Damp State  Pending
inet.0         984        600          0      0          0          0
```

Both IBGP groups are with option Cluster activated.

R1

```
[edit]
lab@R1# run show bgp neighbor 172.30.5.41
Peer: 172.30.5.41+63155 AS 54591 Local: 172.30.5.1+179 AS 54591
  Group: ibgp                   Routing-Instance: master
  Forwarding routing-instance: master
  Type: Internal    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Import: [ rtbh ]
  Options: <Preference LocalAddress AuthKey Refresh>
  Options: <BfdEnabled>
  Authentication key is configured
  Local Address: 172.30.5.1 Holdtime: 90 Preference: 170
  Number of flaps: 5
  Last flap event: RecvNotify
  Error: 'Cease' Sent: 0 Recv: 5
  Peer ID: 172.30.5.41     Local ID: 172.30.5.1        Active Holdtime: 90
  Keepalive Interval: 30         Group index: 1    Peer index: 0    SNMP index: 4
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: enabled, up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer does not support LLGR Restarter functionality
  Peer does not support 4 byte AS extension
  Peer does not support Addpath
  Table inet.0 Bit: 20001
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              217
    Received prefixes:            217
    Accepted prefixes:            217
    Suppressed due to damping:    0
    Advertised prefixes:          383
  ...
```

Note: For demonstration purposes, imagine the route reflector is loaded with old Junos software and does not support 4 byte AS numbers. Here you can see how interoperability between routers supporting only 2 byte AS numbers and routers supporting 2 and 4 byte AS numbers is achieved.

R1 is advertising the route 3.49.128.0/19 to the RR. That route is received from the IX peers and contains 4 byte As in the AS path attribute.

```
[edit]
lab@R1# run show route advertising-protocol bgp 172.30.5.41 3.49.128.0/19

inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)
  Prefix          Nexthop         MED   Lclpref    AS path
* 3.49.128.0/19   192.168.1.3           100        1620 110047427 56422 7923
42335 17359 9293 38097 38522 61520 I
```

Verify the RR is receiving the route and reflects it to the clients. This output is taken from a sample device that does not support or understand the 4-byte AS numbers and is for demonstration purposes only to show that it truly does not recognize the attributes.

© 2020 Juniper Networks, Inc. All rights reserved.

RR sample output:

```
[edit]
lab@route-reflector# run show route 3.49.128.0/19 extensive

inet.0: 667 destinations, 1051 routes (667 active, 0 holddown, 0 hidden)
3.49.128.0/19 (2 entries, 1 announced)
TSI:
KRT in-kernel 3.49.128.0/19 -> {indirect(262143)}
Page 0 idx 0 Type 1 val 8a051f8
    Nexthop: 192.168.1.3
    Localpref: 100
    AS path: [54591] 1620 23456 56422 7923 42335 17359 9293 38097 38522 61520 I
    Unrecognized Attributes: 45 bytes
    Attr flags e0 code 11: 02 0a 00 00 06 54 06 8f 30 c3 00 00 dc 66 00 00 1e f3
00 00 a5 5f 00 00 43 cf 00 00 24 4d 00 00 94 d1 00 00 96 7a 00 00 f0 50
    Communities: 54591:1620
    Cluster ID: 0.0.0.1
    Originator ID: 172.30.5.1
    Advertise: 0000000e
Page 0 idx 1 Type 1 val 8a54168
    Nexthop: 192.168.1.3
    Localpref: 100
    AS path: [54591] 1620 23456 56422 7923 42335 17359 9293 38097 38522 61520 I
    Unrecognized Attributes: 45 bytes
    Attr flags e0 code 11: 02 0a 00 00 06 54 06 8f 30 c3 00 00 dc 66 00 00 1e f3
00 00 a5 5f 00 00 43 cf 00 00 24 4d 00 00 94 d1 00 00 96 7a 00 00 f0 50
    Communities: 54591:1620
    Cluster ID: 0.0.0.1
    Originator ID: 172.30.5.1
Path 3.49.128.0 from 172.30.5.1 Vector len 4.  Val: 0 1
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1532
                Source: 172.30.5.1
                Next hop type: Router, Next hop index: 262142
                Next hop: 172.30.0.65 via ge-0/0/1.206
                Next hop: 172.30.0.69 via ge-0/0/1.207, selected
                Protocol next hop: 192.168.1.3
                Indirect next hop: 89cc000 262143
                State: <Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 10:22      Metric2: 20
                Task: BGP_54591.172.30.5.1+179
                Announcement bits (3): 0-KRT 3-BGP RT Background 4-Resolve tree 2
                AS path: 1620 23456 56422 7923 42335 17359 9293 38097 38522 61520
I Unrecognized Attributes: 45 bytes
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                AS path:  Attr flags e0 code 11: 02 0a 00 00 06 54 06 8f 30 c3 00
00 dc 66 00 00 1e f3 00 00 a5 5f 00 00 43 cf 00 00 24 4d 00 00 94 d1 00 00 96 7a
00 00 f0 50
                Communities: 54591:1620
                Localpref: 100
                Router ID: 172.30.5.1
                Indirect next hops: 1
                        Protocol next hop: 192.168.1.3 Metric: 20
                        Indirect next hop: 89cc000 262143
                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 172.30.0.65 via ge-0/0/1.206
                                Next hop: 172.30.0.69 via ge-0/0/1.207
                        192.168.1.0/24 Originating RIB: inet.0
                          Metric: 20                          Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 172.30.0.65 via ge-0/0/1.206
                                Nexthop: 172.30.0.69 via ge-0/0/1.207
        BGP     Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1532
                Source: 172.30.5.2
                Next hop type: Router, Next hop index: 262142
                Next hop: 172.30.0.65 via ge-0/0/1.206
                Next hop: 172.30.0.69 via ge-0/0/1.207, selected
                Protocol next hop: 192.168.1.3
                Indirect next hop: 89cc000 262143
                State: <NotBest Int Ext>
                Inactive reason: Not Best in its group - Router ID
                Local AS: 54591 Peer AS: 54591
                Age: 4:45       Metric2: 20
                Task: BGP_54591.172.30.5.2+179
                AS path: 1620 23456 56422 7923 42335 17359 9293 38097 38522 615
                AS path:  Attr flags e0 code 11: 02 0a 00 00 06 54 06 8f 30 c3
00 24 4d 00 00 94 d1 00 00 96 7a 00 00 f0 50
                Communities: 54591:1620
                Localpref: 100
                Router ID: 172.30.5.2
                Indirect next hops: 1
                        Protocol next hop: 192.168.1.3 Metric: 20
                        Indirect next hop: 89cc000 262143
                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 172.30.0.65 via ge-0/0/1.206
                                Next hop: 172.30.0.69 via ge-0/0/1.207
                        192.168.1.0/24 Originating RIB: inet.0
                          Metric: 20                          Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 172.30.0.65 via ge-0/0/1.206
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                                Nexthop: 172.30.0.69 via ge-0/0/1.207
```

In this Sample the 4 byte AS in the AS path is replaced by the reserved AS number 23456.
The real AS path with the 4 byte AS number is sent as additional attribute AS4_PATH.
This attribute is not recognized by RR but it is passed to other clients in transitive manner.

Review the route received on R8 from the RR device.

```
[edit]
lab@R8# run show route 3.49.128.0/19 extensive

inet.0: 835 destinations, 1221 routes (673 active, 0 holddown, 548 hidden)
3.49.128.0/19 (2 entries, 1 announced)
TSI:
KRT in-kernel 3.49.128.0/19 -> {indirect(262142)}
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Address: 0x155c64c
                Next-hop reference count: 1149
                Source: 172.30.5.41
                Next hop type: Router, Next hop index: 593
                Next hop: 172.30.0.9 via ge-0/0/4.118, selected
                Protocol next hop: 192.168.1.3
                Indirect next hop: 169c488 262142
                State: <Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 15:43      Metric2: 20
                Task: BGP_54591.172.30.5.41+179
                Announcement bits (2): 0-KRT 6-Resolve tree 4
                AS path:
                AS2 PA[10]: 1620 AS_TRANS 56422 7923 42335 17359 9293 38097 38522
61520
                AS4 PA[10]: 1620 110047427 56422 7923 42335 17359 9293 38097
38522 61520
                Merged[10]: 1620 110047427 56422 7923 42335 17359 9293 38097
38522 61520 I (Originator)
                Cluster list:  0.0.0.1
                Originator ID: 172.30.5.1
                Communities: 54591:1620
                Accepted
                Localpref: 100
                Router ID: 172.30.5.41
                Indirect next hops: 1
                        Protocol next hop: 192.168.1.3 Metric: 20
                        Indirect next hop: 169c488 262142
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 172.30.0.9 via ge-0/0/4.118
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
192.168.1.0/24 Originating RIB: inet.0
  Metric: 20                          Node path count: 1
  Forwarding nexthops: 1
        Nexthop: 172.30.0.9 via ge-0/0/4.118
```

The route is received on R8 and both attributes are recognized. R8 is merging both AS paths and the result is the AS path advertised from R1 containing the 4 byte AS number. Cluster ID attribute is showing that this route is readvertised by the RR.

7) Verify the routing policy operation with RR in the network.

a) R1.

Verify the BGP session status with the RR.

```
[edit]
lab@R1# run show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    1345        600          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591        444        430       0       0       18:24
217/217/217/0        0/0/0/0
192.168.1.3             1620        609        531       0       0     1:05:32
383/564/402/0        0/0/0/0
192.168.1.4             1620        609        529       0       0     1:05:28
0/564/402/0          0/0/0/0
```

Routers have just one IBGP session with the RR and thus receive only the best routes from it.

Verify if unresolved routes are present.

```
[edit]
lab@R1# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify the next-hop for the routes advertised to RR.

```
[edit]
lab@R1# run show route advertising-protocol bgp 172.30.5.41
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)
  Prefix          Nexthop              MED     Lclpref    AS path
* 1.64.0.0/10     192.168.1.3                  100        1620 61671 I
* 1.84.160.0/20   192.168.1.3                  100        1620 33112 I
* 1.96.0.0/11     192.168.1.3                  100        1620 33112 63164 40776
51777 I
* 1.161.192.0/21  192.168.1.3                  100        1620 33112 30404 32138
45045 I
--- (more) ---
```

Routes from R1 is advertised with the next hop of the IX routers. Because they are advertised in IS-IS all routers can resolve the routes without a next-hop self-policy.

Verify for routes with mask shorter than /8 and longer than /24

```
[edit]
lab@R1# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32  B 170    200          >172.30.0.2       2831679853 I
* 172.17.47.3/32  B 170    200          >172.30.0.2       2831679853 I
```

Only shows the routes received from P3-1.

Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@R1# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and send to RR.

```
lab@R1# run show route advertising-protocol bgp 172.30.5.41 aspath-regex "1620
.*" community-name IX

inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)
  Prefix          Nexthop         MED     Lclpref    AS path
* 1.64.0.0/10     192.168.1.3             100        1620 61671 I
* 1.84.160.0/20   192.168.1.3             100        1620 33112 I
* 1.96.0.0/11     192.168.1.3             100        1620 33112 63164 40776 51777 I
* 1.161.192.0/21  192.168.1.3             100        1620 33112 30404 32138 45045 I
* 1.166.89.0/24   192.168.1.3             100        1620 110047427 56422 47123
34908 57085 3090 19128 21300 I
* 1.176.0.0/12    192.168.1.3             100        1620 33112 49129 16320 52954 I
...
```

Verify if customer routes are advertised to the IX peers.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.1.3 aspath-regex
"64512|64513|64514"

inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)
  Prefix              Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24       Self                                    64512 I
* 172.31.1.0/24       Self                                    64512 I
* 172.31.2.0/24       Self                                    64512 I
* 172.31.3.0/24       Self                                    64512 I
* 172.31.4.0/24       Self                                    64512 I
* 172.31.5.0/24       Self                                    64512 I
* 172.31.6.0/24       Self                                    64512 I
...
```

Verify if the local summary route is advertised to IX peers.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.1.3 172.30/16

inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                                    I
```

Verify the local preference of the customer routes received by R1.

```
[edit]
lab@R1# run show route protocol bgp aspath-regex "64512|64513|64514"

inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[BGP/170] 00:01:26, localpref 90, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.2 via ae0.0
172.31.1.0/24      *[BGP/170] 00:01:26, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                       Discard
172.31.2.0/24      *[BGP/170] 00:01:26, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.2 via ae0.0
                       to 172.30.0.6 via ge-0/0/4.114
                       to 172.30.0.10 via ge-0/0/4.118
...
```

Verify the remote triggered black hole routes.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show route protocol bgp community-name rtbh

inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 00:02:19, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                       Discard
```

One route is received from C1 and it is set with next hop of discard.

Verify if the P1, P2 and P3 routes are preferred over the IX routes.

```
[edit]
lab@R1# run show route 172.17.0.0/24

inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.0.0/24      *[BGP/170] 00:04:25, localpref 210, from 172.30.5.41
                       AS path: 110047427 I
                     > to 172.30.0.10 via ge-0/0/4.118
                     [BGP/170] 01:57:55, localpref 100
                       AS path: 1620 110047427 I
                     > to 192.168.1.3 via ge-0/0/5.300
                     [BGP/170] 01:57:51, localpref 100
                       AS path: 1620 110047427 I
                     > to 192.168.1.4 via ge-0/0/5.300
```

Verify if the P1 routes are advertised to IX peers.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.1.3 aspath-regex
"110047427 .*"

inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)

[edit]
lab@R1# run show route advertising-protocol bgp 192.168.1.4 aspath-regex
"110047427 .*"

inet.0: 834 destinations, 1417 routes (672 active, 0 holddown, 324 hidden)
```

No routes coming from the P1 neighbors are advertised to the IX peers.

© 2020 Juniper Networks, Inc. All rights reserved.

b) R2

Verify the BGP session status with the RR.

```
[edit]
lab@R2# run show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table     Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
               1728        600          0          0          0          0
Peer                AS   InPkt   OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41       54591   1019      415       0       1       13:00 217/600/600/0
0/0/0/0
192.168.1.3        1620    603      529       0       0     1:02:13 383/564/402/0
0/0/0/0
192.168.1.4        1620    602      527       0       0     1:02:09 0/564/402/0
0/0/0/0
```

Verify if unresolved routes are present.

```
[edit]
lab@R2# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify the next-hop for the routes advertised to RR.

```
[edit]
lab@R2# run show route advertising-protocol bgp 172.30.5.41

inet.0: 834 destinations, 1800 routes (672 active, 0 holddown, 324 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 1.64.0.0/10       192.168.1.3                    100        1620 61671 I
* 1.84.160.0/20     192.168.1.3                    100        1620 33112 I
* 1.96.0.0/11       192.168.1.3                    100        1620 33112 63164 40776
51777 I
* 1.161.192.0/21    192.168.1.3                    100        1620 33112 30404 32138
45045 I
* 1.166.89.0/24     192.168.1.3                    100        1620 110047427 56422
47123 34908 57085 3090 19128 21300 I
...
```

As with R1, R2 is not using next-hop self-policy.

© 2020 Juniper Networks, Inc. All rights reserved.

Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@R2# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32    B 170    200          >172.30.0.14       2831679853 I
* 172.17.47.3/32    B 170    200          >172.30.0.14       2831679853 I
```

Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@R2# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and send to RR.

```
[edit]
lab@R2# run show route advertising-protocol bgp 172.30.5.41 aspath-regex "1620
.*" community-name IX

inet.0: 834 destinations, 1800 routes (672 active, 0 holddown, 324 hidden)
  Prefix             Nexthop         MED     Lclpref    AS path
* 1.64.0.0/10        192.168.1.3             100        1620 61671 I
* 1.84.160.0/20      192.168.1.3             100        1620 33112 I
* 1.96.0.0/11        192.168.1.3             100        1620 33112 63164 40776 51777 I
```

```
* 1.161.192.0/21     192.168.1.3             100        1620 33112 30404 32138 45045 I
* 1.166.89.0/24      192.168.1.3             100        1620 110047427 56422 47123
34908 57085 3090 19128 21300 I
* 1.176.0.0/12       192.168.1.3             100        1620 33112 49129 16320 52954 I
...
```

Verify if the customer routes are advertised to the IX peers.

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.1.3 aspath-regex ".*
[64512-64514]"

inet.0: 834 destinations, 1800 routes (672 active, 0 holddown, 324 hidden)
  Prefix             Nexthop      MED     Lclpref    AS path
* 172.31.0.0/24      Self                            54591 54591 54591 [54591] 64512 I
* 172.31.1.0/24      Self                            54591 54591 54591 [54591] 64512 I
* 172.31.2.0/24      Self                            54591 54591 54591 [54591] 64512 I
* 172.31.3.0/24      Self                            54591 54591 54591 [54591] 64512 I
* 172.31.4.0/24      Self                            54591 54591 54591 [54591] 64512 I
* 172.31.5.0/24      Self                            54591 54591 54591 [54591] 64512 I
* 172.31.6.0/24      Self                            54591 54591 54591 [54591] 64512 I
* 172.31.7.0/24      Self                            54591 54591 54591 [54591] 64512 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
...
```

Here you can see that the AS-path prepend makes the routes advertised by R1 more preferred then received from the IX peers.

Note: Fourth prepend of the own AS 54591 in the brackets is the automatic prepend added always between EBGP sessions.

Verify if the local summary route is advertised to IX peers.

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.1.3 172.30/16

inet.0: 834 destinations, 1800 routes (672 active, 0 holddown, 324 hidden)
  Prefix          Nexthop      MED     Lclpref    AS path
* 172.30.0.0/16      Self                        54591 54591 54591 [54591] I
```

Verify the local preference of the customer routes received by R2.

```
[edit]
lab@R2# run show route protocol bgp aspath-regex "64512|64513|64514"

inet.0: 834 destinations, 1800 routes (672 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24    *[BGP/170] 00:01:35, localpref 90, from 172.30.5.41
                     AS path: 64512 I
                   > to 172.30.0.18 via ge-0/0/4.127
172.31.1.0/24    *[BGP/170] 00:01:34, localpref 400, from 172.30.5.41
                     AS path: 64512 I
                     Discard
172.31.2.0/24    *[BGP/170] 00:01:34, localpref 400, from 172.30.5.41
                     AS path: 64512 I
                   > to 172.30.0.14 via ge-0/0/4.123
                     to 172.30.0.18 via ge-0/0/4.127
172.31.3.0/24    *[BGP/170] 00:01:34, localpref 400, from 172.30.5.41
                     AS path: 64512 I
                   > to 172.30.0.14 via ge-0/0/4.123
                     to 172.30.0.18 via ge-0/0/4.127
...
```

Verify the remote triggered black hole routes.

```
[edit]
lab@R2# run show route protocol bgp community-name rtbh

inet.0: 834 destinations, 1800 routes (672 active, 0 holddown, 324 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

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

172.31.1.0/24      *[BGP/170] 00:02:24, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                       Discard
```

Verify if the P1, P2 and P3 routes are preferred to IX routes.

```
[edit]
lab@R2# run show route 172.17.0.0/24

inet.0: 834 destinations, 1800 routes (672 active, 0 holddown, 324 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.0.0/24      *[BGP/170] 00:04:33, localpref 210, from 172.30.5.41
                       AS path: 110047427 I
                     > to 172.30.0.1 via ae0.0
                     [BGP/170] 01:54:31, localpref 100
                       AS path: 1620 110047427 I
                     > to 192.168.1.3 via ge-0/0/5.300
                     [BGP/170] 01:54:27, localpref 100
                       AS path: 1620 110047427 I
                     > to 192.168.1.4 via ge-0/0/5.300
```

Verify if P1 routes are advertised to IX peers.

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.1.3 aspath-regex
"110047427 .*"

inet.0: 684 destinations, 1650 routes (522 active, 0 holddown, 324 hidden)

[edit]
lab@R2# run show route advertising-protocol bgp 192.168.1.4 aspath-regex
"110047427 .*"

inet.0: 684 destinations, 1650 routes (522 active, 0 holddown, 324 hidden)
```

c) R3

Verify the BGP session status with the RR.

```
[edit]
lab@R3# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table  Tot Paths  Act Paths Suppressed  History Damp State  Pending
inet.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            660        600          0          0          0          0
inet6.0
             16         16          0          0          0          0
Peer                 AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41      54591        779        148       0       1       12:57
447/447/447/0        0/0/0/0
192.168.0.2   3521382357        212       1301       0       0     1:01:47
69/97/69/0           0/0/0/0
192.168.0.6   2831679853        220       1289       0       0     1:01:43
84/116/84/0          0/0/0/0
fe80::223:9c01:2d8b:6c81  3521382357        134        139       0       0
1:01:39 Establ
  inet6.0: 16/16/16/0
```

Verify if unresolved routes are present.

```
[edit]
lab@R3# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify the next-hop for the routes advertised to RR.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.41

inet.0: 734 destinations, 736 routes (676 active, 0 holddown, 60 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 5.127.0.0/17      Self                           200        2831679853 9726 36659
30705 25538 37414 49276 ?
* 10.128.0.0/11     Self                           200        2831679853 26697 4341
43012 28104 39181 51157 ?
* 12.173.96.0/19    Self                           200        3521382357 4637 7478
16572 33786 ?
* 15.36.0.0/15      Self                           200        3521382357 4637 57359
16881 43174 64323 I
* 17.96.0.0/13      Self                           200        3521382357 4637 22896
I
* 17.205.88.0/21    Self                           200        2831679853 9726 19412
16490 6122 I
```

Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32    B 170    200          >192.168.0.6       2831679853 I
* 172.17.47.3/32    B 170    200          >192.168.0.6       2831679853 I
```

Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@R3# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and send to RR.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.41 community-name P2

inet.0: 734 destinations, 736 routes (676 active, 0 holddown, 60 hidden)
```

Verify if customer routes are advertised to the IX peers.

```
[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.2 aspath-regex
"64512|64513|64514"

inet.0: 734 destinations, 736 routes (676 active, 0 holddown, 60 hidden)
  Prefix          Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24   Self                                    64512 I
* 172.31.1.0/24   Self                                    64512 I
* 172.31.2.0/24   Self                                    64512 I
* 172.31.3.0/24   Self                                    64512 I
* 172.31.4.0/24   Self                                    64512 I
* 172.31.5.0/24   Self                                    64512 I
* 172.31.6.0/24   Self                                    64512 I
* 172.31.7.0/24   Self                                    64512 I
* 172.31.8.0/24   Self                                    64512 I
* 172.31.9.0/24   Self                                    64512 I
...

[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.6 aspath-regex
"64512|64513|64514"

inet.0: 734 destinations, 736 routes (676 active, 0 holddown, 60 hidden)
  Prefix          Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24   Self                                    64512 I
* 172.31.1.0/24   Self                                    64512 I
* 172.31.2.0/24   Self                                    64512 I
* 172.31.3.0/24   Self                                    64512 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.4.0/24    Self                                    64512 I
* 172.31.5.0/24    Self                                    64512 I
* 172.31.6.0/24    Self                                    64512 I
* 172.31.7.0/24    Self                                    64512 I
...
```

Verify if the local summary route is advertised to P peers.

```
[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.6 172.30/16

inet.0: 734 destinations, 736 routes (676 active, 0 holddown, 60 hidden)
  Prefix          Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16   Self                                    I
[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.2 172.30/16

inet.0: 734 destinations, 736 routes (676 active, 0 holddown, 60 hidden)
  Prefix          Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16   Self                                    I
```

Verify the local preference of the customer routes received by R3.

```
[edit]
lab@R3# run show route protocol bgp aspath-regex "64512|64513|64514"

inet.0: 734 destinations, 736 routes (676 active, 0 holddown, 60 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[BGP/170] 00:01:38, localpref 90, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.13 via ge-0/0/4.123
                       to 172.30.0.26 via ge-0/0/4.136
172.31.1.0/24      *[BGP/170] 00:01:38, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                       Discard
172.31.2.0/24      *[BGP/170] 00:01:38, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.26 via ge-0/0/4.136
172.31.3.0/24      *[BGP/170] 00:01:38, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.26 via ge-0/0/4.136
...
```

Verify the remote triggered black hole routes.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show route protocol bgp community-name rtbh

inet.0: 734 destinations, 736 routes (676 active, 0 holddown, 60 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 00:02:27, localpref 400, from 172.30.5.41
                        AS path: 64512 I
                      Discard
```

Verify if the P1, P2 and P3 routes are preferred to IX routes.

```
[edit]
lab@R3# run show route 172.17.0.0/24

inet.0: 734 destinations, 736 routes (676 active, 0 holddown, 60 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.0.0/24      *[BGP/170] 00:04:37, localpref 210, from 172.30.5.41
                        AS path: 110047427 I
                        to 172.30.0.13 via ge-0/0/4.123
                      > to 172.30.0.26 via ge-0/0/4.136
```

On R3 only the best route is advertised by the RR.

d) R4.

Verify the BGP session status with the RR.

```
[edit]
lab@R4# run show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     600        600          0          0          0          0
Peer                 AS    InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41    54591     987         26       0       0       10:25 600/600/600/0
0/0/0/0
```

R4 has just one connection with the RR.

Verify if unresolved routes are present.

```
[edit]
lab@R4# run show route resolution unresolved
Tree Index 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@R4# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32    B 170    200          >172.30.0.21       2831679853 I
* 172.17.47.3/32    B 170    200          >172.30.0.21       2831679853 I
```

Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@R4# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify the local preference of the customer routes received by R4.

```
[edit]
lab@R4# run show route protocol bgp aspath-regex "64512|64513|64514"

inet.0: 676 destinations, 676 routes (676 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[BGP/170] 00:01:39, localpref 90, from 172.30.5.41
                        AS path: 64512 I
                      > to 172.30.0.5 via ge-0/0/4.114
                        to 172.30.0.30 via ge-0/0/4.145
172.31.1.0/24      *[BGP/170] 00:01:38, localpref 400, from 172.30.5.41
                        AS path: 64512 I
                        Discard
172.31.2.0/24      *[BGP/170] 00:01:38, localpref 400, from 172.30.5.41
                        AS path: 64512 I
                      > to 172.30.0.30 via ge-0/0/4.145
172.31.3.0/24      *[BGP/170] 00:01:38, localpref 400, from 172.30.5.41
                        AS path: 64512 I
                      > to 172.30.0.30 via ge-0/0/4.145
...
```

Verify the remote triggered black hole routes.

```
[edit]
lab@R4# run show route protocol bgp community-name rtbh

inet.0: 676 destinations, 676 routes (676 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.31.1.0/24      *[BGP/170] 00:02:27, localpref 400, from 172.30.5.41
                      AS path: 64512 I
                    Discard
```

e) R5

Verify the BGP session status with the RR.

```
[edit]
lab@R5# run show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     616        616          0          0          0          0
inet6.0
                      32         32          0          0          0          0
Peer             AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41    54591        987         32       0       1       12:52
584/584/584/0        0/0/0/0
192.168.0.10   64514        134       1593       0       0     1:00:21 Establ
  inet.0: 16/16/16/0
  inet6.0: 16/16/16/0
192.168.0.14   64514        133       1592       0       0     1:00:17 Establ
  inet.0: 16/16/16/0
  inet6.0: 16/16/16/0
```

Verify if unresolved routes are present.

```
[edit]
lab@R5# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify the next-hop for the routes advertised to RR.

```
[edit]
lab@R5# run show route advertising-protocol bgp 172.30.5.41

inet.0: 678 destinations, 694 routes (678 active, 0 holddown, 0 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 172.31.32.0/24     Self                         300        64514 I
* 172.31.33.0/24     Self                         300        64514 I
* 172.31.34.0/24     Self                         300        64514 I
* 172.31.35.0/24     Self                         300        64514 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.36.0/24  Self                          300        64514 I
* 172.31.37.0/24  Self                          300        64514 I
* 172.31.38.0/24  Self                          300        64514 I
* 172.31.39.0/24  Self                          300        64514 I
...
```

Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@R5# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32   B 170    200       >172.30.0.34      2831679853 I
* 172.17.47.3/32   B 170    200       >172.30.0.34      2831679853 I
```

Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@R5# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and send to RR.

```
[edit]
lab@R5# run show route advertising-protocol bgp 172.30.5.41 community-name C3

inet.0: 678 destinations, 694 routes (678 active, 0 holddown, 0 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 172.31.32.0/24     Self                         300        64514 I
* 172.31.33.0/24     Self                         300        64514 I
* 172.31.34.0/24     Self                         300        64514 I
* 172.31.35.0/24     Self                         300        64514 I
* 172.31.36.0/24     Self                         300        64514 I
* 172.31.37.0/24     Self                         300        64514 I
* 172.31.38.0/24     Self                         300        64514 I
* 172.31.39.0/24     Self                         300        64514 I
* 172.31.40.0/24     Self                         300        64514 I
* 172.31.41.0/24     Self                         300        64514 I
...
```

Verify if customer routes are advertised to the customer AS.

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.10 aspath-regex
"64512|64513|64514"

inet.0: 678 destinations, 694 routes (678 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24           Self                                    64512 I
* 172.31.1.0/24           Self                                    64512 I
* 172.31.2.0/24           Self                                    64512 I
* 172.31.3.0/24           Self                                    64512 I
* 172.31.4.0/24           Self                                    64512 I
* 172.31.5.0/24           Self                                    64512 I
* 172.31.6.0/24           Self                                    64512 I
* 172.31.7.0/24           Self                                    64512 I
* 172.31.8.0/24           Self                                    64512 I
...
```

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.14 aspath-regex
"64512|64513|64514"

inet.0: 678 destinations, 694 routes (678 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24           Self                                    64512 I
* 172.31.1.0/24           Self                                    64512 I
* 172.31.2.0/24           Self                                    64512 I
* 172.31.3.0/24           Self                                    64512 I
* 172.31.4.0/24           Self                                    64512 I
* 172.31.5.0/24           Self                                    64512 I
* 172.31.6.0/24           Self                                    64512 I
* 172.31.7.0/24           Self                                    64512 I
* 172.31.8.0/24           Self                                    64512 I
...
```

Verify if the local summary route is advertised to Customer AS.

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.10 172.30/16

inet.0: 678 destinations, 694 routes (678 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                                    I

[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.14 172.30/16

inet.0: 128 destinations, 144 routes (128 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                                    I
```

Verify the local preference of the customer routes received by R5.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# run show route protocol bgp aspath-regex "64512|64513|64514"

inet.0: 678 destinations, 694 routes (678 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[BGP/170] 00:01:40, localpref 90, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.34 via ae0.0
172.31.1.0/24      *[BGP/170] 00:01:39, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                       Discard
172.31.2.0/24      *[BGP/170] 00:01:39, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.34 via ae0.0
172.31.3.0/24      *[BGP/170] 00:01:39, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.34 via ae0.0
```

Verify the remote triggered black hole routes.

```
[edit]
lab@R5# run show route protocol bgp community-name rtbh

inet.0: 678 destinations, 694 routes (678 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 00:02:29, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                       Discard
```

Verify if the P1, P2 and P3 routes are preferred to IX routes.

```
[edit]
lab@R5# run show route 172.17.0.0/24

inet.0: 678 destinations, 694 routes (678 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.0.0/24      *[BGP/170] 00:04:39, localpref 210, from 172.30.5.41
                       AS path: 110047427 I
                     > to 172.30.0.38 via ge-0/0/4.158
```

Verify the load balancing because of the multipath setting.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# run show route aspath-regex 64514

inet.0: 678 destinations, 694 routes (678 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.32.0/24     *[BGP/170] 01:55:47, localpref 300
                        AS path: 64514 I
                      > to 192.168.0.10 via ge-0/0/5.303
                        to 192.168.0.14 via ge-0/0/5.304
                      [BGP/170] 01:55:43, localpref 300
                        AS path: 64514 I
                      > to 192.168.0.14 via ge-0/0/5.304
172.31.33.0/24     *[BGP/170] 01:55:47, localpref 300, from 192.168.0.10
                        AS path: 64514 I
                        to 192.168.0.10 via ge-0/0/5.303
                      > to 192.168.0.14 via ge-0/0/5.304
                      [BGP/170] 01:55:43, localpref 300
                        AS path: 64514 I
                      > to 192.168.0.14 via ge-0/0/5.304
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

f) R6

Verify the BGP session status with the RR.

```
[edit]
lab@R6# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     601        600          0          0          0          0
Peer                 AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41      54591       1130         48       0       0       18:34
568/569/569/0         0/0/0/0
172.31.31.1      64513        130        202       0       0       59:26
16/16/16/0            0/0/0/0
192.168.0.18     64512        130       1209       0       0       59:22
16/16/16/0            0/0/0/0
```

Verify if unresolved routes are present.

```
[edit]
lab@R6# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify the next-hop for the routes advertised to RR.

```
[edit]
lab@R6# run show route advertising-protocol bgp 172.30.5.41

inet.0: 679 destinations, 680 routes (679 active, 0 holddown, 0 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24      Self                         90         64512 I
* 172.31.1.0/24      Self                         400        64512 I
* 172.31.2.0/24      Self                         400        64512 I
* 172.31.3.0/24      Self                         400        64512 I
* 172.31.4.0/24      Self                         400        64512 I
* 172.31.5.0/24      Self                         400        64512 I
* 172.31.6.0/24      Self                         400        64512 I
* 172.31.7.0/24      Self                         400        64512 I
* 172.31.8.0/24      Self                         400        64512 I
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@R6# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32    B 170    200          >172.30.0.25       2831679853 I
* 172.17.47.3/32    B 170    200          >172.30.0.25       2831679853 I
```

Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@R6# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and send to RR.

```
[edit]
lab@R6# run show route advertising-protocol bgp 172.30.5.41 community-name C1

inet.0: 679 destinations, 680 routes (679 active, 0 holddown, 0 hidden)
  Prefix              Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24       Self                         90         64512 I
* 172.31.1.0/24       Self                         400        64512 I
* 172.31.2.0/24       Self                         400        64512 I
* 172.31.3.0/24       Self                         400        64512 I
* 172.31.4.0/24       Self                         400        64512 I
* 172.31.5.0/24       Self                         400        64512 I
* 172.31.6.0/24       Self                         400        64512 I
...
```

```
[edit]
lab@R6# run show route advertising-protocol bgp 172.30.5.41 community-name C2

inet.0: 679 destinations, 680 routes (679 active, 0 holddown, 0 hidden)
  Prefix              Nexthop              MED     Lclpref    AS path
* 172.31.16.0/24      Self                         300        64513 I
* 172.31.17.0/24      Self                         300        64513 I
* 172.31.18.0/24      Self                         300        64513 I
* 172.31.19.0/24      Self                         300        64513 I
* 172.31.20.0/24      Self                         300        64513 I
* 172.31.21.0/24      Self                         300        64513 I
* 172.31.22.0/24      Self                         300        64513 I
...
```

Verify if customer routes are advertised to the EBGP peers.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# run show route advertising-protocol bgp 172.31.31.1

inet.0: 679 destinations, 680 routes (679 active, 0 holddown, 0 hidden)
  Prefix              Nexthop      MED     Lclpref    AS path
* 0.0.0.0/0           Self         10                 {101 235 339 783 784 1044 1158
… }
```

Note: To C2-1 only the default is advertised as requested.

```
[edit]
lab@R6# run show route advertising-protocol bgp 192.168.0.18 aspath-regex
"64512|64513|64514"

inet.0: 679 destinations, 680 routes (679 active, 0 holddown, 0 hidden)
  Prefix              Nexthop              MED     Lclpref    AS path
* 172.31.16.0/24      Self                 10                 64513 I
* 172.31.17.0/24      Self                 10                 64513 I
* 172.31.18.0/24      Self                 10                 64513 I
* 172.31.19.0/24      Self                 10                 64513 I
* 172.31.20.0/24      Self                 10                 64513 I
* 172.31.21.0/24      Self                 10                 64513 I
* 172.31.22.0/24      Self                 10                 64513 I
* 172.31.23.0/24      Self                 10                 64513 I
...
```

Verify if the local summary route is advertised to C1-1 peer.

```
[edit]
lab@R6# run show route advertising-protocol bgp 192.168.0.18 172.30/16

inet.0: 130 destinations, 145 routes (130 active, 0 holddown, 0 hidden)
  Prefix              Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16       Self                 10                 I
```

Note: Note that the summary route is advertised with a MED value of 10. This will make R6 more preferred for inbound traffic.

Verify the local preference of the customer routes received by R6.

```
[edit]
lab@R6# run show route protocol bgp aspath-regex "64512|64513|64514"

inet.0: 679 destinations, 680 routes (679 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24          *[BGP/170] 01:48:43, localpref 90
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                      AS path: 64512 I
                    > to 192.168.0.18 via ge-0/0/5.305
                    [BGP/170] 00:01:41, localpref 90, from 172.30.5.41
                      AS path: 64512 I
                    > to 172.30.0.42 via ge-0/0/4.167
172.31.1.0/24       *[BGP/170] 01:48:43, localpref 400
                      AS path: 64512 I
                    > to 192.168.0.18 via ge-0/0/5.305
172.31.2.0/24       *[BGP/170] 01:48:43, localpref 400
                      AS path: 64512 I
                    > to 192.168.0.18 via ge-0/0/5.305
...
```

Verify the remote triggered black hole routes.

```
[edit]
lab@R6# run show route protocol bgp community-name rtbh

inet.0: 679 destinations, 680 routes (679 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24       *[BGP/170] 01:50:55, localpref 400
                       AS path: 64512 I
                       > to 192.168.0.18 via ge-0/0/5.305
```

R6 receives the route with community `rtbh` directly from the EBGP session. Because the `rtbh` policy is applied as import on the IBGP session, on R6 the route is active and has a valid next-hop.

Verify the multihop load balancing towards C2-1 peer.

```
[edit]
lab@R6# run show route aspath-regex 64513

inet.0: 679 destinations, 680 routes (679 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.16.0/24    *[BGP/170] 01:55:08, localpref 300, from 172.31.31.1
                       AS path: 64513 I
                     > to 192.168.0.22 via ge-0/0/5.306
                       to 192.168.0.26 via ge-0/0/5.307
172.31.17.0/24    *[BGP/170] 01:55:08, localpref 300, from 172.31.31.1
                       AS path: 64513 I
                     > to 192.168.0.22 via ge-0/0/5.306
                       to 192.168.0.26 via ge-0/0/5.307
172.31.18.0/24    *[BGP/170] 01:55:08, localpref 300, from 172.31.31.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    AS path: 64513 I
                  > to 192.168.0.22 via ge-0/0/5.306
                    to 192.168.0.26 via ge-0/0/5.307
...
```

Verify if the R6 is preferred for outbound traffic to C1.

```
[edit]
lab@R6# run show route 172.31.2.0/24

inet.0: 679 destinations, 680 routes (679 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.2.0/24      *[BGP/170] 02:20:28, localpref 400
                      AS path: 64512 I
                    > to 192.168.0.18 via ge-0/0/5.305
```

g) R7

Verify the BGP session status with the RR.

```
[edit]
lab@R7# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table      Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
               1179        600          0          0          0          0
inet6.0
                 16         16          0          0          0          0
Peer                 AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41      54591       1136         59       0       0       18:35
599/599/599/0        0/0/0/0
192.168.0.30 110047427        615        513       0       0       58:55
0/564/16/0           0/0/0/0
192.168.0.34     64512        130       1211       0       0       58:59
1/16/16/0            0/0/0/0
fc09:c0:ffee::2  110047427        128        132       0       0       58:51
Establ
  inet6.0: 16/16/16/0
```

Verify if unresolved routes are present.

```
[edit]
lab@R7# run show route resolution unresolved
Tree Index 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify the next-hop for the routes advertised to RR.

```
[edit]
lab@R7# run show route advertising-protocol bgp 172.30.5.41

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24           Self                         90         64512 I
```

R7 is advertising only one route from C1-1. The routes received from P1-2 are not advertised, because the same routes are advertised by R8 to the RR with better local preference.

Verify for routes with mask shorter than /8 and longer than /24.

```
lab@R7# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32    B 170    200          >172.30.0.17       2831679853 I
* 172.17.47.3/32    B 170    200          >172.30.0.17       2831679853 I
```

Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@R7# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and send to RR.

```
[edit]
lab@R7# run show route advertising-protocol bgp 172.30.5.41 community-name C1

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24           Self                         90         64512 I

[edit]
lab@R7# run show route advertising-protocol bgp 172.30.5.41 community-name P1

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
```

P1 routes are not advertised because the same routes are received from R8 with better local-preference.

© 2020 Juniper Networks, Inc. All rights reserved.

Verify if customer routes are advertised to the EBGP peers.

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.30 aspath-regex
"64512|64513|64514"

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
  Prefix                Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24         Self                                    64512 I
* 172.31.1.0/24         Self                                    64512 I
* 172.31.2.0/24         Self                                    64512 I
* 172.31.3.0/24         Self                                    64512 I
* 172.31.4.0/24         Self                                    64512 I
* 172.31.5.0/24         Self                                    64512 I
* 172.31.6.0/24         Self                                    64512 I
* 172.31.7.0/24         Self                                    64512 I
* 172.31.8.0/24         Self                                    64512 I
...
```

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.34 aspath-regex
"64512|64513|64514"

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
  Prefix                Nexthop              MED     Lclpref    AS path
* 172.31.16.0/24        Self                 20                 64513 I
* 172.31.17.0/24        Self                 20                 64513 I
* 172.31.18.0/24        Self                 20                 64513 I
* 172.31.19.0/24        Self                 20                 64513 I
* 172.31.20.0/24        Self                 20                 64513 I
* 172.31.21.0/24        Self                 20                 64513 I
* 172.31.22.0/24        Self                 20                 64513 I
* 172.31.23.0/24        Self                 20                 64513 I
...
```

Verify if the local summary route is advertised to EBGP peers.

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.30 172.30/16

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
  Prefix                Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16         Self                                    I
* 172.30.128.0/17       Self                                    I
```


© 2020 Juniper Networks, Inc. All rights reserved.

R7 is advertising part of the internal routes to P1-2 together with the summary route of all internal routes.

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.34 172.30/16

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16           Self                 20                 I
```

R7 is advertising a summary route to C1-1 with a MED value of 20 making R6 preferred for inbound traffic for C1-1.

Verify the local preference of the customer routes received by R7.

```
[edit]
lab@R7# run show route protocol bgp aspath-regex "64512|64513|64514"

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[BGP/170] 01:48:19, localpref 90
                        AS path: 64512 I
                      > to 192.168.0.34 via ge-0/0/5.309
172.31.1.0/24      *[BGP/170] 00:01:42, localpref 400, from 172.30.5.41
                        AS path: 64512 I
                        Discard
                      [BGP/170] 01:48:19, localpref 300
                        AS path: 64512 I
                      > to 192.168.0.34 via ge-0/0/5.309
172.31.2.0/24      *[BGP/170] 00:01:42, localpref 400, from 172.30.5.41
                        AS path: 64512 I
                      > to 172.30.0.41 via ge-0/0/4.167
                      [BGP/170] 01:48:19, localpref 300
                        AS path: 64512 I
                      > to 192.168.0.34 via ge-0/0/5.309
...
```

Verify the remote triggered black hole routes.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show route protocol bgp community-name rtbh

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 00:03:55, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                       Discard
                    [BGP/170] 01:50:32, localpref 300
                       AS path: 64512 I
                    > to 192.168.0.34 via ge-0/0/5.309
```

Verify if R6 is preferred for outbound traffic to C1.

```
[edit]
lab@R7# run show route 172.31.2.0/24

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.2.0/24      *[BGP/170] 00:14:46, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                    > to 172.30.0.41 via ge-0/0/4.167
                    [BGP/170] 02:01:23, localpref 300
                       AS path: 64512 I
                    > to 192.168.0.34 via ge-0/0/5.309
```

Verify that routes not native to P1 AS are not accepted.

```
[edit]
lab@R7# run show route receive-protocol bgp 192.168.0.30 aspath-regex "110047427
.+"

inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)

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

inet6.0: 50 destinations, 52 routes (50 active, 0 holddown, 0 hidden)
```

Verify if locally advertised routes are set with the community no-export.

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.30 172.30/16 extensive
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 837 destinations, 1254 routes (675 active, 0 holddown, 548 hidden)
* 172.30.0.0/16 (1 entry, 1 announced)
 BGP group P1-2 type External
     Nexthop: Self
     AS path: [54591] I (LocalAgg)

* 172.30.128.0/17 (1 entry, 1 announced)
 BGP group P1-2 type External
     Nexthop: Self
     AS path: [54591] I (LocalAgg)
     Communities: no-export
```

h) R8

Verify the BGP session status with the RR.

```
[edit]
lab@R8# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table     Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
               1148        600          0          0          0          0
inet6.0
                 16         16          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41          54591       1139         44       0       0       18:45
584/584/584/0         0/0/0/0
192.168.0.38    110047427        614        515       0       0       58:16
16/564/16/0           0/0/0/0
fc09:c0:ffee::6 110047427        127        130       0       0       58:12
Establ
  inet6.0: 16/16/16/0
```

Verify if unresolved routes are present.

```
[edit]
lab@R8# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

Verify the next-hop for the routes advertised to RR.

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.30.5.41
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 835 destinations, 1221 routes (673 active, 0 holddown, 548 hidden)
  Prefix             Nexthop                 MED     Lclpref    AS path
* 172.17.0.0/24      Self                            210        110047427 I
* 172.17.1.0/24      Self                            210        110047427 I
* 172.17.2.0/24      Self                            210        110047427 I
* 172.17.3.0/24      Self                            210        110047427 I
* 172.17.4.0/24      Self                            210        110047427 I
* 172.17.5.0/24      Self                            210        110047427 I
* 172.17.6.0/24      Self                            210        110047427 I
...
```

Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@R8# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32  B 170    200          >172.30.0.9       2831679853 I
* 172.17.47.3/32  B 170    200          >172.30.0.9       2831679853 I
```

Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@R8# run show route protocol bgp terse | match " 0.0.0.0"
```

Verify if the correct community is set and send to RR.

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.30.5.41 community-name P1

inet.0: 835 destinations, 1221 routes (673 active, 0 holddown, 548 hidden)
  Prefix             Nexthop                 MED     Lclpref    AS path
* 172.17.0.0/24      Self                            210        110047427 I
* 172.17.1.0/24      Self                            210        110047427 I
* 172.17.2.0/24      Self                            210        110047427 I
* 172.17.3.0/24      Self                            210        110047427 I
* 172.17.4.0/24      Self                            210        110047427 I
* 172.17.5.0/24      Self                            210        110047427 I
* 172.17.6.0/24      Self                            210        110047427 I
...
```

Verify if customer routes are advertised to the P1-1 peer.

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.38 aspath-regex
"64512|64513|64514"
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 835 destinations, 1221 routes (673 active, 0 holddown, 548 hidden)
  Prefix              Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24       Self                                    64512 I
* 172.31.1.0/24       Self                                    64512 I
* 172.31.2.0/24       Self                                    64512 I
* 172.31.3.0/24       Self                                    64512 I
* 172.31.4.0/24       Self                                    64512 I
* 172.31.5.0/24       Self                                    64512 I
* 172.31.6.0/24       Self                                    64512 I
...
```

Verify if the local summary route is advertised to P1-1 peer.

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.38 172.30/16

inet.0: 835 destinations, 1221 routes (673 active, 0 holddown, 548 hidden)
  Prefix                Nexthop              MED     Lclpref    AS path
* 172.30.0.0/16         Self                                    I
* 172.30.0.0/17         Self                                    I
```

First part of internal addresses are advertised together with the summary route of all internal routes.

Verify the local preference of the customer routes received by R8.

```
[edit]
lab@R8# run show route protocol bgp aspath-regex "64512|64513|64514"

inet.0: 835 destinations, 1221 routes (673 active, 0 holddown, 548 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[BGP/170] 00:02:03, localpref 90, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.45 via ge-0/0/4.178
172.31.1.0/24      *[BGP/170] 00:02:03, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                       Discard
172.31.2.0/24      *[BGP/170] 00:02:03, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.37 via ge-0/0/4.158
172.31.3.0/24      *[BGP/170] 00:02:03, localpref 400, from 172.30.5.41
                       AS path: 64512 I
                     > to 172.30.0.37 via ge-0/0/4.158
...
```

Verify the remote triggered black hole routes.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# run show route protocol bgp community-name rtbh

inet.0: 835 destinations, 1221 routes (673 active, 0 holddown, 548 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 00:03:57, localpref 400, from 172.30.5.41
                        AS path: 64512 I
                        Discard
```

Verify that routes not originated in AS P1 are not accepted.

```
[edit]
lab@R8# run show route protocol bgp 192.168.0.38 aspath-regex "110047427 .+"

inet.0: 835 destinations, 1221 routes (673 active, 0 holddown, 548 hidden)

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

inet6.0: 50 destinations, 52 routes (50 active, 0 holddown, 0 hidden)
```

Verify if locally advertised routes are set with the community no-export.

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.38 172.30/16 extensive

inet.0: 835 destinations, 1221 routes (673 active, 0 holddown, 548 hidden)
* 172.30.0.0/16 (1 entry, 1 announced)
 BGP group P1-1 type External
     Nexthop: Self
     AS path: [54591] I (LocalAgg)

* 172.30.0.0/17 (1 entry, 1 announced)
 BGP group P1-1 type External
     Nexthop: Self
     AS path: [54591] I (LocalAgg)
     Communities: no-export
```


© 2020 Juniper Networks, Inc. All rights reserved.

Appendix - Chapter 4: MPLS Configuration

This chapter contains tasks that require you to configure advanced transport LSPs for the provider network.

MPLS Overview

Multiprotocol Protocol Label Switching (MPLS) uses labels for forwarding traffic through the network. The concatenation of one hop labels define a Label Switched Paths (LSPs) through the network. This creates clean separation between the control and data plain and provides the ability to multiplex different services over the network. The MPLS Shim header, which contains the label, is inserted between the Layer 2 and Layer 3 headers when packets enter the MPLS network. Based on that label, the packet is label switched through the network. At the egress point of the MPLS network the header is removed and the original packet is forwarded to the destination in its native form.

The following terminology is used in MPLS networks:

- Label Switch Router (LSR) – every router in the MPLS network configured for MPLS forwarding and participates in the creation of LSPs.
- Ingress Router – the router where the packet enters the MPLS network. The packets enter an LSP based on the destination. This is the head-end of the LSP. This router performs a label push operation.
- Transit Router – this router forwards the labeled packet to the next hop of the LSP. It performs swap label operation.
- Penultimate Router – second last router. It is upstream to the egress router. By default, Junos pops the upper label in the stack on this router. If the label stack consists of just one label it is popped at this LSR and a pure IP packet is sent to the egress LSR. The egress LSR should perform a pure IP lookup to forward the packet to the next hop.
- Egress Router – here the packets exit the LSP and MPLS network. This is the tail-end of the LSP. This LSR performs a MPLS pop operation and forwards the packet based on an IP forwarding lookup.

MPLS shim header is 32 bits fixed size. It consists of the following four fields:

- ***Label*** -20 bits containing the label of the packet. The field specifies to which LSP the packet is assigned. It changes from hop to hop while the packet is transported through the network.
- ***Traffic Class Bits (TC)*** – 3 bits field used for quality of service. Previously these bits where called EXP (experimental).

© 2020 Juniper Networks, Inc. All rights reserved.

- ***Bottom of the stack*** – 1 bit field known as 'S' bit is used to indicated if more than one label is used. MPLS supports the use of more than one label. If this bit is set to 1 it means that this is the last label in the stack and an unlabeled packet is following.
- ***Time to live*** – 8 bits TTL field has the same function as with IP packets. It is used to drop the packet when the TTL value reaches 0.

The label field is swapped hop by hop, each time it passes through MPLS enabled router. LSP path can be manually build or with the help of a signaling protocol. The JNCIE exam will tests the candidates' skills on their knowledge of several label signaling protocols.

Label values from 0 to 15 are reserved. Below you will see the most commonly used reserved label values:

- ***Label 0*** – this label is called *IPv4 Explicit null*. It instructs the penultimate router to forward the packet with an MPLS header.
- ***Label 1*** – *Router Alert Label* is something similar like the router alert option in IP packets. It informs the router that the content must be processed by the local router.
- ***Label 2*** – *IPv6 Explicit null* label is used for the same purpose as Label 0 but with IPv6 traffic.
- ***Label 3*** – *Implicit Null* label is by default signaled from the tail-end router to the penultimate router to pop the upper label in the stack.

Junos supports two different Label distribution protocols. Label distribution protocol (LDP) and Resource Reservation Protocol (RSVP).

LDP Overview

The Label Distribution Protocol (LDP) associates a group of prefixes with an LSP. This group is called a Forwarding Equivalence Class (FEC). LDP uses the IGP best path through the network to signal the LSP from the egress LSR to the ingress router. LDP does not have any Traffic Engineering (TE) mechanisms. By default, labels are exchanged in downstream unsolicited advertisement mode. This means that labels for all routes/FECs are advertised and received by all routers. Junos limits this to the loopback interface address only. With additional policy you can force LDP to advertise labels for other routes as well. This way a full mesh LDP network is created, every LSR is ingress for all other routers in the network.

LDP uses UDP multicast on port 646 for the neighbor discovery mechanism. The destination address is 224.0.0.2 which limits the discovery messages to be send only to directly connected routers. After the neighbor discovery phase, LDP builds sessions using TCP on port 646. The LDP transport address is used for initiating the TCP session. If a

© 2020 Juniper Networks, Inc. All rights reserved.

Router ID is not configured, the interface address is used as the Transport address. An LDP transport session is maintained with keepalive messages. LDP also supports neighbor authentication based on the following mechanisms: MD5, SHA1 or AES-128-cipher.

Task 1: LDP Configuration

1) Configure `family mpls` and enable the mpls protocol on the interfaces.

You have to configure LDP as shown in Figure 8. Not all interfaces are enabled for LDP. At a later stage both parts of the core with LDP will be connected with RSVP tunnels!

As with IS-IS you have to enable the interface to accept labeled packets.

Repeat this on all eight devices.

```
[edit]
lab@R1# set groups if-families interfaces ge-0/0/4 unit <*> family mpls


[edit]
lab@R1# set groups if-families interfaces <ae*> unit <*> family mpls
```

An apply-group is already present and can be used to enable `family mpls` on all core interfaces.

Next the interfaces should be specified under `protocols mpls` hierarchy. This way MPLS will know which interfaces it can use to forward traffic. You can use the option *all* to add all interfaces with `family mpls`. If the device has a dedicated management interface fxp0 it is good practices to disable it under `protocols mpls`.

```
[edit]
lab@R1# set protocols mpls interface all

[edit]
lab@R1# set protocols mpls interface fxp0.0 disable

[edit]
lab@R1# commit
commit complete
```

Verify if the interfaces are enabled for MPLS.

```
[edit]
lab@R1# run show mpls interface
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Interface          State       Administrative groups (x: extended)
ae0.0              Up          <none>
ge-0/0/4.114       Up          <none>
ge-0/0/4.118       Up          <none>
ge-0/0/5.300       Dn          <none>
ge-0/0/4.206       Up          <none>
```

Four interfaces on R1 are enabled for MPLS.

2) Configure LDP according to the Figure 8.

Next you must enable LDP as label distribution protocol according to Figure 8. Some additional features like authentication should be enabled on the sessions. Configurations are the same on all routers.

a) R1

```
[edit]
lab@R1# set protocols ldp interface ae0.0

[edit]
lab@R1# set protocols ldp interface ge-0/0/4.114
```

Adding both interfaces to LDP protocol.

```
[edit]
lab@R1# set protocols ldp session 172.30.5.2 authentication-key jncie

[edit]
lab@R1# set protocols ldp session 172.30.5.4 authentication-key jncie
```

LDP authentication is configured for the transport session. Because the routers have configured a Router ID, the sessions are using the loopback IP addresses.

```
[edit]
lab@R1# set protocols ldp track-igp-metric
```

By default, the LDP metric is 1, with a preference of 9. Configure this option to instructs LDP to use the IGP metric for the routes.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set protocols ldp explicit-null

[edit]
lab@R1# commit
commit complete
```

The `explicit-null` knob forces LDP to not advertise label 3 for the FECs which this LSR is the egress node for. LDP advertising a label value 0 (IPv4 explicit null) instead, which forces the penultimate-hop router to not pop the upper label.

b) R2

```
[edit]
lab@R2# set protocols ldp interface ae0.0

[edit]
lab@R2# set protocols ldp interface ge-0/0/4.123

[edit]
lab@R2# set protocols ldp session 172.30.5.1 authentication-key jncie

[edit]
lab@R2# set protocols ldp session 172.30.5.3 authentication-key jncie

[edit]
lab@R2# set protocols ldp track-igp-metric

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols ldp interface ge-0/0/4.123

[edit]
lab@R3# set protocols ldp interface ge-0/0/4.134

[edit]
lab@R3# set protocols ldp session 172.30.5.2 authentication-key jncie

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set protocols ldp session 172.30.5.4 authentication-key jncie

[edit]
lab@R3# set protocols ldp track-igp-metric

[edit]
lab@R3# set protocols ldp explicit-null

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols ldp interface ge-0/0/4.114

[edit]
lab@R4# set protocols ldp interface ge-0/0/4.134

[edit]
lab@R4# set protocols ldp session 172.30.5.1 authentication-key jncie

[edit]
lab@R4# set protocols ldp session 172.30.5.3 authentication-key jncie

[edit]
lab@R4# set protocols ldp track-igp-metric

[edit]
lab@R4# set protocols ldp explicit-null

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols ldp interface ae0.0

[edit]
lab@R5# set protocols ldp interface ge-0/0/4.158
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols ldp session 172.30.5.6 authentication-key jncie

[edit]
lab@R5# set protocols ldp session 172.30.5.8 authentication-key jncie

[edit]
lab@R5# set protocols ldp track-igp-metric

[edit]
lab@R5# set protocols ldp explicit-null

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols ldp interface ae0.0

[edit]
lab@R6# set protocols ldp interface ge-0/0/4.167

[edit]
lab@R6# set protocols ldp session 172.30.5.5 authentication-key jncie

[edit]
lab@R6# set protocols ldp session 172.30.5.7 authentication-key jncie

[edit]
lab@R6# set protocols ldp track-igp-metric

[edit]
lab@R6# set protocols ldp explicit-null

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# set protocols ldp interface ge-0/0/4.167

[edit]
lab@R7# set protocols ldp interface ge-0/0/4.178

[edit]
lab@R7# set protocols ldp session 172.30.5.6 authentication-key jncie

[edit]
lab@R7# set protocols ldp session 172.30.5.8 authentication-key jncie

[edit]
lab@R7# set protocols ldp track-igp-metric

[edit]
lab@R7# set protocols ldp explicit-null

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols ldp interface ge-0/0/4.158

[edit]
lab@R8# set protocols ldp interface ge-0/0/4.178

[edit]
lab@R8# set protocols ldp session 172.30.5.5 authentication-key jncie

[edit]
lab@R8# set protocols ldp session 172.30.5.7 authentication-key jncie

[edit]
lab@R8# set protocols ldp track-igp-metric

[edit]
lab@R8# set protocols ldp explicit-null

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

After committing and adding the necessary LDP configuration to all routers you can verify the LDP sessions.

```
[edit]
lab@R1# run show ldp neighbor
Address              Interface          Label space ID         Hold time
172.30.0.2           ae0.0              172.30.5.2:0             13
172.30.0.6           ge-0/0/4.114       172.30.5.4:0             14
```

Execute the above command with the extensive option reveals more detail.

```
[edit]
lab@R1# run show ldp neighbor extensive
Address              Interface          Label space ID         Hold time
172.30.0.2           ae0.0              172.30.5.2:0             11
  Transport address: 172.30.5.2, Configuration sequence: 1
  Up for 00:52:41
  Reference count: 1
  Hold time: 15, Proposed local/peer: 15/15
  Hello flags: none
  Neighbor types: discovered
Address              Interface          Label space ID         Hold time
172.30.0.6           ge-0/0/4.114       172.30.5.4:0             11
  Transport address: 172.30.5.4, Configuration sequence: 1
  Up for 00:35:40
  Reference count: 1
  Hold time: 15, Proposed local/peer: 15/15
  Hello flags: none
  Neighbor types: discovered
```

You can see the Transport address used for each session.

The LDP label database displays the label received and send to the neighbors.

```
[edit]
lab@R1# run show ldp database
Input label database, 172.30.5.1:0--172.30.5.2:0
  Label     Prefix
 299824     172.30.5.1/32
      3     172.30.5.2/32
 299840     172.30.5.3/32
 299856     172.30.5.4/32

Output label database, 172.30.5.1:0--172.30.5.2:0
  Label     Prefix
      0     172.30.5.1/32
 299776     172.30.5.2/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 299792     172.30.5.3/32
 299808     172.30.5.4/32

Input label database, 172.30.5.1:0--172.30.5.4:0
  Label     Prefix
 299792     172.30.5.1/32
 299808     172.30.5.2/32
 299776     172.30.5.3/32
      0     172.30.5.4/32

Output label database, 172.30.5.1:0--172.30.5.4:0
  Label     Prefix
      0     172.30.5.1/32
 299776     172.30.5.2/32
 299792     172.30.5.3/32
 299808     172.30.5.4/32
```

From the above output you can see the assigned label for all loopback IP addresses that are in the network. Label 0 is advertised for the local routes because of the `explicit-null` configuration. Only the R2 loopback IP address is advertised with label 3. Note that this output is taken before `ldp explicit-null` command was enabled.

R2

```
[edit]
lab@R2# show protocols ldp
track-igp-metric;
interface ge-0/0/4.123;
interface ae0.0;
session 172.30.5.1 {
      authentication-key "$9$mT6AleWXNbEcrvLNY2mfT3/t"; ## SECRET-DATA
}
session 172.30.5.3 {
    authentication-key "$9$wzgGi/9pOIcQF6A0IrlwYgJUH"; ## SECRET-DATA
}
```

Activating the `explicit-null` feature will instruct LDP on R2 to announce a label value of 0 for its loopback IP address.

```
[edit]
lab@R2# set protocols ldp explicit-null

[edit]
lab@R2# commit
commit complete
```

Verify the LDP database on R1 again.

© 2020 Juniper Networks, Inc. All rights reserved.

R1

```
[edit]
lab@R1# run show ldp database
Input label database, 172.30.5.1:0--172.30.5.2:0
  Label      Prefix
 299776      172.30.5.1/32
      0      172.30.5.2/32
 299792      172.30.5.3/32
 299808      172.30.5.4/32

Output label database, 172.30.5.1:0--172.30.5.2:0
  Label      Prefix
      0      172.30.5.1/32
 299776      172.30.5.2/32
 299792      172.30.5.3/32
 299808      172.30.5.4/32

Input label database, 172.30.5.1:0--172.30.5.4:0
  Label      Prefix
 299792      172.30.5.1/32
 299808      172.30.5.2/32
 299776      172.30.5.3/32
      0      172.30.5.4/32

Output label database, 172.30.5.1:0--172.30.5.4:0
  Label      Prefix
      0      172.30.5.1/32
 299776      172.30.5.2/32
 299792      172.30.5.3/32
 299808      172.30.5.4/32
```

By default, Label distribution protocols insert the routes in the `inet.3` table. All LSP routes are installed in this table on the ingress node.

Note: You will likely see slightly different label values since they are dynamically created and assigned by each device in the path.

```
[edit]
lab@R1# run show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.2/32       *[LDP/9] 00:04:07, metric 5
                     > to 172.30.0.2 via ae0.0
172.30.5.3/32       *[LDP/9] 00:04:07, metric 15
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                     > to 172.30.0.2 via ae0.0, Push 299840
172.30.5.4/32      *[LDP/9] 00:04:10, metric 10
                     > to 172.30.0.6 via ge-0/0/4.114, Push 0
```

Three LDP routes are present in routing table `inet.3`. The metric is taken from the IGP route. The label operation and the labels for the LSP are also displayed. This output is displayed before the `explicit-null` option was enabled on R2. This is the reason why for the R2's loopback IP address no label will be pushed. R2 is sending label 3 for its loopback IP address. An Implicit Null label (3) is used for signaling only. R4 is also directly connected to R1 as R2 but has `explicit-null` enabled and you can see that R1 will push label 0 for the packets to R4 loopback IP address. Packets to R3 loopback IP address will be labeled with label 299840 and send to R2. R2 will swap it with label 0 and send it to R3.

Label information base or label switching table is stored in `mpls.0` routing table.

```
[edit]
lab@R1# run show route table mpls.0

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 02:55:28, metric 1
                      Receive
1                  *[MPLS/0] 02:55:28, metric 1
                      Receive
2                  *[MPLS/0] 02:55:28, metric 1
                      Receive
299776             *[LDP/9] 00:04:29, metric 1
                     > to 172.30.0.2 via ae0.0, Pop
299776(S=0)        *[LDP/9] 00:04:29, metric 1
                     > to 172.30.0.2 via ae0.0, Pop
299792             *[LDP/9] 00:04:29, metric 1
                     > to 172.30.0.2 via ae0.0, Swap 299840
299808             *[LDP/9] 00:04:32, metric 1
                     > to 172.30.0.6 via ge-0/0/4.114, Swap 0
299808(S=0)        *[LDP/9] 00:04:32, metric 1
                     > to 172.30.0.6 via ge-0/0/4.114, Pop
```

The first three labels are automatically added after you enable MPLS on the router. Those are the IPv4 and IPv6 explicit null labels. Basically, these are the labels which cause the RE to process the packet. This table is used by LSR routers to switch the packets. Incoming label 299792 will be swapped with label 299840 and send via interface `ae0.0` to next-hop `172.30.0.2`. The other two labels have entries with `(S=0)`. This router is the penultimate-hop router for these two label entries.

© 2020 Juniper Networks, Inc. All rights reserved.

3) Configure IS-IS LDP synchronization.

This step requires to configure the LDP IGP synchronization feature. In this network IS-IS is used as the IGP. Enabling the synchronization feature ensures that LDP must be fully established before IS-IS paths are used for.

Enable LDP synchronization for IS-IS on all interfaces running LDP.

IS-IS configuration is applied to `interfaces all` from previous chapters. The exception are the interfaces to IX on R1 and R2.

```
[edit]
lab@R1# show protocols isis
reference-bandwidth 10g;
topologies ipv6-unicast;
level 2 disable;
level 1 {
    authentication-key "$9$BpLElMg4ZDHmVw2aUH5TBIEyeW"; ## SECRET-DATA
    authentication-type md5;
    wide-metrics-only;
}
interface ge-0/0/5.300 {
    passive;
}
interface all {
    point-to-point;
    bfd-liveness-detection {
        minimum-interval 2000;
        multiplier 3;
    }
}
```

Adding LDP synchronization only to the LDP enabled interfaces requires specific configuration for each interface. Also add the other IS-IS configuration. In this particular case BFD and the IS-IS interface type must be set to point-to-point.

Note: LDP synchronization is supported only on point-to-point interfaces.

a) R1

```
[edit]
lab@R1# set protocols isis interface ge-0/0/4.114 point-to-point


[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set protocols isis interface ge-0/0/4.114 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R1# set protocols isis interface ge-0/0/4.114 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R1# set protocols isis interface ge-0/0/4.114 ldp-synchronization

[edit]
lab@R1# set protocols isis interface ae0.0 point-to-point

[edit]
lab@R1# set protocols isis interface ae0.0 family inet bfd-liveness-detection minimum-
interval 2000

[edit]
lab@R1# set protocols isis interface ae0.0 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R1# set protocols isis interface ae0.0 ldp-synchronization

[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols isis interface ge-0/0/4.123 point-to-point

[edit]
lab@R2# set protocols isis interface ge-0/0/4.123 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R2# set protocols isis interface ge-0/0/4.123 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R2# set protocols isis interface ge-0/0/4.123 ldp-synchronization


[edit]
lab@R2# set protocols isis interface ae0.0 point-to-point
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set protocols isis interface ae0.0 family inet bfd-liveness-detection minimum-
interval 2000

[edit]
lab@R2# set protocols isis interface ae0.0 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R2# set protocols isis interface ae0.0 ldp-synchronization

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols isis interface ge-0/0/4.123 point-to-point

[edit]
lab@R3# set protocols isis interface ge-0/0/4.123 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R3# set protocols isis interface ge-0/0/4.123 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R3# set protocols isis interface ge-0/0/4.123 ldp-synchronization

[edit]
lab@R3# set protocols isis interface ge-0/0/4.134 point-to-point

[edit]
lab@R3# set protocols isis interface ge-0/0/4.134 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R3# set protocols isis interface ge-0/0/4.134 family inet bfd-liveness-detection
multiplier 3


[edit]
lab@R3# set protocols isis interface ge-0/0/4.134 ldp-synchronization
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols isis interface ge-0/0/4.114 point-to-point

[edit]
lab@R4# set protocols isis interface ge-0/0/4.114 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R4# set protocols isis interface ge-0/0/4.114 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R4# set protocols isis interface ge-0/0/4.114 ldp-synchronization

[edit]
lab@R4# set protocols isis interface ge-0/0/4.134 point-to-point

[edit]
lab@R4# set protocols isis interface ge-0/0/4.134 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R4# set protocols isis interface ge-0/0/4.134 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R4# set protocols isis interface ge-0/0/4.134 ldp-synchronization

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols isis interface ge-0/0/4.158 point-to-point

[edit]
lab@R5# set protocols isis interface ge-0/0/4.158 family inet bfd-liveness-detection
minimum-interval 2000
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols isis interface ge-0/0/4.158 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R5# set protocols isis interface ge-0/0/4.158 ldp-synchronization

[edit]
lab@R5# set protocols isis interface ae0.0 point-to-point

[edit]
lab@R5# set protocols isis interface ae0.0 family inet bfd-liveness-detection minimum-
interval 2000

[edit]
lab@R5# set protocols isis interface ae0.0 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R5# set protocols isis interface ae0.0 ldp-synchronization

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols isis interface ge-0/0/4.167 point-to-point

[edit]
lab@R6# set protocols isis interface ge-0/0/4.167 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R6# set protocols isis interface ge-0/0/4.167 family inet bfd-liveness-detection
multiplier 3


[edit]
lab@R6# set protocols isis interface ge-0/0/4.167 ldp-synchronization

[edit]
lab@R6# set protocols isis interface ae0.0 point-to-point

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R6# set protocols isis interface ae0.0 family inet bfd-liveness-detection minimum-
interval 2000

[edit]
lab@R6# set protocols isis interface ae0.0 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R6# set protocols isis interface ae0.0 ldp-synchronization

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set protocols isis interface ge-0/0/4.158 point-to-point

[edit]
lab@R7# set protocols isis interface ge-0/0/4.158 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R7# set protocols isis interface ge-0/0/4.158 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R7# set protocols isis interface ge-0/0/4.158 ldp-synchronization

[edit]
lab@R7# set protocols isis interface ge-0/0/4.167 point-to-point

[edit]
lab@R7# set protocols isis interface ge-0/0/4.167 family inet bfd-liveness-detection
minimum-interval 2000


[edit]
lab@R7# set protocols isis interface ge-0/0/4.167 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R7# set protocols isis interface ge-0/0/4.167 ldp-synchronization

[edit]
lab@R7# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols isis interface ge-0/0/4.158 point-to-point

[edit]
lab@R8# set protocols isis interface ge-0/0/4.158 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R8# set protocols isis interface ge-0/0/4.158 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R8# set protocols isis interface ge-0/0/4.158 ldp-synchronization

[edit]
lab@R8# set protocols isis interface ge-0/0/4.178 point-to-point

[edit]
lab@R8# set protocols isis interface ge-0/0/4.178 family inet bfd-liveness-detection
minimum-interval 2000

[edit]
lab@R8# set protocols isis interface ge-0/0/4.178 family inet bfd-liveness-detection
multiplier 3

[edit]
lab@R8# set protocols isis interface ge-0/0/4.178 ldp-synchronization

[edit]
lab@R8# commit
commit complete
```

Display the IS-IS interface details.

```
[edit]
lab@R1# run show isis interface ae0.0 extensive
IS-IS interface database:
ae0.0
  Index: 327, State: 0x6, Circuit id: 0x1, Circuit type: 1
  LSP interval: 100 ms, CSNP interval: 20 s, Loose Hello padding, IIH max size:
1492
  Adjacency advertisement: Advertise, Layer2-map: Disabled
  LDP sync state: in sync, for: 00:09:31, reason: LDP session up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
config holdtime: infinity
Interface Group Holddown Delay: 20 s, remaining: 0 s
Level 1
  Adjacencies: 1, Priority: 64, Metric: 5
  Hello Interval: 9.000 s, Hold Time: 27 s
  IPV6 UnicastMetric: 5
```

R1 is in Sync on interface ae0.0

4) Configure R1 and R2 to inject the IX facing subnet to LDP.

As it was stated in the overview, by default LDP only announces /32 routes loopback IP addresses.

To accomplish this, you have to use an LDP egress policy. Where you will match the routes in inet.0 for which you want to build an LSP.

Note: LDP has a capability to filter label bindings between LDP neighbors. This can be achieved with import and export policies. The difference is that with export policy you can control which bindings are advertised to a specific or all neighbors, whereas with egress policy you say which routes to have label bindings. You can have egress and export policy at the same time.

Creating egress policy on R1 and R2.

a) R1

```
 [edit]
lab@R1# set policy-options policy-statement ldp-routes term 1 from route-filter
192.168.1.0/24 exact

[edit]
lab@R1# set policy-options policy-statement ldp-routes term 1 from route-filter
172.30.5.1/32 exact

[edit]
lab@R1# set policy-options policy-statement ldp-routes term 1 then accept

[edit]
lab@R1# set protocols ldp egress-policy ldp-routes

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) R2

```
[edit]
lab@R2# set policy-options policy-statement ldp-routes term 1 from route-filter
192.168.1.0/24 exact

[edit]
lab@R2# set policy-options policy-statement ldp-routes term 1 from route-filter
172.30.5.2/32 exact

[edit]
lab@R2# set policy-options policy-statement ldp-routes term 1 then accept

[edit]
lab@R2# set protocols ldp egress-policy ldp-routes

[edit]
lab@R2# commit
commit complete
```

You must include the loopback IP address in the match statement of the policy. Otherwise only the direct subnet to IX will be advertised with label.

Verify the result of the policy.

```
[edit]
lab@R1# run show ldp database
Input label database, 172.30.5.1:0--172.30.5.2:0
  Label     Prefix
 299968     172.30.5.1/32
      0     172.30.5.2/32
 299936     172.30.5.3/32
 299952     172.30.5.4/32
      0     192.168.1.0/24

Output label database, 172.30.5.1:0--172.30.5.2:0
  Label     Prefix
      0     172.30.5.1/32
 299824     172.30.5.2/32
 299840     172.30.5.3/32
 299856     172.30.5.4/32
      0     192.168.1.0/24

Input label database, 172.30.5.1:0--172.30.5.4:0
  Label     Prefix
 299840     172.30.5.1/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 299808     172.30.5.2/32
 299776     172.30.5.3/32
      0     172.30.5.4/32
 299840     192.168.1.0/24

Output label database, 172.30.5.1:0--172.30.5.4:0
  Label     Prefix
      0     172.30.5.1/32
 299824     172.30.5.2/32
 299840     172.30.5.3/32
 299856     172.30.5.4/32
      0     192.168.1.0/24
```

The direct subnet 192.168.1.0/24 is advertised as labeled to the LDP neighbors. Also, the same route is received with label value 0 from R2 as well.

5) You must make sure that each FEC advertised by R1 and R2 are reachable by a separate LSP.

By default, Junos advertises the same label for all prefixes with the same next hop and/or label advertised from the next-hop router.

To change the default behavior, you must configure all routers in the network with `ldp deaggregate` to signal different labels for those prefixes. This helps with load-balancing of the MPLS packets.

Before we activate LDP deaggregate we first check the label advertised to R4.

```
[edit]
lab@R4# run show ldp database
Input label database, 172.30.5.4:0--172.30.5.1:0
  Label     Prefix
      0     172.30.5.1/32
 299824     172.30.5.2/32
 299840     172.30.5.3/32
 299856     172.30.5.4/32
      0     192.168.1.0/24

Output label database, 172.30.5.4:0--172.30.5.1:0
  Label     Prefix
 299840     172.30.5.1/32
 299808     172.30.5.2/32
 299776     172.30.5.3/32
      0     172.30.5.4/32
 299840     192.168.1.0/24
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Input label database, 172.30.5.4:0--172.30.5.3:0
  Label     Prefix
 299824     172.30.5.1/32
 299776     172.30.5.2/32
      0     172.30.5.3/32
 299808     172.30.5.4/32
 299776     192.168.1.0/24

Output label database, 172.30.5.4:0--172.30.5.3:0
  Label     Prefix
 299840     172.30.5.1/32
 299808     172.30.5.2/32
 299776     172.30.5.3/32
      0     172.30.5.4/32
 299840     192.168.1.0/24
```

You can see that R3 is advertising the same label for both prefixes advertised by R2 to R3.

Activate the LDP deaggregate feature on all routers in the network.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols ldp deaggregate

[edit]
lab@R1# commit
commit complete
```

Verify again the LDP database on R4.

```
[edit]
lab@R4# run show ldp database
Input label database, 172.30.5.4:0--172.30.5.1:0
  Label     Prefix
      0     172.30.5.1/32
 299824     172.30.5.2/32
 299840     172.30.5.3/32
 299856     172.30.5.4/32
      0     192.168.1.0/24

Output label database, 172.30.5.4:0--172.30.5.1:0
  Label     Prefix
 299904     172.30.5.1/32
 299872     172.30.5.2/32
 299856     172.30.5.3/32
      0     172.30.5.4/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 299888     192.168.1.0/24

Input label database, 172.30.5.4:0--172.30.5.3:0
  Label     Prefix
 299872     172.30.5.1/32
 299888     172.30.5.2/32
      0     172.30.5.3/32
 299856     172.30.5.4/32
 299840     192.168.1.0/24

Output label database, 172.30.5.4:0--172.30.5.3:0
  Label     Prefix
 299904     172.30.5.1/32
 299872     172.30.5.2/32
 299856     172.30.5.3/32
      0     172.30.5.4/32
 299888     192.168.1.0/24
```

After the change R3 is advertising different label for the prefixes received from R2.

Task 2: RSVP Configuration

This task introduces RSVP to the network. RSVP provides Traffic-Engineering and MPLS protection capabilities for the services within the network.

RSVP Overview

The Resource Reservation Protocol (RSVP) is an Internet standard for the reservation of resource in the network. It is adopted to be used as a label distribution protocol. RSVP is not a routing protocol; its use is limited only to label distribution. IGP paths are used to discovering nodes in the network.

RSVP request resources in a unidirectional manner. Reservations are signaled on an ingress LSR (sender) towards an egress LSR (receiver).

RSVP is capable of building traffic-engineering path through the network based on given restrictions. To support this functionality RSVP uses the following messages:

- *Path* – path message is sent by the ingress router to request LSP to be created. It contains the destination address of the egress LSR but is processed by all routes on the path. The Path through the network can be manually configured, can be computed based on different constrains or can follow the IGP path. Mandatory and optional objects are carried by the Path message.

© 2020 Juniper Networks, Inc. All rights reserved.

- *Resv* – A resv message is sent back to the ingress LSR in request of the Path message. Resv message is reserving the resources for the LSP, including the label. This message follows the recorded path from the path message. As with Path message, Resv messages can carry different optional objects.
- *PathTear* – this message deletes the path and all relevant reservations. It is sent by the sender and is always directed downstream to the receiver.
- *PathErr* – the purpose of this message is to report errors with the Path messages along the path.
- *ResvErr* – reports errors of processing of the Resv messages.

Path and Resv messages are sent periodically to refresh the soft state. The state is deleted if the Path or Resv messages are not received before the expiration cleanup interval.

The Junos RSVP implementation also supports MD5 authentication and hello keepalives. Hello messages account for faster router failures detection. They are sent by default every nine seconds.

As it was stated RSVP provided traffic-engineering capabilities to the network. This is done with help of Explicit Route Objects (ERO), which are included in RSVP Path messages. The ERO can be configured manually where the administrator can specify each hop explicitly. The other option supported by RSVP is automatic based on traffic-engineering information provided by the IGP. IS-IS and OSPF support the exchange of additional information for the links in the network. These are available bandwidth and link coloring. The information is stored in traffic-engineering database (TED). The algorithm used to calculate the Constrained Path through the network is a modified version of the SPF algorithm, called Constrained Shortest Path First (CSPF) algorithm.

The following steps are used by the CSPF algorithm:

- Prune links with insufficient bandwidth.
- Prune links that do not contain the included colors.
- Prune links that contain an excluded color.
- Calculate shortest path from ingress to egress respecting the configured ERO if present.
- If equal cost paths exist, choose the path whose last hop address equals the LSP destination.
- If there are still remaining equal cost paths choose the one with fewer hops. If again more than one paths remain apply the CSPF load-balancing rule configured on the LSP (least fill, most fill or random).

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure RSVP on all core interfaces and define interface bandwidths.

As with LDP protocol you have to enable RSVP on the desired interfaces. You already enabled all interfaces to accept labeled packets with Configure `family mpls`.

Note: Routers on the path of the LSP should have enough bandwidth on the ingress interface where the Path message arrives. By default, the physical speed of the interface is used as RSVP interface bandwidth.

In this task, you are asked to configure a RSVP bandwidth value lower than the physical interface speed, with the exception for Aggregated Ethernet interfaces. You must configure RSVP session authentication as well.

a) R1

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.114 authentication-key jncie

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.114 bandwidth 333m

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.118 authentication-key jncie

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.118 bandwidth 333m

[edit]
lab@R1# set protocols rsvp interface ae0.0 authentication-key jncie

[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.123 authentication-key jncie

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.123 bandwidth 333m

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.127 authentication-key jncie

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# set protocols rsvp interface ge-0/0/4.127 bandwidth 333m

[edit]
lab@R2# set protocols rsvp interface ae0.0 authentication-key jncie

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.123 authentication-key jncie

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.123 bandwidth 333m

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.134 authentication-key jncie

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.134 bandwidth 333m

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.136 authentication-key jncie

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.136 bandwidth 333m

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.114 authentication-key jncie

[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.114 bandwidth 333m

[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.134 authentication-key jncie

[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.134 bandwidth 333m
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.145 authentication-key jncie

[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.145 bandwidth 333m

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.145 authentication-key jncie

[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.145 bandwidth 333m

[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.158 authentication-key jncie

[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.158 bandwidth 333m

[edit]
lab@R5# set protocols rsvp interface ae0.0 authentication-key jncie

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.136 authentication-key jncie

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.136 bandwidth 333m

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.167 authentication-key jncie

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.167 bandwidth 333m
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols rsvp interface ae0.0 authentication-key jncie

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.127 authentication-key jncie

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.127 bandwidth 333m

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.167 authentication-key jncie

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.167 bandwidth 333m

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.178 authentication-key jncie

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.178 bandwidth 333m

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.118 authentication-key jncie

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.118 bandwidth 333m

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.158 authentication-key jncie

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.158 bandwidth 333m

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# set protocols rsvp interface ge-0/0/4.178 authentication-key jncie

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.178 bandwidth 333m

[edit]
lab@R8# commit
commit complete
```

Verify the RSVP interface database.

```
[edit]
lab@R1# run show rsvp interface
RSVP interface: 3 active
                  Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
ae0.0       Up           0   100%  2Gbps       2Gbps       0bps        0bps
ge-0/0/4.114Up           0   100%  333Mbps     333Mbps     0bps        0bps
ge-0/0/4.118Up           0   100%  333Mbps     333Mbps     0bps        0bps
```

You can see that the static bandwidth on the GE interfaces is changed to the configured value of 333 Mbps. For interface ae0.0 all the bandwidth can be used.

The extensive option of the RSVP interface database shows more information for the interfaces like the hello interval and protection.

```
[edit]
lab@R1# run show rsvp interface ae0.0 extensive
RSVP interface: 3 active
ae0.0 Index 69, State Ena/Up
  Authentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.30.0.1
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 2Gbps
  ct0: StaticBW 2Gbps, AvailableBW 2Gbps
    MaxAvailableBW 2Gbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off
```

All routers should be configured in the same way. RSVP sessions have a soft state, which means that an LSP must be signaled/ configured before this show command will display output.

© 2020 Juniper Networks, Inc. All rights reserved.

2) Configure administrative groups on all routers.

In Table 10 the values are given for the administrative groups that have to be configured on the MPLS enabled interfaces. Administrative groups provide the ability to mark the interfaces with similar characteristics for the same group. Based on these groups you can restrict the LSPs to use a certain path in the network. Marking the interfaces is known as coloring but is actually based on setting the affinity bits on interfaces. Up to 32 administrative groups are supported, ranging from 0 to 31. transferred by the IGP. The administrative groups / affinity bits are used by the CSPF algorithm for traffic engineering purposes.

You must define a name for the group and assign an affinity bit to it. Different groups can be assigned to interfaces. Usually colors are used for names of the groups.

For this task you must configure green and red interfaces. The group configuration must be identical on all routers in the MPLS network.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols mpls admin-groups green 0

[edit]
lab@R1# set protocols mpls admin-groups red 1
```

Next you must assign the groups to the interfaces according to Figure 15. Previously you configured `protocols mpls - interfaces all`. However, now you must make the statement more specific as for each interface a different group might be needed.

Router	Interface	Admin. Group
R1	ge-0/0/4.114	green
	ge-0/0/4.118	red
	ae0.0	green, red
R2	ge-0/0/4.127	green
	ge-0/0/4.123	red
	ae0.0	green, red
R3	ge-0/0/4.134	green, red
	ge-0/0/4.136	green
	ge-0/0/4.123	red
R4	ge-0/0/4.134	green, red
	ge-0/0/4.114	green

© 2020 Juniper Networks, Inc. All rights reserved.

R5	ge-0/0/4.145	red
	ge-0/0/4.158	green
	ge-0/0/4.145	red
	ae0.0	green, red
R6	ge-0/0/4.136	green
	ge-0/0/4.167	red
	ae0.0	green, red
R7	ge-0/0/4.178	green, red
	ge-0/0/4.127	green
	ge-0/0/4.167	red
R8	ge-0/0/4.178	green, red
	ge-0/0/4.158	green
	ge-0/0/4.118	red

Table 37

a) R1

```
 [edit]
lab@R1# set protocols mpls interface ge-0/0/4.114 admin-group green

[edit]
lab@R1# set protocols mpls interface ge-0/0/4.118 admin-group red

[edit]
lab@R1# set protocols mpls interface ae0.0 admin-group [green red]


[edit]
lab@R1# commit
commit complete
```

Interface **ae0.0** is assigned with both groups.

b) R2

```
[edit]
lab@R2# set protocols mpls interface ge-0/0/4.123 admin-group red

[edit]
lab@R2# set protocols mpls interface ge-0/0/4.127 admin-group green

[edit]
lab@R2# set protocols mpls interface ae0.0 admin-group [green red]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols mpls interface ge-0/0/4.123 admin-group red

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.136 admin-group green

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.134 admin-group [green red]

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols mpls interface ge-0/0/4.145 admin-group red

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.114 admin-group green

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.134 admin-group [green red]

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols mpls interface ge-0/0/4.145 admin-group red

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.158 admin-group green

[edit]
lab@R5# set protocols mpls interface ae0.0 admin-group [green red]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols mpls interface ge-0/0/4.167 admin-group red

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.136 admin-group green

[edit]
lab@R6# set protocols mpls interface ae0.0 admin-group [green red]

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set protocols mpls interface ge-0/0/4.167 admin-group red

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.127 admin-group green

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.178 admin-group [green red]

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols mpls interface ge-0/0/4.118 admin-group red

[edit]
lab@R8# set protocols mpls interface ge-0/0/4.158 admin-group green

[edit]
lab@R8# set protocols mpls interface ge-0/0/4.178 admin-group [green red]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# commit
commit complete
```

Verify the administrative group on the interfaces.

```
[edit]
lab@R1# run show mpls interface
Interface          State       Administrative groups (x: extended)
ae0.0              Up          red
                               green
ge-0/0/4.114       Up          green
ge-0/0/4.118       Up          red
...
```

Administrative groups and interface assignments/coloring should be configured on all routers according Table 10 and Figure 9.

After configuring all routers, you can check the traffic-engineering database (TED). Traffic engineering support is enabled by default for IS-IS.

```
[edit]
lab@R1# run show ted database extensive
TED database: 8 ISIS nodes 8 INET nodes
NodeID: R1.00(172.30.5.1)
  Type: Rtr, Age: 37 secs, LinkIn: 3, LinkOut: 3
  Protocol: IS-IS(1)
    To: R8.00(172.30.5.8), Local: 172.30.0.9, Remote: 172.30.0.10
      Local interface index: 74, Remote interface index: 70
      Color: 0x2 red
      Metric: 10
      Static BW: 333Mbps
      Reservable BW: 333Mbps
      Available BW [priority] bps:
          [0] 333Mbps      [1] 333Mbps      [2] 333Mbps      [3] 333Mbps
          [4] 333Mbps      [5] 333Mbps      [6] 333Mbps      [7] 333Mbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 333Mbps      [1] 333Mbps      [2] 333Mbps      [3] 333Mbps
          [4] 333Mbps      [5] 333Mbps      [6] 333Mbps      [7] 333Mbps
    To: R4.00(172.30.5.4), Local: 172.30.0.5, Remote: 172.30.0.6
      Local interface index: 73, Remote interface index: 70
      Color: 0x1 green
      Metric: 10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Static BW: 333Mbps
  Reservable BW: 333Mbps
  Available BW [priority] bps:
      [0] 333Mbps      [1] 333Mbps      [2] 333Mbps      [3] 333Mbps
      [4] 333Mbps      [5] 333Mbps      [6] 333Mbps      [7] 333Mbps
  Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
      [0] 333Mbps      [1] 333Mbps      [2] 333Mbps      [3] 333Mbps
      [4] 333Mbps      [5] 333Mbps      [6] 333Mbps      [7] 333Mbps
To: R2.00(172.30.5.2), Local: 172.30.0.1, Remote: 172.30.0.2
  Local interface index: 69, Remote interface index: 69
  Color: 0x3 red green
  Metric: 5
  Static BW: 2Gbps
  Reservable BW: 2Gbps
  Available BW [priority] bps:
      [0] 2Gbps        [1] 2Gbps        [2] 2Gbps        [3] 2Gbps
      [4] 2Gbps        [5] 2Gbps        [6] 2Gbps        [7] 2Gbps
  Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
      [0] 2Gbps        [1] 2Gbps        [2] 2Gbps        [3] 2Gbps
      [4] 2Gbps        [5] 2Gbps        [6] 2Gbps        [7] 2Gbps
NodeID: R2.00(172.30.5.2)
  Type: Rtr, Age: 37 secs, LinkIn: 3, LinkOut: 3
  Protocol: IS-IS(1)
    To: R1.00(172.30.5.1), Local: 172.30.0.2, Remote: 172.30.0.1
      Local interface index: 69, Remote interface index: 69
      Color: 0x3 red green
      Metric: 5
      Static BW: 2Gbps
      Reservable BW: 2Gbps
      Available BW [priority] bps:
      [0] 2Gbps        [1] 2Gbps        [2] 2Gbps        [3] 2Gbps
      [4] 2Gbps        [5] 2Gbps        [6] 2Gbps        [7] 2Gbps
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
      [0] 2Gbps        [1] 2Gbps        [2] 2Gbps        [3] 2Gbps
      [4] 2Gbps        [5] 2Gbps        [6] 2Gbps        [7] 2Gbps
To: R7.00(172.30.5.7), Local: 172.30.0.17, Remote: 172.30.0.18
  Local interface index: 74, Remote interface index: 70
  Color: 0x1 green
  Metric: 10
  Static BW: 333Mbps
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Reservable BW: 333Mbps
  Available BW [priority] bps:
      [0] 333Mbps      [1] 333Mbps      [2] 333Mbps      [3] 333Mbps
      [4] 333Mbps      [5] 333Mbps      [6] 333Mbps      [7] 333Mbps
  Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
      [0] 333Mbps      [1] 333Mbps      [2] 333Mbps      [3] 333Mbps
      [4] 333Mbps      [5] 333Mbps      [6] 333Mbps      [7] 333Mbps
To: R3.00(172.30.5.3), Local: 172.30.0.13, Remote: 172.30.0.14
  Local interface index: 73, Remote interface index: 70
  Color: 0x2 red
  Metric: 10
  Static BW: 333Mbps
  Reservable BW: 333Mbps
  Available BW [priority] bps:
      [0] 333Mbps      [1] 333Mbps      [2] 333Mbps      [3] 333Mbps
      [4] 333Mbps      [5] 333Mbps      [6] 333Mbps      [7] 333Mbps
  Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
      [0] 333Mbps      [1] 333Mbps      [2] 333Mbps      [3] 333Mbps
      [4] 333Mbps      [5] 333Mbps      [6] 333Mbps      [7] 333Mbps
```

The output shows the detailed view of the TED on R1. It is populated with all routers running IS-IS and have RSVP enabled interfaces. Detailed traffic engineering information is included for each link like, administrative groups, metrics, static bandwidth, reservable and available bandwidth by priority level.

Let's show the content of the IS-IS LSP advertised by R1 after configuring administrative groups.

```
[edit]
lab@R2# run show isis database level 1 R1.00-00 extensive
IS-IS level 1 link-state database:

R1.00-00 Sequence: 0x1b0, Checksum: 0x7eda, Lifetime: 1188 secs
   IPV4 Unicast IS neighbor: R2.00        Metric:         5
     Two-way fragment: R2.00-00, Two-way first fragment: R2.00-00
   IPV4 Unicast IS neighbor: R4.00      Metric:        10
     Two-way fragment: R4.00-00, Two-way first fragment: R4.00-00
   IPV4 Unicast IS neighbor: R8.00       Metric:        10
     Two-way fragment: R8.00-00, Two-way first fragment: R8.00-00
   IPV4 Unicast IS neighbor: route-reflector.00 Metric:        10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Two-way fragment: route-reflector.00-00, Two-way first fragment: route-
reflector.00-00
  IPV6 Unicast IS neighbor: R2.00        Metric:        5
    Two-way fragment: R2.00-00, Two-way first fragment: R2.00-00
  IPV6 Unicast IS neighbor: R8.00       Metric:       10
    Two-way fragment: R8.00-00, Two-way first fragment: R8.00-00
  IP IPV4 Unicast prefix: 172.30.0.0/30      Metric:        5 Internal Up
  IP IPV4 Unicast prefix: 172.30.0.4/30      Metric:       10 Internal Up
  IP IPV4 Unicast prefix: 172.30.0.8/30      Metric:       10 Internal Up
  IP IPV4 Unicast prefix: 172.30.0.64/30     Metric:       10 Internal Up
  IP IPV4 Unicast prefix: 172.30.5.1/32      Metric:        0 Internal Up
  IP IPV4 Unicast prefix: 192.168.1.0/24     Metric:       10 Internal Up
  V6 IPV6 Unicast prefix: fd17:f0f4:f691:5::1/128 Metric:        0 Internal Up

Header: LSP ID: R1.00-00, Length: 482 bytes
  Allocated length: 482 bytes, Router ID: 172.30.5.1
  Remaining lifetime: 1188 secs, Level: 1, Interface: 69
  Estimated free bytes: 70, Actual free bytes: 0
  Aging timer expires in: 1188 secs
  Protocols: IP, IPv6

Packet: LSP ID: R1.00-00, Length: 482 bytes, Lifetime : 1196 secs
  Checksum: 0x7eda, Sequence: 0x1b0, Attributes: 0x1 <L1>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 18, Packet version: 1, Max area: 0

TLVs:
  Area address: 49.0001 (3)
  LSP Buffer Size: 1492
  Speaks: IP
  Speaks: IPV6
  Topology: unicast
  Topology: ipv6 unicast
  IP router id: 172.30.5.1
  IP address: 172.30.5.1
  Hostname: R1
  IS extended neighbor: R2.00, Metric: default 5
    IP address: 172.30.0.1
    Neighbor's IP address: 172.30.0.2
    Local interface index: 69, Remote interface index: 69
    Current reservable bandwidth:
      Priority 0 : 2Gbps
      Priority 1 : 2Gbps
      Priority 2 : 2Gbps
      Priority 3 : 2Gbps
      Priority 4 : 2Gbps
      Priority 5 : 2Gbps
      Priority 6 : 2Gbps
      Priority 7 : 2Gbps
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Maximum reservable bandwidth: 2Gbps
    Maximum bandwidth: 2Gbps
    Administrative groups:  0x3 red green
  IS extended neighbor: R4.00, Metric: default 10
    IP address: 172.30.0.5
    Neighbor's IP address: 172.30.0.6
    Local interface index: 73, Remote interface index: 70
    Current reservable bandwidth:
      Priority 0 : 333Mbps
      Priority 1 : 333Mbps
      Priority 2 : 333Mbps
      Priority 3 : 333Mbps
      Priority 4 : 333Mbps
      Priority 5 : 333Mbps
      Priority 6 : 333Mbps
      Priority 7 : 333Mbps
    Maximum reservable bandwidth: 333Mbps
    Maximum bandwidth: 333Mbps
    Administrative groups:  0x1 green
  IS extended neighbor: R8.00, Metric: default 10
    IP address: 172.30.0.9
    Neighbor's IP address: 172.30.0.10
    Local interface index: 74, Remote interface index: 70
    Current reservable bandwidth:
      Priority 0 : 333Mbps
      Priority 1 : 333Mbps
      Priority 2 : 333Mbps
      Priority 3 : 333Mbps
      Priority 4 : 333Mbps
      Priority 5 : 333Mbps
      Priority 6 : 333Mbps
      Priority 7 : 333Mbps
    Maximum reservable bandwidth: 333Mbps
    Maximum bandwidth: 333Mbps
    Administrative groups:  0x2 red
  IS extended neighbor: route-reflector.00, Metric: default 10
    IP address: 172.30.0.65
    Neighbor's IP address: 172.30.0.66
    Local interface index: 75, Remote interface index: 64
  IPV6 UnicastIS neighbor: R2.00, Metric: default 5
  IPV6 UnicastIS neighbor: R8.00, Metric: default 10
  IP extended prefix: 172.30.5.1/32 metric 0 up
  IP extended prefix: 172.30.0.64/30 metric 10 up
  IP extended prefix: 172.30.0.8/30 metric 10 up
  IP extended prefix: 172.30.0.4/30 metric 10 up
  IP extended prefix: 192.168.1.0/24 metric 10 up
  IP extended prefix: 172.30.0.0/30 metric 5 up
  IPV6 UnicastIPv6 prefix: fd17:f0f4:f691:5::1/128 Metric 0 Up
  Authentication data: 17 bytes
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
No queued transmissions
```

Available bandwidth for reservation and the administrative group for each interface is advertised by R1.

3) Configure RSVP signaled LSP.

Now you must configure the RSVP signaled LSPs according Table 11 and Figure 10. CSPF will be used to create the ERO. BFD must to be configured as well.

Ingress	Egress	LSP ID
R1	R8	A
R1	R6	C
R2	R7	E
R2	R5	G
R3	R8	J
R3	R8	L
R3	R6	Q
R4	R7	N
R4	R7	P
R4	R5	S
R5	R2	H
R5	R4	T
R6	R1	D
R6	R3	R
R7	R2	F
R7	R4	M
R7	R4	O
R8	R1	B
R8	R3	I
R8	R3	K

Table 11

a) R1

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 to 172.30.5.8

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set protocols mpls label-switched-path R1-to-R8 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 to 172.30.5.6

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R1# commit
commit complete
```

RSVP signaled LSP are configured with `label-switched-path` under `family mpls` hierarchy. The `to` statement is mandatory as it tells RSVP the IP address of the egress router. BFD is configured as the OAM feature for the LSP.

After you commit, ERO is calculated by CSPF and the result is passed to RSVP for signaling the LSP based. The RSVP sessions are created between the MPLS routers.

```
[edit]
lab@R1# run show rsvp session
Ingress RSVP: 2 sessions
To              From            State Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.1      Up    0  1 FF       -    299840 R1-to-R6
172.30.5.8      172.30.5.1      Up    0  1 FF       -         3 R1-to-R8
Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Both ingress sessions for the configured LSPs are UP. Display the extensive view of the command above provides more details for the both LSPs.

```
[edit]
lab@R1# run show rsvp session extensive
Ingress RSVP: 2 sessions

172.30.5.6
  From: 172.30.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Recovery label received: -, Recovery label sent: 299824
  Resv style: 1 FF, Label in: -, Label out: 299824
  Time left:    -, Since: Thu Sep 19 21:53:42 2013
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 59321 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 172.30.0.2 (ae0.0) 13 pkts
  RESV rcvfrom: 172.30.0.2 (ae0.0) 13 pkts
  Explct route: 172.30.0.2 172.30.0.18 172.30.0.41
  Record route: <self> 172.30.0.2 172.30.0.18 172.30.0.41

172.30.5.8
  From: 172.30.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R8, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left:    -, Since: Thu Sep 19 21:48:05 2013
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 59318 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 172.30.0.10 (ge-0/0/4.118) 22 pkts
  RESV rcvfrom: 172.30.0.10 (ge-0/0/4.118) 21 pkts
  Explct route: 172.30.0.10
  Record route: <self> 172.30.0.10
Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Many details are described by the extensive view of the RSVP sessions. The `Explct route` object is showing the calculated / configured ERO. The `Record route` object is the path that the path message traverses. The `Label out` and `Label in` objects are shown as well.

Another command which is useful to receive information about the RSVP signaled LSP is the `show mpls lsp`.

[edit]

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# run show mpls lsp
Ingress LSP: 2 sessions
To              From            State Rt P    ActivePath    LSPname
172.30.5.6      172.30.5.1      Up     0 *                  R1-to-R6
172.30.5.8      172.30.5.1      Up     0 *                  R1-to-R8
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Verify the routing table for routes inserted by RSVP for the two LSPs.

```
[edit]
lab@R1# run show route protocol rsvp

inet.0: 832 destinations, 1415 routes (670 active, 0 holddown, 324 hidden)

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.6/32       *[RSVP/7/1] 00:01:11, metric 25
                     > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R6
172.30.5.8/32       *[RSVP/7/1] 00:06:48, metric 10
                     > to 172.30.0.10 via ge-0/0/4.118, label-switched-path R1-to-
R8

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

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

inet6.0: 31 destinations, 32 routes (31 active, 0 holddown, 0 hidden)
```

As expected, the loopback IP addresses of R6 and R8 are inserted in the routing table `inet.3` by RSVP. The next hop is through the LSPs.

The output is very much the same as the output from the command `show rsvp session`. It displays the configured and active MPLS LSPs. The extensive output shows details for the LSPs and the log of the events related to them.

```
[edit]
lab@R1# run show mpls lsp extensive
Ingress LSP: 2 sessions
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.6 S 172.30.0.30 S 172.30.0.34 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.6 172.30.0.30 172.30.0.34
    OAM state : BFD session not up LSP-ping not up
    5 Sep 19 21:48:05.683 Selected as active path
    4 Sep 19 21:48:05.682 Record Route:  172.30.0.6 172.30.0.30 172.30.0.34
    3 Sep 19 21:48:05.682 Up
    2 Sep 19 21:48:05.605 Originate Call
    1 Sep 19 21:48:05.605 CSPF: computation result accepted  172.30.0.6
172.30.0.30 172.30.0.34
  Created: Thu Sep 19 21:48:05 2013

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R8
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 10)
 172.30.0.10 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.10
    OAM state : BFD session not up LSP-ping not up
    5 Sep 19 21:48:05.616 Selected as active path
    4 Sep 19 21:48:05.615 Record Route:  172.30.0.10
    3 Sep 19 21:48:05.615 Up
    2 Sep 19 21:48:05.603 Originate Call
    1 Sep 19 21:48:05.603 CSPF: computation result accepted  172.30.0.10
  Created: Thu Sep 19 21:48:05 2013
```

The LSPs are UP after the CSPF computed route is sent to RSVP and the LSP is signaled.

© 2020 Juniper Networks, Inc. All rights reserved.

However, the BFD session is DOWN. Let's verify the BFD sessions table.

```
[edit]
lab@R1# run show bfd session
                                                  Detect   Transmit
Address          State     Interface      Time     Interval  Multiplier
127.0.0.1        Down      ae0.0          9.000     3.000         3
127.0.0.1        Down      ge-0/0/4.118   9.000     3.000         3
172.30.0.2       Up        ae0.0          9.000     3.000         3
172.30.0.6       Up        ge-0/0/4.114   9.000     3.000         3
172.30.0.10      Up        ge-0/0/4.118   9.000     3.000         3
172.30.0.66      Up        ge-0/0/4.206   9.000     3.000         3
172.30.5.41      Up        ge-0/0/4.206   9.000     3.000                 3
```

The reason for that could be that the BFD packets are dropped. In Chapter One you configured a firewall filter on the RE to log dropped packet. Let's check the log of the filter.

```
[edit]
lab@R1# run show firewall log detail
Time of Log: 2013-09-19 23:34:26 CEST, Filter: pfe, Filter action: discard, Name
of interface: ge-0/0/4.118
Name of protocol: UDP, Packet Length: 100, Source address: 172.30.5.8:49832,
Destination address: 127.0.0.1:3503
```

Note: The output is taken after the LSP to R1 is configured on R8 and it is trying to bring up BFD session.

From the log, it is clear that the packets sent from R8 for BFD are dropped by the filter. BFD for MPLS LSPs is using MPLS ping to verify the data path of the LSPs. For proper operation of MPLS ping you must configure the IP address from the reserved subnet 127/8. By default, MPLS ping uses the address 127.0.0.1 as is seen from the log. MPLS ping uses UDP for transporting messages and the port number is 3503.

First the IP address 127.0.0.1 must be configured on the loopback interface on all routers. Secondly the filter protecting the RE must be changed to accept the MPLS ping packets.

Repeat this on all eight devices. Substitute the correct loopback address for the primary interface.

```
[edit]
lab@R1# set interfaces lo0.0 family inet address 127.0.0.1/32

[edit]
lab@R1# set interfaces lo0.0 family inet address 172.30.5.1/32 primary
```


© 2020 Juniper Networks, Inc. All rights reserved.

Making the address primary ensures that this IP address would be used for sourcing packets from the RE.

```
[edit]
lab@R1# set firewall family inet filter protect-re term lsp-ping from protocol udp

[edit]
lab@R1# set firewall family inet filter protect-re term lsp-ping from port 3503

[edit]
lab@R1# set firewall family inet filter protect-re term lsp-ping then policer re-
policer

[edit]
lab@R1# set firewall family inet filter protect-re term lsp-ping then accept

[edit]
lab@R1# insert firewall family inet filter protect-re term lsp-ping before term last

[edit]
lab@R1# commit
commit complete
```

The new term must be inserted before the **`term last`**, which is rejecting all the packets not matched by the terms before.

After all routers are configured in the same way, you can verify again the MPLS LSP and the BFD sessions.

Note: The actual path the LSPs use through the network might be different than the ones shown in the examples.

Note: You may need to wait a minute for BFD communication and synchronization to complete.

```
[edit]
lab@R1# run show mpls lsp name R1-to-R8 extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R8
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.18 S 172.30.0.46 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.2 172.30.0.18 172.30.0.46
    OAM state : BFD session up  LSP-ping up
    8 Sep 21 09:53:16.770 LSP ping failure on path[35 times]
    7 Sep 21 08:27:27.106 Selected as active path
    6 Sep 21 08:27:27.104 Record Route:  172.30.0.2 172.30.0.18 172.30.0.46
    5 Sep 21 08:27:27.104 Up
    4 Sep 21 08:27:27.027 Originate Call
    3 Sep 21 08:27:27.027 CSPF: computation result accepted  172.30.0.2
172.30.0.18 172.30.0.46
    2 Sep 21 08:26:57.483 CSPF failed: no route toward 172.30.5.8[6 times]
    1 Sep 21 08:24:02.218 CSPF: could not determine self
  Created: Sat Sep 21 08:22:07 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 4 sessions
Total 0 displayed, Up 0, Down 0
```

```
[edit]
lab@R1# run show bfd session
                                                  Detect   Transmit
Address          State     Interface              Time     Interval  Multiplier
127.0.0.1        Up        ae0.0                  0.900     0.300         3
127.0.0.1        Up        ae0.0                  0.900     0.300         3
172.30.0.2       Up        ae0.0                  0.450     0.150         3
172.30.0.6       Up        ge-0/0/4.114           0.450     0.150         3
172.30.0.10      Up        ge-0/0/4.118           0.450     0.150         3
172.30.0.66      Up        ge-0/0/4.206           0.450     0.150         3
172.30.5.6       Up                               0.900     0.300         3
172.30.5.8       Up                               0.900     0.300         3
172.30.5.41      Up                               0.450     0.150         3
...
```

You can see that **LSP-ping** is also UP. This shows that you can use LSP ping to test the connectivity via the LSP.

```
[edit]
lab@R1# run ping mpls rsvp R1-to-R8
!!!!!
--- lsping statistics ---
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
5 packets transmitted, 5 packets received, 0% packet loss
```

b) R2.

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R5 to 172.30.5.5

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R5 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 to 172.30.5.7

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R2# commit
commit complete
```

Verify the RSVP session for the configured LSPs.

```
[edit]
lab@R2# run show rsvp session
Ingress RSVP: 2 sessions
To              From            State Rt Style Labelin Labelout LSPname
172.30.5.5      172.30.5.2      Up    0  1 FF       -    299824 R2-to-R5
172.30.5.7      172.30.5.2      Up    0  1 FF       -         3 R2-to-R7
Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 1 sessions
To              From            State Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.1      Up    0  1 FF  299824   299824 R1-to-R6
Total 1 displayed, Up 1, Down 0
```

Both sessions are UP and signaled successfully. In this sample output for R2 you can see that there is one Transit session as well. The name of the session is carried with the RSVP object, so you see that this is the LSP from R1 to R6.

Verify the MPLS LSPs.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# run show mpls lsp
Ingress LSP: 2 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.5      172.30.5.2      Up    0 *                       R2-to-R5
172.30.5.7      172.30.5.2      Up    0 *                       R2-to-R7
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 1 sessions
To              From            State Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.1      Up    0  1 FF  299824   299824 R1-to-R6
Total 1 displayed, Up 1, Down 0
```

```
[edit]
lab@R2# run show route table inet.3

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32      *[LDP/9] 03:44:43, metric 5
                    > to 172.30.0.1 via ae0.0, Push 0
172.30.5.3/32      *[LDP/9] 03:48:54, metric 10
                    > to 172.30.0.14 via ge-0/0/4.123, Push 0
172.30.5.4/32      *[LDP/9] 03:44:16, metric 15
                    > to 172.30.0.1 via ae0.0, Push 299808
172.30.5.5/32      *[RSVP/7/1] 00:03:21, metric 25
                    > to 172.30.0.1 via ae0.0, label-switched-path R2-to-R5
172.30.5.7/32      *[RSVP/7/1] 00:03:21, metric 10
                    > to 172.30.0.18 via ge-0/0/4.127, label-switched-path R2-to-
R7
```

c) R3.

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 to 172.30.5.6

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 to 172.30.5.8

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-2 to 172.30.5.8

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-2 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R3# commit
commit complete
```

R3 has two LSPs to R8. Listing the RSVP session and MPLS LSP table.

```
[edit]
lab@R3# run show rsvp session
Ingress RSVP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.3      Up       0  1 FF       -        3 R3-to-R6
172.30.5.8      172.30.5.3      Up       0  1 FF       -   299840 R3-to-R8-1
172.30.5.8      172.30.5.3      Up       0  1 FF       -   299856 R3-to-R8-2
Total 3 displayed, Up 3, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
[edit]
lab@R3# run show mpls lsp
Ingress LSP: 3 sessions
To              From            State Rt P  ActivePath   LSPname
172.30.5.6      172.30.5.3      Up     0 *               R3-to-R6
172.30.5.8      172.30.5.3      Up     0 *               R3-to-R8-1
172.30.5.8      172.30.5.3      Up     0 *               R3-to-R8-2
Total 3 displayed, Up 3, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

The `inet.3` routing table proves that the route for R8 loopback IP address is reachable via two MPLS LSPs.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show route table inet.3 protocol rsvp

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.6/32   *[RSVP/7/1] 00:05:28, metric 10
                 > to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-R6
172.30.5.8/32   *[RSVP/7/1] 00:05:28, metric 25
                 > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-R8-
1
                   to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-R8-
2
```

d) R4

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 to 172.30.5.7

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-2 to 172.30.5.7


[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-2 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R5 to 172.30.5.5

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R5 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R4# commit
commit complete
```

R4 has two LSPs to R7 loopback IP address.

```
[edit]
lab@R4# run show rsvp session
Ingress RSVP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.5    172.30.5.4  Up       0  1 FF       -         3 R4-to-R5
172.30.5.7    172.30.5.4  Up       0  1 FF       -    299872 R4-to-R7-1
172.30.5.7    172.30.5.4  Up       0  1 FF       -    299888 R4-to-R7-2
Total 3 displayed, Up 3, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
[edit]
lab@R4# run show mpls lsp
Ingress LSP: 3 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.5      172.30.5.4      Up     0 *                      R4-to-R5
172.30.5.7      172.30.5.4      Up     0 *                      R4-to-R7-1
172.30.5.7      172.30.5.4      Up     0 *                      R4-to-R7-2
Total 3 displayed, Up 3, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
[edit]
lab@R4# run show route table inet.3 protocol rsvp

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.5/32      *[RSVP/7/1] 00:03:44, metric 10
                    > to 172.30.0.30 via ge-0/0/4.145, label-switched-path R4-to-
R5
172.30.5.7/32      *[RSVP/7/1] 00:03:44, metric 25
                    > to 172.30.0.5 via ge-0/0/4.114, label-switched-path R4-to-
R7-1
                      to 172.30.0.30 via ge-0/0/4.145, label-switched-path R4-to-
R7-2
```

e) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 to 172.30.5.4

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set protocols mpls label-switched-path R5-to-R4 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R2 to 172.30.5.2

[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R2 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R5# commit
commit complete
```

Here you can see that there are Ingress and Egress LSPs. LSPs are unidirectional, this means that an LSP should be configured in both directions for bidirectional communication.

```
[edit]
lab@R5# run show rsvp session
Ingress RSVP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.2      172.30.5.5      Up       0  1 FF       -    299840 R5-to-R2
172.30.5.4      172.30.5.5      Up       0  1 FF       -         3 R5-to-R4
Total 2 displayed, Up 2, Down 0

Egress RSVP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.5      172.30.5.2      Up       0  1 FF       3         - R2-to-R5
172.30.5.5      172.30.5.4      Up       0  1 FF       3         - R4-to-R5
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.4      Up       0  1 FF  299888   299856 R4-to-R7-2
Total 1 displayed, Up 1, Down 0
```

R5 is advertising an Implicit Null label 3.

```
[edit]
lab@R5# run show mpls lsp
Ingress LSP: 2 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.2      172.30.5.5      Up     0 *                      R5-to-R2
172.30.5.4      172.30.5.5      Up     0 *                      R5-to-R4
Total 2 displayed, Up 2, Down 0

Egress LSP: 2 sessions
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.5      172.30.5.2      Up       0  1 FF       3        - R2-to-R5
172.30.5.5      172.30.5.4      Up       0  1 FF       3        - R4-to-R5
Total 2 displayed, Up 2, Down 0

Transit LSP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.4      Up       0  1 FF  299888   299856 R4-to-R7-2
Total 1 displayed, Up 1, Down 0
```

f) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 to 172.30.5.1

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 to 172.30.5.3

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R6# commit
commit complete
```

Verify the RSVP sessions and MPLS LSP table.

```
[edit]
lab@R6# run show rsvp session
Ingress RSVP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.6      Up       0  1 FF       -   299904 R6-to-R1
172.30.5.3      172.30.5.6      Up       0  1 FF       -        3 R6-to-R3
Total 2 displayed, Up 2, Down 0

Egress RSVP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.3      Up       0  1 FF       3        - R3-to-R6
172.30.5.6      172.30.5.1      Up       0  1 FF       3        - R1-to-R6
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.7    172.30.5.4  Up       0  1 FF  299856        3 R4-to-R7-2
Total 1 displayed, Up 1, Down 0
```

```
[edit]
lab@R6# run show mpls lsp
Ingress LSP: 2 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.1      172.30.5.6      Up     0 *                       R6-to-R1
172.30.5.3      172.30.5.6      Up     0 *                       R6-to-R3
Total 2 displayed, Up 2, Down 0

Egress LSP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.3      Up       0  1 FF       3        - R3-to-R6
172.30.5.6      172.30.5.1      Up       0  1 FF       3        - R1-to-R6
Total 2 displayed, Up 2, Down 0

Transit LSP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.4      Up       0  1 FF  299856        3 R4-to-R7-2
Total 1 displayed, Up 1, Down 0
```

Although there are Transit and Egress LSPs, only the Ingress LSP routes are present in the routing table `inet.3`. As it was stated only the Ingress router uses the `inet.3` table, other routers on the path are using the `mpls.0` table.

```
[edit]
lab@R6# run show route table inet.3

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32      *[RSVP/7/1] 00:02:07, metric 25
                    > to 172.30.0.33 via ae0.0, label-switched-path R6-to-R1
172.30.5.3/32      *[RSVP/7/1] 00:02:07, metric 10
                    > to 172.30.0.25 via ge-0/0/4.136, label-switched-path R6-to-
R3
172.30.5.5/32      *[LDP/9] 04:37:47, metric 5
                    > to 172.30.0.33 via ae0.0, Push 0
172.30.5.7/32      *[LDP/9] 04:33:07, metric 10
                    > to 172.30.0.42 via ge-0/0/4.167, Push 0
172.30.5.8/32      *[LDP/9] 04:32:22, metric 15
                    > to 172.30.0.33 via ae0.0, Push 299840
```

g) R7

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 to 172.30.5.2

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 to 172.30.5.4

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-2 to 172.30.5.4

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-2 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R7# commit
commit complete
```

R7 is configured with three LSPs, two of them are towards R4. They are needed to ensure a bidirectional path between both routers exists.

```
[edit]
lab@R7# run show rsvp session
Ingress RSVP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.2      172.30.5.7      Up       0  1 FF       -         3 R7-to-R2
172.30.5.4      172.30.5.7      Up       0  1 FF       -    299872 R7-to-R4-1
172.30.5.4      172.30.5.7      Up       0  1 FF       -    299888 R7-to-R4-2
Total 3 displayed, Up 3, Down 0

Egress RSVP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.2      Up       0  1 FF       3         - R2-to-R7
172.30.5.7      172.30.5.4      Up       0  1 FF       3         - R4-to-R7-1
172.30.5.7      172.30.5.4      Up       0  1 FF       3         - R4-to-R7-2
Total 3 displayed, Up 3, Down 0

Transit RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.1      Up       0  1 FF  299824         3 R1-to-R6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Total 1 displayed, Up 1, Down 0
```

```
[edit]
lab@R7# run show mpls lsp
Ingress LSP: 3 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.2      172.30.5.7      Up     0 *                      R7-to-R2
172.30.5.4      172.30.5.7      Up     0 *                      R7-to-R4-1
172.30.5.4      172.30.5.7      Up     0 *                      R7-to-R4-2
Total 3 displayed, Up 3, Down 0

Egress LSP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.2      Up       0  1 FF       3         - R2-to-R7
172.30.5.7      172.30.5.4      Up       0  1 FF       3         - R4-to-R7-1
172.30.5.7      172.30.5.4      Up       0  1 FF       3         - R4-to-R7-2
Total 3 displayed, Up 3, Down 0

Transit LSP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.1      Up       0  1 FF  299824         3 R1-to-R6
Total 1 displayed, Up 1, Down 0
```

You can see that the LSP to R2 is configured for bidirectional connectivity.

```
[edit]
lab@R7# run show route table inet.3 protocol rsvp

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.2/32      *[RSVP/7/1] 00:01:49, metric 10
                    > to 172.30.0.17 via ge-0/0/4.127, label-switched-path R7-to-
R2
172.30.5.4/32      *[RSVP/7/1] 00:01:49, metric 25
                       to 172.30.0.41 via ge-0/0/4.167, label-switched-path R7-to-
R4-1
                    > to 172.30.0.41 via ge-0/0/4.167, label-switched-path R7-to-
R4-2
```

h) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-1 to 172.30.5.3

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-1 oam bfd-liveness-detection
minimum-interval 3000
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-2 to 172.30.5.3

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-2 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 to 172.30.5.1

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 oam bfd-liveness-detection
minimum-interval 3000

[edit]
lab@R8# commit
commit complete
```

Verify the RSVP session table and MPLS LSP table.

```
[edit]
lab@R8# run show rsvp session
Ingress RSVP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.8      Up       0  1 FF       -         3 R8-to-R1
172.30.5.3      172.30.5.8      Up       0  1 FF       -    299952 R8-to-R3-1
172.30.5.3      172.30.5.8      Up       0  1 FF       -    299904 R8-to-R3-2
Total 3 displayed, Up 3, Down 0

Egress RSVP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.8      172.30.5.3      Up       0  1 FF       3         - R3-to-R8-1
172.30.5.8      172.30.5.3      Up       0  1 FF       3         - R3-to-R8-2
172.30.5.8      172.30.5.1      Up       0  1 FF       3         - R1-to-R8
Total 3 displayed, Up 3, Down 0

Transit RSVP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.6      Up       0  1 FF  299856         3 R6-to-R1
172.30.5.2      172.30.5.5      Up       0  1 FF  299840    299888 R5-to-R2
172.30.5.5      172.30.5.2      Up       0  1 FF  299824         3 R2-to-R5
Total 3 displayed, Up 3, Down 0
```

```
lab@R8# run show mpls lsp
Ingress LSP: 3 sessions
To              From            State Rt P     ActivePath       LSPname
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.1  172.30.5.8   Up     0 *                    R8-to-R1
172.30.5.3  172.30.5.8   Up     0 *                    R8-to-R3-1
172.30.5.3  172.30.5.8   Up     0 *                    R8-to-R3-2
Total 3 displayed, Up 3, Down 0

Egress LSP: 3 sessions
To          From         State   Rt Style Labelin Labelout LSPname
172.30.5.8  172.30.5.3   Up       0  1 FF       3         - R3-to-R8-1
172.30.5.8  172.30.5.3   Up       0  1 FF       3         - R3-to-R8-2
172.30.5.8  172.30.5.1   Up       0  1 FF       3         - R1-to-R8
Total 3 displayed, Up 3, Down 0

Transit LSP: 3 sessions
To          From         State   Rt Style Labelin Labelout LSPname
172.30.5.1  172.30.5.6   Up       0  1 FF  299856        3 R6-to-R1
172.30.5.2  172.30.5.5   Up       0  1 FF  299840   299888 R5-to-R2
172.30.5.5  172.30.5.2   Up       0  1 FF  299824        3 R2-to-R5
Total 3 displayed, Up 3, Down 0
```

Confirm all LSPs are configured to ensure bidirectional paths to R1 and R3.

```
[edit]
lab@R8# run show route table inet.3

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32      *[RSVP/7/1] 00:00:56, metric 10
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R1
172.30.5.3/32      *[RSVP/7/1] 00:00:56, metric 25
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-1
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-2
```

4) Configure the LSPs to use links which have the required administrative groups.

The next step is to instruct CSPF to use links with the link colors specified in the task requirements.

a) R1

On R1 the LSP to R8 should include links which have the green color assigned.

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 admin-group include-any green
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

To configure a given LSP to be restricted based on link colors you can use the following three options:

- *include-any* – performs logical OR to the configured colors. Any links matching one or many of the configured colors are included in CSPF calculations.
- *include-all* – performs logical AND. Only links matching all configured colors are taken in account for calculations.
- *exclude* – links with colors specified with exclude are automatically excluded from the calculations. If more than one color is specified logical OR is performed.

Verify the LSP parameters.

```
[edit]
lab@R1# run show mpls lsp name R1-to-R8 extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R8
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
         Include Any: green
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.18 S 172.30.0.46 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
         172.30.0.2 172.30.0.18 172.30.0.46
    OAM state : BFD session up LSP-ping up
   12 Sep 21 12:09:04.278 Selected as active path
   11 Sep 21 12:09:04.278 Record Route:  172.30.0.2 172.30.0.18 172.30.0.46
   10 Sep 21 12:09:04.278 Up
    9 Sep 21 12:09:04.206 Originate Call
    8 Sep 21 12:09:04.206 CSPF: computation result accepted  172.30.0.2
172.30.0.18 172.30.0.46
    7 Sep 21 12:09:04.198 Clear Call
    6 Sep 21 12:09:04.193 Deselected as active
    5 Sep 21 12:04:32.267 Selected as active path
    4 Sep 21 12:04:32.266 Record Route:  172.30.0.2 172.30.0.18 172.30.0.46
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    3 Sep 21 12:04:32.264 Up
    2 Sep 21 12:04:32.176 Originate Call
    1 Sep 21 12:04:32.176 CSPF: computation result accepted  172.30.0.2
172.30.0.18 172.30.0.46
  Created: Sat Sep 21 12:04:31 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 3 sessions
Total 0 displayed, Up 0, Down 0
```

b) R2

On R2 LSP to R7 should pass only over links with color red.

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 admin-group include-any red

[edit]
lab@R2# commit
commit complete
```

Verify the LSP parameters.

```
[edit]
lab@R2# run show mpls lsp name R2-to-R7 extensive
Ingress LSP: 2 sessions

172.30.5.7
  From: 172.30.5.2, State: Up, ActiveRoute: 0, LSPname: R2-to-R7
  ActivePath: (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Revert timer: 0
 *Primary   primary-2        State: Up, No-decrement-ttl
    Priorities: 7 0
    SmartOptimizeTimer: 180
          Include Any: red
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.1 S 172.30.0.10 S 172.30.0.45 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.1 172.30.0.10 172.30.0.45
    OAM state : BFD session up  LSP-ping up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    5 Sep 22 13:18:32.210 Selected as active path
    4 Sep 22 13:18:32.209 Record Route:  172.30.0.1 172.30.0.10 172.30.0.45
    3 Sep 22 13:18:32.209 Up
    2 Sep 22 13:18:32.124 Originate Call
    1 Sep 22 13:18:32.124 CSPF: computation result accepted  172.30.0.1
172.30.0.10 172.30.0.45

Total 1 displayed, Up 1, Down 0

Egress LSP: 3 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 9 sessions
Total 0 displayed, Up 0, Down 0
```

c) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 admin-group include-any red


[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R5 admin-group include-any green

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 admin-group include-any green

[edit]
lab@R5# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

f) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 admin-group include-any red

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 admin-group include-any green

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 admin-group include-any green

[edit]
lab@R8# commit
commit complete
```

5) Configure LSPs I, K, J and L with static configured EROs.

Four of the LSPs have to be configured with static predefined paths through the network. Static EROs are configured under `protocols mpls path`. You can specify strict or loose hops that the LSPs should transit. When a loose hop is configured, it identifies one or more transit LSRs through which the LSP must be signaled. The IGP is used for the exact route from the inbound router to the first loose hop, or from one loose hop to the next. The loose hop specifies only that a particular LSR hop must be included in the LSP. Whereas when a strict hop is configured, it identifies an exact path through which the LSP must be routed. You can have at the same time loose and strict hops in a single ERO. You can specify parameters within path like reserved bandwidth and priority that affect only that path. If not specified they are inherited from the `label-switched-path` configuration.

R3

```
[edit]
lab@R3# set protocols mpls path path-1 172.30.5.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols mpls path path-1 172.30.5.1

[edit]
lab@R3# set protocols mpls path path-1 172.30.5.8

[edit]
lab@R3# set protocols mpls path path-2 172.30.5.6

[edit]
lab@R3# set protocols mpls path path-2 172.30.5.7

[edit]
lab@R3# set protocols mpls path path-2 172.30.5.8


[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 primary path-1

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-2 primary path-2

[edit]
lab@R3# commit
commit complete
```

Define two different paths, one is applied to the first LSP to R8, the other to the second LSP to R8. This way both LSPs will have different paths statically configured.

Note: Only a single primary path can be configured for a given LSP.

Verify the path of both LSPs.

```
[edit]
lab@R3# run show mpls lsp name R3-to-R8-1 extensive
Ingress LSP: 3 sessions

172.30.5.8
  From: 172.30.5.3, State: Up, ActiveRoute: 0, LSPname: R3-to-R8-1
  ActivePath: path-1 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-1              State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 172.30.0.13 S 172.30.0.1 S 172.30.0.10 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.13 172.30.0.1 172.30.0.10
    OAM state : BFD session up LSP-ping up
    5 Sep 21 14:07:11.685 Selected as active path
    4 Sep 21 14:07:11.684 Record Route:  172.30.0.13 172.30.0.1 172.30.0.10
    3 Sep 21 14:07:11.684 Up
    2 Sep 21 14:07:11.594 Originate Call
    1 Sep 21 14:07:11.594 CSPF: computation result accepted  172.30.0.13
172.30.0.1 172.30.0.10
```

```
[edit]
lab@R3# run show mpls lsp name R3-to-R8-2 extensive
Ingress LSP: 3 sessions

172.30.5.8
  From: 172.30.5.3, State: Up, ActiveRoute: 0, LSPname: R3-to-R8-2
  ActivePath: path-2 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-2              State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 172.30.0.26 S 172.30.0.42 S 172.30.0.46 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.26 172.30.0.42 172.30.0.46
    OAM state : BFD session not up LSP-ping not up
    5 Sep 21 14:07:11.784 Selected as active path
    4 Sep 21 14:07:11.783 Record Route:  172.30.0.26 172.30.0.42 172.30.0.46
    3 Sep 21 14:07:11.782 Up
    2 Sep 21 14:07:11.608 Originate Call
    1 Sep 21 14:07:11.608 CSPF: computation result accepted  172.30.0.26
172.30.0.42 172.30.0.46
```

You can see that after the commit both LSP are using the configured paths and they reach R8 loopback IP address via different routes defined.

On R8 you have to configure the LSP.

R8

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# set protocols mpls path path-1 172.30.5.1

[edit]
lab@R8# set protocols mpls path path-1 172.30.5.2

[edit]
lab@R8# set protocols mpls path path-1 172.30.5.3

[edit]
lab@R8# set protocols mpls path path-2 172.30.5.5

[edit]
lab@R8# set protocols mpls path path-2 172.30.5.4

[edit]
lab@R8# set protocols mpls path path-2 172.30.5.3

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-1 primary path-1

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-2 primary path-2

[edit]
lab@R8# commit
commit complete
```

Verify the route of the LSPs.

```
[edit]
lab@R8# run show mpls lsp name R8-to-R3-1 extensive
Ingress LSP: 3 sessions

172.30.5.3
  From: 172.30.5.8, State: Up, ActiveRoute: 0, LSPname: R8-to-R3-1
  ActivePath: path-1 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-1              State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.9 S 172.30.0.2 S 172.30.0.14 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.9 172.30.0.2 172.30.0.14
    OAM state : BFD session up LSP-ping up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    5 Sep 21 14:10:31.286 Selected as active path
    4 Sep 21 14:10:31.286 Record Route:  172.30.0.9 172.30.0.2 172.30.0.14
    3 Sep 21 14:10:31.286 Up
    2 Sep 21 14:10:31.204 Originate Call
    1 Sep 21 14:10:31.204 CSPF: computation result accepted  172.30.0.9
172.30.0.2 172.30.0.14
  Created: Sat Sep 21 14:10:30 2013
Total 1 displayed, Up 1, Down 0
```

```
[edit]
lab@R8# run show mpls lsp name R8-to-R3-2 extensive
Ingress LSP: 3 sessions

172.30.5.3
  From: 172.30.5.8, State: Up, ActiveRoute: 0, LSPname: R8-to-R3-2
  ActivePath: path-2 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-2           State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 172.30.0.37 S 172.30.0.29 S 172.30.0.21 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.37 172.30.0.29 172.30.0.21
    OAM state : BFD session up  LSP-ping up
    5 Sep 21 14:10:31.389 Selected as active path
    4 Sep 21 14:10:31.388 Record Route:  172.30.0.37 172.30.0.29 172.30.0.21
    3 Sep 21 14:10:31.388 Up
    2 Sep 21 14:10:31.207 Originate Call
    1 Sep 21 14:10:31.207 CSPF: computation result accepted  172.30.0.37
172.30.0.29 172.30.0.21
  Created: Sat Sep 21 14:10:30 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 3 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions
Total 0 displayed, Up 0, Down 0
```

6) Configure LSPs M, O, N and P with static configured EROs.

The four LSPs should be configured with static EROs as well. The difference from the four in the previous step is the administrative group configuration, which is transferred to the

© 2020 Juniper Networks, Inc. All rights reserved.

static paths. The CSPF algorithm will prune all links that do not match the administrative group and will calculate the shortest path to the destination while taking the EROs into account.

R4

```
[edit]
lab@R4# set protocols mpls path path-1 172.30.5.3

[edit]
lab@R4# set protocols mpls path path-2 172.30.5.5

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 admin-group include-any red


[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 primary path-1

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-2 admin-group include-any red

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-2 primary path-2

[edit]
lab@R4# commit
commit complete
```

Verify the LSPs paths.

```
[edit]
lab@R4# run show mpls lsp name R4-to-R7-1 extensive
Ingress LSP: 3 sessions

172.30.5.7
  From: 172.30.5.4, State: Up, ActiveRoute: 0, LSPname: R4-to-R7-1
  ActivePath: path-1 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-1           State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
          Include Any: red
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 45)
 172.30.0.21 S 172.30.0.13 S 172.30.0.1 S 172.30.0.10 S 172.30.0.45 S
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.21 172.30.0.13 172.30.0.1 172.30.0.10 172.30.0.45
    OAM state : BFD session up LSP-ping up
    5 Sep 21 14:27:39.539 Selected as active path
    4 Sep 21 14:27:39.539 Record Route:  172.30.0.21 172.30.0.13 172.30.0.1
172.30.0.10 172.30.0.45
    3 Sep 21 14:27:39.537 Up
    2 Sep 21 14:27:39.347 Originate Call
    1 Sep 21 14:27:39.347 CSPF: computation result accepted  172.30.0.21
172.30.0.13 172.30.0.1 172.30.0.10 172.30.0.45
  Created: Sat Sep 21 14:27:39 2013
Total 1 displayed, Up 1, Down 0
```

```
[edit]
lab@R4# run show mpls lsp name R4-to-R7-2 extensive
Ingress LSP: 3 sessions

172.30.5.7
  From: 172.30.5.4, State: Up, ActiveRoute: 0, LSPname: R4-to-R7-2
  ActivePath: path-2 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-2              State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
          Include Any: red
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.30 S 172.30.0.34 S 172.30.0.42 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.30 172.30.0.34 172.30.0.42
    OAM state : BFD session up LSP-ping up
    5 Sep 21 14:27:39.555 Selected as active path
    4 Sep 21 14:27:39.554 Record Route:  172.30.0.30 172.30.0.34 172.30.0.42
    3 Sep 21 14:27:39.552 Up
    2 Sep 21 14:27:39.353 Originate Call
    1 Sep 21 14:27:39.353 CSPF: computation result accepted  172.30.0.30
172.30.0.34 172.30.0.42
  Created: Sat Sep 21 14:27:39 2013
Total 1 displayed, Up 1, Down 0
```

On R7 both LSPs to R5 should be configured to use links with color green and with static EROs.

R7

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols mpls path path-1 172.30.5.2

[edit]
lab@R7# set protocols mpls path path-2 172.30.5.8

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 admin-group include-any green

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 primary path-1

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-2 admin-group include-any green

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-2 primary path-2

[edit]
lab@R7# commit
commit complete
```

Verify the changed path of both LSPs.

```
[edit]
lab@R7# run show mpls lsp name R7-to-R4-1 extensive
Ingress LSP: 3 sessions

172.30.5.4
  From: 172.30.5.7, State: Up, ActiveRoute: 0, LSPname: R7-to-R4-1
  ActivePath: path-1 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-1                State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
          Include Any: green
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.17 S 172.30.0.1 S 172.30.0.6 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.17 172.30.0.1 172.30.0.6
    OAM state : BFD session up  LSP-ping up
    5 Sep 21 14:41:04.558 Selected as active path
    4 Sep 21 14:41:04.558 Record Route:  172.30.0.17 172.30.0.1 172.30.0.6
    3 Sep 21 14:41:04.557 Up
    2 Sep 21 14:41:04.470 Originate Call
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show mpls lsp name R7-to-R4-2 extensive
Ingress LSP: 3 sessions

172.30.5.4
  From: 172.30.5.7, State: Up, ActiveRoute: 0, LSPname: R7-to-R4-2
  ActivePath: path-2 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-2              State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
         Include Any: green
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 45)
 172.30.0.46 S 172.30.0.37 S 172.30.0.34 S 172.30.0.25 S 172.30.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
         172.30.0.46 172.30.0.37 172.30.0.34 172.30.0.25 172.30.0.22
    OAM state : BFD session up  LSP-ping up
    5 Sep 21 14:41:04.604 Selected as active path
    4 Sep 21 14:41:04.604 Record Route:  172.30.0.46 172.30.0.37 172.30.0.34
172.30.0.25 172.30.0.22
    3 Sep 21 14:41:04.604 Up
    2 Sep 21 14:41:04.472 Originate Call
    1 Sep 21 14:41:04.472 CSPF: computation result accepted  172.30.0.46
172.30.0.37 172.30.0.34 172.30.0.25 172.30.0.22
```

7) Next all LSPs except LSP A, B, S and T should be configured with bandwidth reservation.

LSPs can be configured with a reserved bandwidth along the path. RSVP signals the reservation. During the setup process each downstream router will receive objects for the reservation. Then each router determines if enough bandwidth is available on the egress interface. Please note that bandwidth reservations are based on a logical value, the actual traffic passing through interface is not considered. Also keep in mind that it is possible to reserve bandwidth without using a TED.

Configure the required LSPs with bandwidth reservation of 60Mbps.

a) R1

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 bandwidth 60m

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R5 bandwidth 60m

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 bandwidth 60m

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 bandwidth 60m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 bandwidth 60m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-2 bandwidth 60m

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 bandwidth 60m

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-2 bandwidth 60m

[edit]
lab@R4# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

e) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R2 bandwidth 60m

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 bandwidth 60m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 bandwidth 60m

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 bandwidth 60m

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 bandwidth 60m

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-2 bandwidth 60m

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-1 bandwidth 60m

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-2 bandwidth 60m

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# commit
commit complete
```

Verify the LSP after the commit.

```
[edit]
lab@R1# run show mpls lsp name R1-to-R6 extensive
Ingress LSP: 2 sessions

172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    Bandwidth: 60Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.10 S 172.30.0.37 S 172.30.0.34 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.10 172.30.0.37 172.30.0.34
    OAM state : BFD session up  LSP-ping up
   13 Sep 21 14:56:32.369 Record Route:  172.30.0.10 172.30.0.37 172.30.0.34
   12 Sep 21 14:56:32.368 Up
   11 Sep 21 14:56:32.285 Originate Call
   10 Sep 21 14:56:32.284 Clear Call
    9 Sep 21 14:56:32.284 CSPF: computation result accepted  172.30.0.10
172.30.0.37 172.30.0.34
    8 Sep 21 09:53:31.177 LSP ping failure on path[15 times]
    7 Sep 21 08:26:27.933 Selected as active path
```

The LSP is signaled with a 60Mbps bandwidth reservation. All other LSPs on the other routers should be configured the same way.

8) Configure LSPs A, B, S and T auto bandwidth.

Regardless of the static reservation, Junos provides a capability for automatic bandwidth allocation. The routers can resignal the LSPs in case a bandwidth change is required. MPLS statistics are collected and the utilization is monitored over a predefined time. At the end of this time period the LSP is resignaled using the make-before-break feature. This ensures that traffic is not dropped during resignaling of the LSP.

Configure autobandwidth on the required LSPs.

© 2020 Juniper Networks, Inc. All rights reserved.

a) R1

```
[edit]
lab@R1# set protocols mpls statistics file Auto-Bandwidth

[edit]
lab@R1# set protocols mpls statistics auto-bandwidth

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 auto-bandwidth adjust-interval
172800
```

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 auto-bandwidth minimum-
bandwidth 30m

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 auto-bandwidth maximum-
bandwidth 120m

[edit]
lab@R1# commit
commit complete
```

MPLS statistics must be configured for the automatic bandwidth feature to work. The Adjust-interval is the period that the ingress router monitors the utilization. By default, it is 5 minutes. The minimum-bandwidth and maximum-bandwidth is the range in which the bandwidth allocation for the LSP must stay.

b) R4

```
[edit]
lab@R4# set protocols mpls statistics file Auto-Bandwidth

[edit]
lab@R4# set protocols mpls statistics auto-bandwidth

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R5 auto-bandwidth adjust-interval
172800

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R5 auto-bandwidth minimum-
bandwidth 30m

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# set protocols mpls label-switched-path R4-to-R5 auto-bandwidth maximum-
bandwidth 120m

[edit]
lab@R4# commit
commit complete
```

c) R5

```
[edit]
lab@R5# set protocols mpls statistics file Auto-Bandwidth

[edit]
lab@R5# set protocols mpls statistics auto-bandwidth

[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 auto-bandwidth adjust-interval
172800

[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 auto-bandwidth minimum-
bandwidth 30m

[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 auto-bandwidth maximum-
bandwidth 120m

[edit]
lab@R5# commit
commit complete
```

d) R8

```
[edit]
lab@R8# set protocols mpls statistics file Auto-Bandwidth

[edit]
lab@R8# set protocols mpls statistics auto-bandwidth

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 auto-bandwidth adjust-interval
172800

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 auto-bandwidth minimum-
bandwidth 30m
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 auto-bandwidth maximum-
bandwidth 120m

[edit]
lab@R8# commit
commit complete
```

Verify the result of the autobandwidth LSP configuration.

```
[edit]
lab@R1# run show mpls lsp name R1-to-R8 extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R8
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 30Mbps MaxBW: 120Mbps
  AdjustTimer: 172800 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 172764 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    Bandwidth: 30Mbps
    SmartOptimizeTimer: 180
          Include Any: green
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.18 S 172.30.0.46 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
```

In the output is seen that the LSP is configured for autobandwidth optimization. The current signaled bandwidth is 30Mbps, which is the minimum as no real traffic is passing.

The LSPs on the other routers should be configured in the same way.

9) Configure LSPs A, B, E, F, I, J, Q, R, S, and T with higher priorities.

RSVP signaled LSPs support setup and hold priorities. These priorities are used by the routers on the LSP path to evaluate the setup of new LSPs against the hold priority of

© 2020 Juniper Networks, Inc. All rights reserved.

existing LSPs. When there are not enough resources available, LSPs with a stronger setup priority cause LSPs with weaker hold priority to tear down.

LSP setup and hold priorities are in range from 0 to 7, where 0 means strongest and 7 is weakest. By default, preemption is disabled. All LSP are signaled with a setup priority of 7 and hold priority of 0. Junos will not commit if the LSP hold priority is weaker than the setup priority.

Configure higher priorities.

a) R1

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 priority 6 6

[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 priority 6 6

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 priority 6 6

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 priority 6 6

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# set protocols mpls label-switched-path R4-to-R5 priority 6 6

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 priority 6 6

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 priority 6 6

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 priority 6 6

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 priority 6 6

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-1 priority 6 6

[edit]
lab@R8# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

Note: First digit in the command is the setup priority, second is the hold priority.

After committing the configuration, the LSP to R8 have a higher setup priority.

```
[edit]
lab@R1# run show mpls lsp name R1-to-R8 extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R8
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 30Mbps MaxBW: 120Mbps
  AdjustTimer: 172800 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 172022 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 6 6
    Bandwidth: 30Mbps
    SmartOptimizeTimer: 180
          Include Any: green
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.18 S 172.30.0.46 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
```

All other LSPs in the list should be configured in the same way with higher priorities.

10) Configure all other LSP with lower priorities.

The rest of the LSPs have to be configured with lower priority, which means that they could be preempted by the LSPs configured in the previous step. In fact, only the hold priority will be changed from 0 to 7 as by default the setup priority is 7, which is the lowest possible value.

a) R1

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 priority 7 7
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R5 priority 7 7

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-2 priority 7 7

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 priority 7 7

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-2 priority 7 7

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R2 priority 7 7

[edit]
lab@R5# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

f) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 priority 7 7

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 priority 7 7

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-2 priority 7 7

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-2 priority 7 7

[edit]
lab@R8# commit
commit complete
```

```
[edit]
lab@R1# run show mpls lsp name R1-to-R6 extensive
Ingress LSP: 2 sessions

172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath:  (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 7
    Bandwidth: 60Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.18 S 172.30.0.41 S
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.2 172.30.0.18 172.30.0.41
    OAM state : BFD session up  LSP-ping up
    5 Sep 21 15:22:59.366 Selected as active path
    4 Sep 21 15:22:59.365 Record Route:  172.30.0.2 172.30.0.18 172.30.0.41
    3 Sep 21 15:22:59.362 Up
    2 Sep 21 15:22:59.283 Originate Call
```

The hold priority of the LSP to R6 is changed to 7.

11) Configure soft-preemption for LSPs K, L, O, and P.

By default, an LSP which is preempted is torn down before a new path is calculated. For example, if an LSP is preempted, the ingress LSR tries to find new path to the egress LSR. During that time traffic for the preempted LSP is dropped. To avoid this situation Junos is able to specify `soft-preempt` to the RSVP route object. When this feature is configured the ingress LSR tries to establish a new LSP before it tears down the preempted LSP.

Configure soft-preemption for the requested LSPs.

a) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-2 soft-preemption

[edit]
lab@R3# commit
commit complete
```

b) R4

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-2 soft-preemption

[edit]
lab@R4# commit
commit complete
```

c) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-2 soft-preemption
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# commit
commit complete
```

d) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-2 soft-preemption

[edit]
lab@R8# commit
commit complete
```

Verify if soft-preemption is signaled.

```
 [edit]
lab@R3# run show rsvp session name R3-to-R8-2 extensive
Ingress RSVP: 8 sessions

172.30.5.8
  From: 172.30.5.3, LSPstate: Up, ActiveRoute: 0
  LSPname: R3-to-R8-2, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 332208
  Resv style: 1 SE, Label in: -, Label out: 332208
  Time left:    -, Since: Wed Oct 23 07:37:55 2013
  Tspec: rate 60Mbps size 60Mbps peak Infbps m 20 M 1500
  Port number: sender 12 receiver 7255 protocol 0
  Soft preemption desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 172.30.0.26 (ge-0/0/4.136) 14 pkts
  RESV rcvfrom: 172.30.0.26 (ge-0/0/4.136) 9 pkts
  Explct route: 172.30.0.26 172.30.0.42 172.30.0.46
  Record route: <self> 172.30.0.26 172.30.0.42 172.30.0.46
Total 1 displayed, Up 1, Down 0
```

12) Configure LSPs I, J, K, L, M, N, O and P for automatic optimization.

By default, when an LSP is established only an IGP change or manual clearing of the LSP can force a recomputation of the path. If you enable automatic optimization, the LSP path can change as result of periodic CSPF recomputations. By default, the `optimize-timer` is set to 0, which in fact disables the optimization.

© 2020 Juniper Networks, Inc. All rights reserved.

Another requirement for the LSPs is to enable the make-before-brake capability to ensure that an optimized LSP is created before the old LSP is torn down. This is accomplished with the knob `adaptive`. The `adaptive` feature changes the LSP reservation style. By default, Junos is using the FF (fixed filter) reservation style, which means that if one LSP has two established paths through the network and if both cross the same link, the LSP bandwidth is reserved twice over that link. This might cause an issue if there is not enough logical bandwidth left for the LSP to be established. The `adaptive` feature changes the reservation to SE (shared explicit). This way the bandwidth will be shared and only accounted once on the common links.

Note: If `adaptive` is applied to primary or secondary LSP paths, the reservation style is only changed for that specific path where is configured for. In this case the bandwidth is still reserved twice.

Configure LSP optimization.

a) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 optimize-timer 28800

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 adaptive

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-2 optimize-timer 28800

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-2 adaptive

[edit]
lab@R3# commit
commit complete
```

b) R4

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 optimize-timer 28800

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 adaptive

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-2 optimize-timer 28800
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-2 adaptive

[edit]
lab@R4# commit
commit complete
```

c) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 optimize-timer 28800

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 adaptive

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-2 optimize-timer 28800

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-2 adaptive

[edit]
lab@R7# commit
commit complete
```

d) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-1 optimize-timer 28800

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-1 adaptive

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-2 optimize-timer 28800

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-2 adaptive

[edit]
lab@R8# commit
commit complete
```

After committing, both LSPs should be recomputed and the optimization period is considered.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show mpls lsp name R3-to-R8-1 extensive
Ingress LSP: 3 sessions

172.30.5.8
  From: 172.30.5.3, State: Up, ActiveRoute: 0, LSPname: R3-to-R8-1
  ActivePath: path-1 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-1           State: Up
    Priorities: 6 6
    Bandwidth: 60Mbps
    OptimizeTimer: 28800
    SmartOptimizeTimer: 180
    Reoptimization in 27344 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.13 S 172.30.0.1 S 172.30.0.10 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
         172.30.0.13 172.30.0.1 172.30.0.10
    OAM state : BFD session not up LSP-ping not up
   11 Sep 22 11:12:46.059 Selected as active path
   10 Sep 22 11:12:46.058 Record Route:  172.30.0.13 172.30.0.1 172.30.0.10
    9 Sep 22 11:12:46.058 Up
    8 Sep 22 11:12:45.988 Originate Call
    7 Sep 22 11:12:45.988 CSPF: computation result accepted  172.30.0.13
172.30.0.1 172.30.0.10
    6 Sep 22 11:12:16.357 CSPF failed: no route toward 172.30.5.8[11 times]
    5 Sep 22 11:07:48.457 Clear Call: CSPF computation failed
    4 Sep 22 11:07:48.456 CSPF: link down/deleted:
172.30.0.9(172.30.5.1:74)(R1.00/172.30.5.1)-
>172.30.0.10(172.30.5.8:71)(R8.00/172.30.5.8)
    3 Sep 22 11:07:22.329 Originate Call
    2 Sep 22 11:07:22.329 CSPF: computation result accepted  172.30.0.13
172.30.0.1 172.30.0.10
    1 Sep 22 11:06:53.339 CSPF failed: no route toward 172.30.5.8[68 times]
  Created: R1 Sep 22 10:34:08 2013
Total 1 displayed, Up 1, Down 0
```

Reoptimization should take place in 27344 seconds.

```
[edit]
lab@R3# run show rsvp session ingress
Ingress RSVP: 4 sessions
To              From            State  Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.3      Up      0  1 FF       -   300656 R3-to-R6
172.30.5.8      172.30.5.3      Up      0  1 SE       -   301216 R3-to-R8-1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.8  172.30.5.3   Up      0  1 SE       -    300464 R3-to-R8-2
Total 4 displayed, Up 4, Down 0
```

Reservation style for both LSPs is changed to SE (Shared Explicit).

13) Configure R5 and R6 to install specific prefix to inet.3 routing table.

By default, only the /32 prefix is installed in the `inet.3` table. Using the `install` option, you can add additional prefixes to `inet.3`. By including the `active` parameter, the prefix is also added to the `inet.0` table, which means that the IGP can also use this prefix. Remember that by default only BGP can use prefixes in the `inet.3` table.

Configure R5 to use an LSP for forwarding traffic towards prefixes advertised by the IX peers. To accomplish this, you must ensure that the next-hop of those prefixes will be resolved by an RSVP LSP. The IX prefixes are advertised by R1 and R2 with the original next-hop of the external link, which are advertised in IS-IS to be resolved. R5 has an LSP to R2 only, so installing subnet 192.168.1.0/24 to be resolved under that LSP will accomplish the Task.

```
[edit]
lab@R5# run show route community-name IX 1.64.0.0/10 extensive

inet.0: 677 destinations, 693 routes (676 active, 0 holddown, 1 hidden)
1.64.0.0/10 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.64.0.0/10 -> {indirect(262145)}
Aggregated into 0.0.0.0/0
Page 0 idx 0 Type 1 val 182a824
    Nexthop: Self
    AS path: [54591] 1620 61671 I
    Communities: 54591:1620
Path 1.64.0.0 from 172.30.5.41 Vector len 4.  Val: 0
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Address: 0x155eeac
                Next-hop reference count: 1149
                Source: 172.30.5.41
                Next hop type: Router, Next hop index: 262146
                Next hop: 172.30.0.29 via ge-0/0/4.145, selected
                Next hop: 172.30.0.38 via ge-0/0/4.158
                Protocol next hop: 192.168.1.4
                Indirect next hop: 18102b8 262145
                State: <Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 55:22      Metric2: 30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                Task: BGP_54591.172.30.5.41+179
                Announcement bits (4): 0-KRT 9-Aggregate 11-BGP_RT_Background 12-
Resolve tree 4
                AS path: 1620 61671 I (Originator)
                Cluster list:  0.0.0.1
                Originator ID: 172.30.5.1
                Communities: 54591:1620
                Accepted
                Localpref: 100
                Router ID: 172.30.5.41
                Indirect next hops: 1
                        Protocol next hop: 192.168.1.4 Metric: 30
                        Indirect next hop: 18102b8 262145
                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 172.30.0.29 via ge-0/0/4.145
                                Next hop: 172.30.0.38 via ge-0/0/4.158
                192.168.1.0/24 Originating RIB: inet.0
                  Metric: 30                          Node path count: 1
                  Forwarding nexthops: 2
                        Nexthop: 172.30.0.29 via ge-0/0/4.145
                        Nexthop: 172.30.0.38 via ge-0/0/4.158
```

```
lab@R5# set protocols mpls label-switched-path R5-to-R2 install 192.168.1.0/24

[edit]
lab@R5# commit
commit complete
```

Verify if the route is added to inet.3

```
[edit]
lab@R5# run show route table inet.3 192.168.1.0

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.0/24     *[RSVP/7/1] 00:02:15, metric 25
                    > to 172.30.0.29 via ge-0/0/4.145, label-switched-path R5-to-
R2
```

Verify if the next-hop of prefix advertised from IX peers is now resolved through the LSP.

```
[edit]
lab@R5# run show route community-name IX 1.64.0.0/10 extensive

inet.0: 677 destinations, 693 routes (676 active, 0 holddown, 1 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
1.64.0.0/10 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.64.0.0/10 -> {indirect(262145)}
Aggregated into 0.0.0.0/0
Page 0 idx 0 Type 1 val 182a824
    Nexthop: Self
    AS path: [54591] 1620 61671 I
    Communities: 54591:1620
Path 1.64.0.0 from 172.30.5.41 Vector len 4.  Val: 0
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Address: 0x155eeac
                Next-hop reference count: 1149
                Source: 172.30.5.41
                Next hop type: Router, Next hop index: 672
                Next hop: 172.30.0.29 via ge-0/0/4.145 weight 0x1, selected
                Label-switched-path R5-to-R2
                Label operation: Push 299984
                Label TTL action: prop-ttl
                Protocol next hop: 192.168.1.4
                Indirect next hop: 18102b8 262145
                State: <Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 58:18      Metric2: 25
                Task: BGP_54591.172.30.5.41+179
                Announcement bits (4): 0-KRT 9-Aggregate 11-BGP_RT_Background 12-
Resolve tree 4
                AS path: 1620 61671 I (Originator)
                Cluster list:  0.0.0.1
                Originator ID: 172.30.5.1
                Communities: 54591:1620
                Accepted
                Localpref: 100
                Router ID: 172.30.5.41
                Indirect next hops: 1
                        Protocol next hop: 192.168.1.4 Metric: 25
                        Indirect next hop: 18102b8 262145
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 172.30.0.29 via ge-0/0/4.145 weight 0x1
                        192.168.1.0/24 Originating RIB: inet.3
                          Metric: 25                        Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 172.30.0.29 via ge-0/0/4.145
...
```

Configure the LSP to R1 on R6 to install 192.168.1.0/24 prefix.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 install 192.168.1.0/24

[edit]
lab@R6# commit
commit complete
```

Verify if prefix 192.168.1.0/24 is present in inet.3, This may take a minute to change.

```
[edit]
lab@R6# run show route table inet.3 192.168.1.0

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.1.0/24     *[RSVP/7/1] 00:01:33, metric 25
                    > to 172.30.0.33 via ae0.0, label-switched-path R6-to-R1
```

14) Configure LDP tunneling on the LSP as shown in the Figure on Page 155. Make sure that any router has LDP signaled LSP to any other router.

For this requirement LDP tunneling must be enabled on the RSVP LSP and extended LDP discovery must be enabled on the router's loopback interfaces. LDP tunneling over the RSVP LSP ensures that LDP LSPs are signaled between each and every router.

First you must enable LDP on loopback interface on all routers.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols ldp interface lo0.0

[edit]
lab@R1# commit
commit complete
```

Next you must activate the `ldp-tunneling` feature under the desired RSVP LSPs. On each router there is one LSP which has to be configured for LDP tunneling.

a) R1

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 ldp-tunneling

[edit]
lab@R1# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 ldp-tunneling

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 ldp-tunneling

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R5 ldp-tunneling

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 ldp-tunneling

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 ldp-tunneling

[edit]
lab@R6# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

g) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 ldp-tunneling

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 ldp-tunneling

[edit]
lab@R8# commit
commit complete
```

After you enable `ldp-tunneling` on the RSVP LSPs, you can verify the operation of this feature. R3 is a good candidate for this.

Verify the LDP neighbor table.

```
[edit]
lab@R3# run show ldp neighbor
Address              Interface            Label space ID         Hold time
172.30.5.6           lo0.0                172.30.5.6:0             39
172.30.0.13          ge-0/0/4.123         172.30.5.2:0             12
172.30.0.22          ge-0/0/4.134         172.30.5.4:0             11
```

R3 has an LDP neighbor with R6 loopback IP address. The LDP session is established over the RSVP tunnel. This can be seen from the routes in `inet.3` table.

```
[edit]
lab@R3# run show route table inet.3

inet.3: 15 destinations, 21 routes (9 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32   *[LDP/9] 00:10:05, metric 15
                > to 172.30.0.13 via ge-0/0/4.123, Push 299792
172.30.5.2/32   *[LDP/9] 00:08:35, metric 10
                > to 172.30.0.13 via ge-0/0/4.123, Push 0
172.30.5.4/32   *[LDP/9] 00:09:04, metric 10
                > to 172.30.0.22 via ge-0/0/4.134, Push 0
172.30.5.5/32   *[LDP/9] 00:05:11, metric 15
                > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-R6
172.30.5.6/32       *[RSVP/7/1] 00:06:38, metric 10
                    > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                    [LDP/9] 00:05:11, metric 10
                     > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
172.30.5.7/32       *[LDP/9] 00:01:00, metric 20
                    > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-R6
172.30.5.8/32       *[RSVP/7/1] 00:01:42, metric 25
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-R8-2
                    [LDP/9] 00:00:41, metric 25
                    > to 172.30.0.13 via ge-0/0/4.123, Push 300672
172.30.5.41/32      *[IS-IS/15] 00:00:04, metric 20
                    > to 172.30.0.13 via ge-0/0/4.123
192.168.1.0/24      *[LDP/9] 00:08:35, metric 20
                    > to 172.30.0.13 via ge-0/0/4.123, Push 0
```

You can see that loopback IP addresses on the other LDP island have LDP routes with a RSVP LSP as the next-hop. This RSVP LSP is enabled for LDP tunneling.

Note that there are 10 hidden routes in the output above. Those hidden routes are transparently created by Junos to support LDP tunneling. Junos enables the IGP shortcuts feature, which ensures that these IS-IS routes are resolved through RSVP LSPs. As this feature is not configured manually but by Junos itself (for LDP tunneling support) they are hidden so that they are not used to resolve BGP next-hops.

```
[edit]
lab@R3# run show route table inet.3 hidden

inet.3: 15 destinations, 21 routes (9 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.0.32/30  [IS-IS] 00:05:35, metric 15
                > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-path Bypass-
>172.30.0.22
172.30.0.36/30  [IS-IS] 00:05:35, metric 25
                > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                 to 172.30.0.13 via ge-0/0/4.123, label-switched-path Bypass-
>172.30.0.22
172.30.0.40/30  [IS-IS] 00:05:35, metric 20
                > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-path Bypass-
>172.30.0.22
172.30.0.44/30  [IS-IS] 00:01:24, metric 30
                > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-path Bypass-
>172.30.0.22
172.30.0.52/30  [IS-IS] 00:05:35, metric 17
                > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-path Bypass-
>172.30.0.22
172.30.0.56/30  [IS-IS] 00:05:35, metric 25
                > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-path Bypass-
>172.30.0.22
172.30.5.5/32   [IS-IS] 00:05:35, metric 15
                > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-path Bypass-
>172.30.0.22
172.30.5.6/32   [IS-IS] 00:05:35, metric 10
                > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-path Bypass-
>172.30.0.22
172.30.5.7/32   [IS-IS] 00:01:24, metric 20
                 > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-R6
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-path Bypass-
>172.30.0.22
172.30.5.8/32   [IS-IS] 00:01:05, metric 25
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-R8-1
                > to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-R8-2
                to 172.30.0.22 via ge-0/0/4.134, label-switched-path Bypass-
>172.30.0.13->172.30.0.1
```

15) You must ensure that IPv4 traffic from P1 to P2 uses LSP I. Traffic from P1 to P3 must use LSP K.

To accomplish this step, you must configure a forwarding policy which matches prefixes based on their attached community and change the next-hop based on the requirements.

First let's check the routing table before we make any change.

```
[edit]
lab@R8# run show route community-name P2

inet.0: 834 destinations, 1236 routes (671 active, 0 holddown, 549 hidden)
+ = Active Route, - = Last Active, * = Both

12.173.96.0/19     *[BGP/170] 00:18:57, localpref 200, from 172.30.5.41
                      AS path: 3521382357 4637 7478 16572 33786 ?
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
15.36.0.0/15       *[BGP/170] 00:18:57, localpref 200, from 172.30.5.41
                      AS path: 3521382357 4637 57359 16881 43174 64323 I
                      to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                    > to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
17.96.0.0/13       *[BGP/170] 00:18:57, localpref 200, from 172.30.5.41
                      AS path: 3521382357 4637 22896 I
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
26.0.0.0/8         *[BGP/170] 00:18:57, localpref 200, from 172.30.5.41
...

[edit]
lab@R8# run show route community-name P3

inet.0: 834 destinations, 1236 routes (671 active, 0 holddown, 549 hidden)
+ = Active Route, - = Last Active, * = Both

5.127.0.0/17       *[BGP/170] 00:18:27, localpref 200, from 172.30.5.41
                      AS path: 2831679853 9726 36659 30705 25538 37414 49276 ?
                      to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                    > to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
10.128.0.0/11      *[BGP/170] 00:18:27, localpref 200, from 172.30.5.41
                      AS path: 2831679853 26697 4341 43012 28104 39181 51157 ?
                      to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                    > to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
17.205.88.0/21     *[BGP/170] 00:18:27, localpref 200, from 172.30.5.41
                      AS path: 2831679853 9726 19412 16490 6122 I
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
18.100.0.0/17      *[BGP/170] 00:18:27, localpref 200, from 172.30.5.41
...
```

Junos randomly chooses the next-hop for the traffic to P2 and P3.

Define the policy to change the next-hop based on the community.

```
[edit]
lab@R8# set policy-options policy-statement lsp-map term 1 from protocol bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set policy-options policy-statement lsp-map term 1 from community P2

[edit]
lab@R8# set policy-options policy-statement lsp-map term 1 then install-nexthop lsp R8-
to-R3-1

[edit]
lab@R8# set policy-options policy-statement lsp-map term 2 from protocol bgp

[edit]
lab@R8# set policy-options policy-statement lsp-map term 2 from community P3

[edit]
lab@R8# set policy-options policy-statement lsp-map term 2 then install-nexthop lsp R8-
to-R3-2
```

Apply the policy to the forwarding-table.

```
[edit]
lab@R8# set routing-options forwarding-table export lsp-map

[edit]
lab@R8# commit
commit complete
```

The policy is applied to the forwarding table, the same way as the load-balancing policy.

Verify the result of this change.

```
[edit]
lab@R8# run show route community-name P2

inet.0: 834 destinations, 1236 routes (671 active, 0 holddown, 549 hidden)
+ = Active Route, - = Last Active, * = Both

12.173.96.0/19     *[BGP/170] 00:01:59, localpref 200, from 172.30.5.41
                      AS path: 3521382357 4637 7478 16572 33786 ?
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-R3-1
15.36.0.0/15       *[BGP/170] 00:01:59, localpref 200, from 172.30.5.41
                      AS path: 3521382357 4637 57359 16881 43174 64323 I
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-R3-
1
17.96.0.0/13       *[BGP/170] 00:01:59, localpref 200, from 172.30.5.41
                      AS path: 3521382357 4637 22896 I
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-R3-
1
26.0.0.0/8         *[BGP/170] 00:01:59, localpref 200, from 172.30.5.41
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    AS path: 3521382357 46848 I
                    to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-R3-
1
...
```

```
[edit]
lab@R8# run show route community-name P3

inet.0: 834 destinations, 1236 routes (671 active, 0 holddown, 549 hidden)
+ = Active Route, - = Last Active, * = Both

5.127.0.0/17       *[BGP/170] 00:00:32, localpref 200, from 172.30.5.41
                      AS path: 2831679853 9726 36659 30705 25538 37414 49276 ?
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
10.128.0.0/11      *[BGP/170] 00:00:32, localpref 200, from 172.30.5.41
                      AS path: 2831679853 26697 4341 43012 28104 39181 51157 ?
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
17.205.88.0/21     *[BGP/170] 00:00:32, localpref 200, from 172.30.5.41
                      AS path: 2831679853 9726 19412 16490 6122 I
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
18.100.0.0/17      *[BGP/170] 00:00:32, localpref 200, from 172.30.5.41
                      AS path: 2831679853 9726 19412 40487 63999 13005 28592 5369
9061 I
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
```

16) Configure a per flow load balancing policy on R4 and R7.

This task requires that you configure load balancing over two LSPs from R4 to R7 and vice versa. This is accomplished with a policy applied to the forwarding table (the same way as in previous tasks).

Display the forwarding table for R7 loopback IP address before applying the policy.

```
[edit]
lab@R4# run show route forwarding-table destination 172.30.5.7
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
172.30.5.7/32      user     0 172.30.0.5         ucst   586    26 ge-0/0/4.114

Routing table: __master.anon__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
default             perm     0                      rjct   528     1
...
```

A single next hop is present in the forwarding table, although in the inet.3 table two next hops are available.

```
[edit]
lab@R4# run show route table inet.3 172.30.5.7

inet.3: 13 destinations, 19 routes (8 active, 0 holddown, 9 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.7/32       *[RSVP/7/1] 00:00:51, metric 25
                     > to 172.30.0.21 via ge-0/0/4.134, label-switched-path R4-to-R7-1
                       to 172.30.0.30 via ge-0/0/4.145, label-switched-path R4-to-R7-2
                     [LDP/9] 00:00:40, metric 25
                     > to 172.30.0.5 via ge-0/0/4.114, Push 300400
```

Create a policy to enable per flow load-balancing.

Repeat this on all eight devices.

```
[edit]
lab@R4# set policy-options policy-statement load-balance term 1 then load-balance per-
packet

[edit]
lab@R4# set routing-options forwarding-table export load-balance

[edit]
lab@R4# commit
commit complete
```

Note: This policy will have load balancing effect on all prefixes not only for the LSPs to R7.
The policy has just one term with a then statement without any match criteria.

Verify the result of the policy.

```
[edit]
lab@R4# run show route forwarding-table destination 172.30.5.7
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
172.30.5.7/32      user     0                    ulst 262144     4
                                172.30.0.5       ucst   586    24 ge-0/0/4.114
                                172.30.0.30      ucst   591    11 ge-0/0/4.145
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Routing table: __master.anon__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   528     1
...
```

The prefix has now both LSPs as next hop.

17) Configure all routers to not decrement the TTL.

By default, the TTL value in the packet is decremented by every LSR on the LSP the packet traverses. This prevents loops and used for topology discovery. Most Service Providers do not want to reveal the topology of their network.

The `no-decrement-ttl` feature alters the default behavior. It is a proprietary mechanism of Juniper Networks and signaled by RSVP. It can be configured per LSP or for all LSPs and it is configured only on the ingress router. It keeps the IP packet TTL value intact and does not copy it the MPLS label. Instead it sets a value of 255 on the label. The TTL of the IP packet is only decremented on the egress LSR.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols mpls no-decrement-ttl

[edit]
lab@R1# commit
commit complete
```

Remember that only BGP is using `inet.3` for resolving next-hops. To force other prefixes to use LSPs for forwarding traffic you can move routes from `inet.3` to `inet.0`. This can be accomplished with following option.

```
[edit]
lab@R1# set protocols mpls traffic-engineering mpls-forwarding

[edit]
lab@R1# commit
commit complete
```

That knob tells Junos to copy all `inet.3` routes to `inet.0`.

```
[edit]
lab@R1# run show route 172.30.5.5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 833 destinations, 1425 routes (670 active, 0 holddown, 325 hidden)
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both

172.30.5.5/32      @[IS-IS/15] 00:00:48, metric 20
                       to 172.30.0.6 via ge-0/0/4.114
                     > to 172.30.0.10 via ge-0/0/4.118
                    #[LDP/9] 00:00:47, metric 20
                     > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R8

inet.3: 11 destinations, 16 routes (7 active, 0 holddown, 7 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.5/32      *[LDP/9] 00:00:47, metric 20
                     > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R8
```

Note: IS-IS routes are with "@" in front. This marks the routes eligible for policy evaluation but not for forwarding traffic. The LDP route is used for forwarding.

Note: The following traceroute was completed prior to configuring the `no-decrement-ttl` feature.

```
[edit]
lab@R1# run traceroute 172.30.5.5 source 172.30.5.1
traceroute to 172.30.5.5 (172.30.5.5) from 172.30.5.1, 30 hops max, 40 byte
packets
 1  172.30.0.2 (172.30.0.2)  2.426 ms  2.325 ms  7.663 ms
     MPLS Label=300512 CoS=0 TTL=1 S=0
     MPLS Label=299776 CoS=0 TTL=1 S=1
 2  172.30.0.18 (172.30.0.18)  2.806 ms  2.720 ms  2.172 ms
     MPLS Label=300096 CoS=0 TTL=1 S=0
     MPLS Label=299776 CoS=0 TTL=2 S=1
 3  172.30.0.46 (172.30.0.46)  2.299 ms  2.395 ms  2.404 ms
     MPLS Label=299776 CoS=0 TTL=1 S=1
 4  172.30.5.5 (172.30.5.5)  3.579 ms  2.736 ms  2.413 ms
```

When you configure the `no-decrement-ttl` feature, and issue again the same traceroute command the result is much different.

```
[edit]
lab@R1# run traceroute 172.30.5.5 source 172.30.5.1
traceroute to 172.30.5.5 (172.30.5.5) from 172.30.5.1, 30 hops max, 40 byte packets
 1  172.30.5.5 (172.30.5.5)  3.222 ms  8.728 ms  8.158 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Only the egress router is responding to the traceroute message. If you issue a traceroute from a router connected to the ingress PE, you will see the ingress PE as well as the egress PE.

Task 3: RSVP Protection

By default, when network failures occur along the path of RSVP signaled LSPs, traffic will be dropped. The ingress LSR continues to send traffic until a PathTear message arrives. Then the ingress LSR will try to signal new path if possible. RSVP provides different traffic protection mechanisms, which can significantly improve network restoration.

1) Configure secondary paths for all LSP but K, L, O and P.

Secondary paths are established only if the primary path fails. After failure the ingress router always tries to establish the primary path again. An LSP can only have one primary path, but multiple secondary paths are allowed. Secondary paths are signaled in the order they appear in the configuration.

On R1 both LSP should be protected with a primary and a secondary path. Paths are configured to use CSPF routes without strict ERO configuration.

a) R1

```
[edit]
lab@R1# set protocols mpls path primary-1

[edit]
lab@R1# set protocols mpls path primary-2

[edit]
lab@R1# set protocols mpls path secondary-1

[edit]
lab@R1# set protocols mpls path secondary-2

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 primary primary-1

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 secondary secondary-1

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 primary primary-2

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 secondary secondary-2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols mpls path primary-1

[edit]
lab@R2# set protocols mpls path primary-2

[edit]
lab@R2# set protocols mpls path secondary-1

[edit]
lab@R2# set protocols mpls path secondary-2

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R5 primary primary-1

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R5 secondary secondary-1

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 primary primary-2

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 secondary secondary-2

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols mpls path primary-1

[edit]
lab@R3# set protocols mpls path secondary-1

[edit]
lab@R3# set protocols mpls path secondary-2

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 primary primary-1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 secondary secondary-1

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 secondary secondary-2

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols mpls path primary-1

[edit]
lab@R4# set protocols mpls path secondary-1

[edit]
lab@R4# set protocols mpls path secondary-2

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R5 primary primary-1

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R5 secondary secondary-1

[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 secondary secondary-2

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols mpls path primary-1

[edit]
lab@R5# set protocols mpls path primary-2

[edit]
lab@R5# set protocols mpls path secondary-1

[edit]
lab@R5# set protocols mpls path secondary-2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 primary primary-1

[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 secondary secondary-1

[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R2 primary primary-2

[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R2 secondary secondary-2

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols mpls path primary-1

[edit]
lab@R6# set protocols mpls path primary-2

[edit]
lab@R6# set protocols mpls path secondary-1

[edit]
lab@R6# set protocols mpls path secondary-2

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 primary primary-1

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 secondary secondary-1

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 primary primary-2

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 secondary secondary-2

[edit]
lab@R6# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

g) R7

```
[edit]
lab@R7# set protocols mpls path primary-1

[edit]
lab@R7# set protocols mpls path secondary-1

[edit]
lab@R7# set protocols mpls path secondary-2

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 primary primary-1

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 secondary secondary-1

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 secondary secondary-2

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols mpls path primary-1

[edit]
lab@R8# set protocols mpls path secondary-1

[edit]
lab@R8# set protocols mpls path secondary-2

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 primary primary-1

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 secondary secondary-1

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-1 secondary secondary-2

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify the LSP to R8 after the changes are committed.

```
[edit]
lab@R1# run show mpls lsp name R1-to-R8 extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R8
  ActivePath: primary-2 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 30Mbps MaxBW: 120Mbps
  AdjustTimer: 172800 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 170933 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary   primary-2         State: Up, No-decrement-ttl
    Priorities: 6 6
    Bandwidth: 30Mbps
    SmartOptimizeTimer: 180
          Include Any: green
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.18 S 172.30.0.46 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.2 172.30.0.18 172.30.0.46
    OAM state : BFD session up  LSP-ping up
    5 Sep 22 13:15:05.356 Selected as active path
    4 Sep 22 13:15:05.356 Record Route:  172.30.0.2 172.30.0.18 172.30.0.46
    3 Sep 22 13:15:05.356 Up
    2 Sep 22 13:15:05.235 Originate Call
    1 Sep 22 13:15:05.235 CSPF: computation result accepted  172.30.0.2
172.30.0.18 172.30.0.46
  Secondary secondary-2       State: Dn, No-decrement-ttl
    Priorities: 6 6
    Bandwidth: 30Mbps
    SmartOptimizeTimer: 180
          Include Any: green
        No computed ERO.
    OAM state : BFD session not up LSP-ping not up
  Created: R1 Sep 22 13:15:05 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

The secondary path is in the Down state. It is waiting for primary path to fail. All parameters are inherited from the configuration under label-switched-path.

2) Configure the standby option for LSPs C, D, G and H secondary paths.

As you saw in the previous Step the secondary path is Down and waiting for the primary to fail. This leads to dropped traffic until secondary path is signaled and forward the traffic. Specifying **`standby`** under the secondary path forces the LSR to signal the path even though the primary has not failed yet. The Ingress router pushes traffic only to the primary path. The secondary path will already be established, but still waits until the primary fails. This reduces failover times as the secondary path is pre-established, but also requires additional state in the LSRs for the standby LSPs.

Configure a secondary path with the `standby` option for the LSPs.

a) R1

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 secondary secondary-1 standby

[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R5 secondary secondary-1 standby

[edit]
lab@R2# commit
commit complete
```

c) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R2 secondary secondary-1 standby

[edit]
lab@R5# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

d) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 secondary secondary-1 standby

[edit]
lab@R6# commit
commit complete
```

After committing the changes, verify the LSP to R6.

```
[edit]
lab@R1# run show mpls lsp name R1-to-R6 extensive
Ingress LSP: 2 sessions

172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: primary-1 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary-1          State: Up, No-decrement-ttl
    Priorities: 7 7
    Bandwidth: 60Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.6 S 172.30.0.30 S 172.30.0.34 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.6 172.30.0.30 172.30.0.34
    OAM state : BFD session up  LSP-ping up
   15 Sep 22 13:32:37.475 Selected as active path: due to 'primary'
   14 Sep 22 13:31:36.625 Record Route:  172.30.0.6 172.30.0.30 172.30.0.34
   13 Sep 22 13:31:36.625 Up
   12 Sep 22 13:31:36.559 Originate Call
   11 Sep 22 13:31:36.558 Clear Call
   10 Sep 22 13:31:36.558 CSPF: computation result accepted  172.30.0.6
172.30.0.30 172.30.0.34
    9 Sep 22 13:31:14.686 Deselected as active
    8 Sep 22 13:31:14.685 Requested bandwidth unavailable
    7 Sep 22 13:31:14.683 Session preempted
    6 Sep 22 13:31:14.682 172.30.0.9: Down
    5 Sep 22 13:15:05.378 Selected as active path
    4 Sep 22 13:15:05.377 Record Route:  172.30.0.10 172.30.0.37 172.30.0.34
    3 Sep 22 13:15:05.377 Up
    2 Sep 22 13:15:05.244 Originate Call
    1 Sep 22 13:15:05.244 CSPF: computation result accepted  172.30.0.10
172.30.0.37 172.30.0.34
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Standby    secondary-1       State: Up, No-decrement-ttl
    Priorities: 7 7
    Bandwidth: 60Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.18 S 172.30.0.41 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.2 172.30.0.18 172.30.0.41
    OAM state : BFD session up  LSP-ping up
   11 Sep 22 13:49:11.849 Record Route:  172.30.0.2 172.30.0.18 172.30.0.41
   10 Sep 22 13:49:11.849 Up
    9 Sep 22 13:49:11.788 Originate Call
    8 Sep 22 13:49:11.788 CSPF: computation result accepted  172.30.0.2
172.30.0.18 172.30.0.41
    7 Sep 22 13:34:05.497 Clear Call
    6 Sep 22 13:32:37.475 Deselected as active: due to 'primary'
    5 Sep 22 13:31:14.760 Selected as active path
    4 Sep 22 13:31:14.760 Record Route:  172.30.0.6 172.30.0.30 172.30.0.34
    3 Sep 22 13:31:14.760 Up
    2 Sep 22 13:31:14.687 Originate Call
    1 Sep 22 13:31:14.687 CSPF: computation result accepted  172.30.0.6
172.30.0.30 172.30.0.34
  Created: R1 Sep 22 13:15:06 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 3 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 9 sessions
Total 0 displayed, Up 0, Down 0
```

The secondary path is signaled and in state UP. In this specific case, the path taken by the secondary LSP is different than the primary LSP, because there was not enough bandwidth on the path used for the primary LSP. It is recommended for the primary and secondary path to have EROs defined so that it is sure that they pass different LSRs in the network.

The other ingress LSRs/routers should be configured in the same way.

3) Configure the `adaptive` option for LSPs C, D, G, and H to change the reservation from (FF) Fixed Filter to (SE) Shared Explicit.

As you saw in step 12 in the previous task, configuring the `adaptive` option for LSPs enable the make-before-brake capability and it changes the RSVP reservation style. In this Step you have to change the reservation style for the given LSPs.

© 2020 Juniper Networks, Inc. All rights reserved.

The reservation style for the LSP from R1 to R6 should be changed

Let's verify the reservation style before we change it.

```
[edit]
lab@R1# run show rsvp session name R1-to-R6
Ingress RSVP: 3 sessions
To              From            State  Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.1      Up      0  1 FF       -    300576 R1-to-R6
172.30.5.6      172.30.5.1      Up      0  1 FF       -    300560 R1-to-R6
Total 2 displayed, Up 2, Down 0

Egress RSVP: 3 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 8 sessions
Total 0 displayed, Up 0, Down 0
```

Note: When you display the RSVP sessions you see both RSVP sessions, one for primary path and second for the secondary. They belong to a single LSP.

a) R1

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 adaptive

[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R5 adaptive

[edit]
lab@R2# commit
commit complete
```

c) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R2 adaptive

[edit]
lab@R5# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

d) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 adaptive

[edit]
lab@R6# commit
commit complete
```

Verify the reservation after the change.

```
[edit]
lab@R1#
Ingress RSVP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.1      Up       0  1 SE       -    300816 R1-to-R6
172.30.5.6      172.30.5.1      Up       0  1 SE       -    301296 R1-to-R6
Total 2 displayed, Up 2, Down 0

Egress RSVP: 2 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 8 sessions
Total 0 displayed, Up 0, Down 0
```

Shared Explicit (SE) means that if both paths must pass over single link, the reserved bandwidth will be counted just once.

4) Configure the revert timer for LSPs E, F, Q and R.

As it was stated if a secondary path is configured the ingress LSR fails over to it even though there is another physical path that complies the primary path constraints. When the ingress LSR signals the primary path on different physical path, traffic is reverted back from the secondary path to the new primary path. By default, the minimum time that the primary path must be UP before the traffic to revert to it is 60 seconds. By Configure that time to 0 ensures the LSP will not revert back to the primary path. Another way to disable the revert-timer is to configure two secondary paths instead of configuring a primary and a secondary path.

Configure revert-timer on the respective devices.

a) R2

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 revert-timer 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# commit
commit complete
```

b) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 revert-timer 0

[edit]
lab@R3# commit
commit complete
```

c) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 revert-timer 0

[edit]
lab@R6# commit
commit complete
```

d) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 revert-timer 0

[edit]
lab@R7# commit
commit complete
```

Verify the result of the changed revert-timer.

```
lab@R7# run show mpls lsp extensive name R7-to-R2
Ingress LSP: 3 sessions

172.30.5.2
  From: 172.30.5.7, State: Up, ActiveRoute: 0, LSPname: R7-to-R2
  ActivePath: primary-2 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Revert timer: 0
 *Primary   primary-2         State: Up, No-decrement-ttl
    Priorities: 6 6
    Bandwidth: 60Mbps
    SmartOptimizeTimer: 180
          Include Any: red
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.46 S 172.30.0.9 S 172.30.0.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.46 172.30.0.9 172.30.0.2
    OAM state : BFD session up  LSP-ping up
   20 Sep 22 14:22:38.454 Selected as active path
   19 Sep 22 14:22:38.453 Record Route:  172.30.0.46 172.30.0.9 172.30.0.2
   18 Sep 22 14:22:38.453 Up
   17 Sep 22 14:22:38.347 Originate Call
   16 Sep 22 14:22:38.347 CSPF: computation result accepted  172.30.0.46
172.30.0.9 172.30.0.2
   15 Sep 22 14:22:31.473 CSPF failed: no route toward 172.30.5.2
   14 Sep 22 14:22:31.472 Clear Call: CSPF computation failed
   13 Sep 22 14:22:31.471 CSPF: link down/deleted:
172.30.0.14(172.30.5.3:70)(R3.00/172.30.5.3)-
>172.30.0.13(172.30.5.2:73)(R2.00/172.30.5.2)
   12 Sep 22 14:22:31.422 Originate Call
   11 Sep 22 14:22:31.421 Clear Call
   10 Sep 22 14:22:31.421 CSPF: computation result accepted  172.30.0.41
172.30.0.33 172.30.0.29 172.30.0.21 172.30.0.13
    9 Sep 22 14:22:31.420 172.30.0.9: Session preempted
    8 Sep 22 14:22:31.419 Deselected as active
    7 Sep 22 14:22:31.416 ResvTear received
    6 Sep 22 14:22:31.415 172.30.0.45: Down
    5 Sep 22 13:41:13.424 Selected as active path
    4 Sep 22 13:41:13.424 Record Route:  172.30.0.46 172.30.0.9 172.30.0.2
    3 Sep 22 13:41:13.423 Up
    2 Sep 22 13:41:13.300 Originate Call
    1 Sep 22 13:41:13.300 CSPF: computation result accepted  172.30.0.46
172.30.0.9 172.30.0.2
  Secondary secondary-2      State: Dn, No-decrement-ttl
    Priorities: 6 6
    Bandwidth: 60Mbps
    SmartOptimizeTimer: 180
          Include Any: red
        No computed ERO.
    OAM state : BFD session not up LSP-ping not up
   10 Sep 22 14:23:36.589 Clear Call
    9 Sep 22 14:22:38.459 Record Route:  172.30.0.46 172.30.0.9 172.30.0.2
    8 Sep 22 14:22:38.459 Up
    7 Sep 22 14:22:38.348 Originate Call
    6 Sep 22 14:22:38.348 CSPF: computation result accepted  172.30.0.46
172.30.0.9 172.30.0.2
    5 Sep 22 14:22:31.480 CSPF failed: no route toward 172.30.5.2
    4 Sep 22 14:22:31.478 Clear Call: CSPF computation failed
    3 Sep 22 14:22:31.478 CSPF: link down/deleted:
172.30.0.1(172.30.5.1:69)(R1.00/172.30.5.1)-
>172.30.0.2(172.30.5.2:69)(R2.00/172.30.5.2)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    2 Sep 22 14:22:31.432 Originate Call
    1 Sep 22 14:22:31.432 CSPF: computation result accepted  172.30.0.46
172.30.0.9 172.30.0.2
  Created: R1 Sep 22 13:41:12 2013
```

5) Configure fast reroute for LSPs C, D, G and H.

Fast reroute can dramatically reduce packet loss in the case a primary path fails. Even if you have a secondary path signaled with the **`standby`** feature there is still the possibility of many packet drops. This happens because the message for the failed path must still travel back towards the ingress LSR so it can be informed and switchover the traffic. Fast reroute provides a way for intermediate LSRs to immediately start to transmit traffic on an alternate path and simultaneously inform the Ingress LSR about the downstream link or node failure. Fast reroute paths are called detours. Fast reroute is configured on the Ingress LSR and is signaled with a RSVP Object. By default, the TED database is used to calculate the detour paths. The administrative groups are inherited from the LSP and up to six hops can be added to the LSP in attempt to reach the egress LSR.

Fast reroute is a short-term solution. A PathErr message is send to the Ingress LSR to resignal the primary LSP or use the secondary. If a detour is not available a ResvTear is sent and label reservations are removed and the LSP is deleted. On the other side if a new primary path is not available, the detour path could stay UP for an indefinite time.

Because only the administrative groups inherit it is possible that detours have less bandwidth available. You can configure bandwidth, hop limit, include and exclude administrative groups.

Configure the fast reroute feature for LSP from R1 to R6.

a) R1

```
[edit]
lab@R1# set protocol mpls label-switched-path R1-to-R6 fast-reroute hop-limit 5

[edit]
lab@R1# set protocol mpls label-switched-path R1-to-R6 fast-reroute no-include-any

[edit]
lab@R1# commit
commit complete
```

b) R6

```
[edit]
lab@R6# set protocol mpls label-switched-path R6-to-R1 fast-reroute hop-limit 5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocol mpls label-switched-path R6-to-R1 fast-reroute no-include-any

[edit]
lab@R6# commit
commit complete
```

Fast reroute is configured with a hop limit of 5, which is lower than the default of 6 hops. no-include-any is added so that the administrative groups are not inherit for the LSP.

Verify the MPLS LSP table for the LSP to R6.

Note: You might need to wait a few moments for the Fast reroute detour to be established.

```
[edit]
lab@R1# run show mpls lsp name R1-to-R6 extensive
Ingress LSP: 6 sessions

172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: primary-1 (primary)
  FastReroute desired
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary-1         State: Up, No-decrement-ttl
    Priorities: 7 7
    Bandwidth: 60Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.14 S 172.30.0.26 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.2(flag=9) 172.30.0.14(flag=1) 172.30.0.26
    OAM state : BFD session up  LSP-ping up
    9 Nov  1 04:23:00.666 Fast-reroute Detour Up
    8 Nov  1 04:22:54.821 Record Route:  172.30.0.2(flag=9) 172.30.0.14(flag=1)
172.30.0.26
    7 Nov  1 04:22:54.799 Record Route:  172.30.0.2 172.30.0.14(flag=1)
172.30.0.26
    6 Nov  1 04:22:51.725 Selected as active path
    5 Nov  1 04:22:51.725 Record Route:  172.30.0.2 172.30.0.14 172.30.0.26
    4 Nov  1 04:22:51.724 Up
    3 Nov  1 04:22:51.649 Originate Call
    2 Nov  1 04:22:51.649 CSPF: computation result accepted  172.30.0.2
172.30.0.14 172.30.0.26
    1 Nov  1 04:22:22.811 CSPF failed: no route toward 172.30.5.6[5 times]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Standby   secondary-1      State: Up, No-decrement-ttl
    Priorities: 7 7
    Bandwidth: 60Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.14 S 172.30.0.26 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.2(flag=9) 172.30.0.14(flag=1) 172.30.0.26
    OAM state : BFD session up  LSP-ping up
   12 Nov  1 04:24:38.693 Fast-reroute Detour Up
   11 Nov  1 04:22:54.821 Record Route:  172.30.0.2(flag=9) 172.30.0.14(flag=1)
172.30.0.26
   10 Nov  1 04:22:54.799 Record Route:  172.30.0.2 172.30.0.14(flag=1)
172.30.0.26
    9 Nov  1 04:22:51.823 Record Route:  172.30.0.2 172.30.0.14 172.30.0.26
    8 Nov  1 04:22:51.822 Up
    7 Nov  1 04:22:51.663 Originate Call
    6 Nov  1 04:22:51.662 Clear Call
    5 Nov  1 04:22:51.662 CSPF: computation result accepted  172.30.0.2
172.30.0.14 172.30.0.26
    4 Nov  1 04:22:51.661 172.30.0.6: Requested bandwidth unavailable
    3 Nov  1 04:22:51.651 Originate Call
    2 Nov  1 04:22:51.651 CSPF: computation result accepted  172.30.0.6
172.30.0.30 172.30.0.34
    1 Nov  1 04:22:22.811 CSPF failed: no route toward 172.30.5.6[4 times]
  Created: Fri Nov  1 04:18:04 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 8 sessions, 2 detours
Total 0 displayed, Up 0, Down 0

Transit LSP: 14 sessions, 2 detours
Total 0 displayed, Up 0, Down 0
```

It is marked that the fast reroute is desired and that the detour is UP. To verify the detour, you must check a transit LSRs for that LSP. For both primary and secondary paths fast reroute detours are signaled. For example, let's look at the primary path. The physical path calculated by CSPF is from R1, R2, R3 towards R6.

The primary LSP transits R3. Verify the Transit LSP on R3.

```
[edit]
lab@R3# run show mpls lsp transit name R1-to-R6 extensive
Transit LSP: 28 sessions, 4 detours

172.30.5.6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
From: 172.30.5.1, LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 SE, Label in: 300288, Label out: 3
Time left:  152, Since: Fri Nov  1 04:22:49 2013
Tspec: rate 60Mbps size 60Mbps peak Infbps m 20 M 1500
Port number: sender 1 receiver 18138 protocol 0
FastReroute desired
PATH rcvfrom: 172.30.0.13 (ge-0/0/4.123) 25 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 172.30.0.26 (ge-0/0/4.136) 24 pkts
RESV rcvfrom: 172.30.0.26 (ge-0/0/4.136) 29 pkts
Explct route: 172.30.0.26
Record route: 172.30.0.1 172.30.0.13 <self> 172.30.0.26
  Detour is Up
  Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Detour adspec: received MTU 1500 sent MTU 1500
  Path MTU: received 1500
  Detour PATH sentto: 172.30.0.22 (ge-0/0/4.134) 22 pkts
  Detour RESV rcvfrom: 172.30.0.22 (ge-0/0/4.134) 21 pkts
  Detour Explct route: 172.30.0.22 172.30.0.30 172.30.0.34
  Detour Record route: 172.30.0.1 172.30.0.13 <self> 172.30.0.22
  172.30.0.30 172.30.0.34
  Detour Label out: 300128

Total 1 displayed, Up 1, Down 0
```

On R3 a detour is created through R4 to protect the direct link between R3 and R6.

```
[edit]
lab@R4# run show mpls lsp transit name R1-to-R6 extensive
Transit LSP: 29 sessions, 4 detours

172.30.5.6
  From: 172.30.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299904
  Resv style: 1 SE, Label in: 300112, Label out: 299904
  Time left:  153, Since: Fri Nov  1 04:22:54 2013
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 18138 protocol 0
  Detour branch from 172.30.0.5, to skip 172.30.5.2, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
  Detour branch from 172.30.0.2, to skip 172.30.5.3, Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 172.30.0.5 (ge-0/0/4.114) 35 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 172.30.0.30 (ge-0/0/4.145) 15 pkts
    RESV rcvfrom: 172.30.0.30 (ge-0/0/4.145) 23 pkts
    Explct route: 172.30.0.30 172.30.0.34
    Record route: 172.30.0.5 <self> 172.30.0.30 172.30.0.34
    Label in: 300112, Label out: 299904
  Detour branch from 172.30.0.21, to skip 172.30.5.6, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 172.30.0.21 (ge-0/0/4.134) 34 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 172.30.0.30 (ge-0/0/4.145) 18 pkts
    RESV rcvfrom: 172.30.0.30 (ge-0/0/4.145) 34 pkts
    Explct route: 172.30.0.30 172.30.0.34
    Record route: 172.30.0.1 172.30.0.13 172.30.0.21 <self> 172.30.0.30
    172.30.0.34
    Label in: 300128, Label out: 299904

172.30.5.6
  From: 172.30.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Secondary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299920
  Resv style: 1 SE, Label in: 300144, Label out: 299920
  Time left:  153, Since: Fri Nov  1 04:22:54 2013
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 3 receiver 18138 protocol 0
  Detour branch from 172.30.0.5, to skip 172.30.5.2, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
  Detour branch from 172.30.0.2, to skip 172.30.5.3, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 172.30.0.5 (ge-0/0/4.114) 34 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 172.30.0.30 (ge-0/0/4.145) 17 pkts
    RESV rcvfrom: 172.30.0.30 (ge-0/0/4.145) 32 pkts
    Explct route: 172.30.0.30 172.30.0.34
    Record route: 172.30.0.5 <self> 172.30.0.30 172.30.0.34
    Label in: 300144, Label out: 299920
  Detour branch from 172.30.0.21, to skip 172.30.5.6, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 172.30.0.21 (ge-0/0/4.134) 34 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 172.30.0.30 (ge-0/0/4.145) 16 pkts
    RESV rcvfrom: 172.30.0.30 (ge-0/0/4.145) 28 pkts
    Explct route: 172.30.0.30 172.30.0.34
    Record route: 172.30.0.1 172.30.0.13 172.30.0.21 <self> 172.30.0.30
    172.30.0.34
    Label in: 300160, Label out: 299920
Total 2 displayed, Up 2, Down 0
```

The detour on R4 transits R5 and reaches the Egress of the LSP, R6, over a different physical path.

6) Configure link protection.

The Junos implementation of facility backup is `link protection`. This feature is interface based rather than LSP based. It should be configured on each RSVP enabled interface where protection is desired and on the LSP itself to make it eligible to make use of the facility backup. Whereas fast reroute is trying to protect the whole LSP (1:1 protection), link protection (1:n) protects only the link. However, it can protect multiple LSPs. A bypass LSP must terminate on the downstream LSR. Bypass LSPs are calculated by CSPF. You can specify administrative groups and bandwidth for Bypass LSPs. This configuration is done on individual RSVP interfaces under link-protection.

Node protection is similar to the link protection, but in this case whole downstream node is protected.

Enable link-protection on all RSVP interfaces.

a) R1

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.114 link-protection

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.118 link-protection

[edit]
lab@R1# set protocols rsvp interface ae0.0 link-protection

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) R2

```
[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.123 link-protection

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.127 link-protection

[edit]
lab@R2# set protocols rsvp interface ae0.0 link-protection

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.123 link-protection

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.134 link-protection

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.136 link-protection

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.114 link-protection

[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.134 link-protection

[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.145 link-protection

[edit]
lab@R4# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

e) R5

```
[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.145 link-protection

[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.158 link-protection

[edit]
lab@R5# set protocols rsvp interface ae0.0 link-protection

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.136 link-protection

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.167 link-protection

[edit]
lab@R6# set protocols rsvp interface ae0.0 link-protection

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.127 link-protection


[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.167 link-protection

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.178 link-protection

[edit]
lab@R7# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

h) R8

```
[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.118 link-protection

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.158 link-protection

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.178 link-protection

[edit]
lab@R8# commit
commit complete
```

Configure link protection for LSPs A, B, E, F, Q, R, S and T.

a) R1

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 link-protection

[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 link-protection

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 link-protection

[edit]
lab@R3# commit
commit complete
```

d) R4

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R5 link-protection

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path R5-to-R4 link-protection

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 link-protection

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 link-protection

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 link-protection

[edit]
lab@R8# commit
commit complete
```

Verify the LSP state after enabling link-protection.

```
lab@R1# run show mpls lsp name R1-to-R8 extensive
Ingress LSP: 2 sessions
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: R1-to-R8
  ActivePath: primary-2 (primary)
  Link protection desired
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 30Mbps MaxBW: 120Mbps
  AdjustTimer: 172800 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 171607 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary-2         State: Up, No-decrement-ttl
    Priorities: 6 6
    Bandwidth: 30Mbps
    SmartOptimizeTimer: 180
          Include Any: green
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.2 S 172.30.0.18 S 172.30.0.46 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.5.2(flag=0x20) 172.30.0.2(Label=301664) 172.30.5.7(flag=0x20)
172.30.0.18(Label=300896) 172.30.5.8(flag=0x20) 172.30.0.46(Label=3)
    OAM state : BFD session up  LSP-ping up
    5 Sep 22 15:01:46.788 Selected as active path
    4 Sep 22 15:01:46.787 Record Route:  172.30.5.2(flag=0x20)
172.30.0.2(Label=301664) 172.30.5.7(flag=0x20) 172.30.0.18(Label=300896)
172.30.5.8(flag=0x20) 172.30.0.46(Label=3)
    3 Sep 22 15:01:46.787 Up
    2 Sep 22 15:01:46.661 Originate Call
    1 Sep 22 15:01:46.661 CSPF: computation result accepted  172.30.0.2
172.30.0.18 172.30.0.46
  Secondary secondary-2       State: Dn, No-decrement-ttl
    Priorities: 6 6
    Bandwidth: 30Mbps
    SmartOptimizeTimer: 180
          Include Any: green
        No computed ERO.
    OAM state : BFD session not up LSP-ping not up
  Created: R1 Sep 22 15:01:46 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 3 sessions, 3 detours
Total 0 displayed, Up 0, Down 0

Transit LSP: 9 sessions, 1 detours
Total 0 displayed, Up 0, Down 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify the RSVP interface which is enabled for link-protection.

```
[exit]
lab@R1# run show rsvp interface ge-0/0/4.118 extensive
ge-0/0/4.118 Index 74, State Ena/Up
  Authentication, NoAggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 172.30.0.9
  ActiveResv 11, PreemptionCnt 4, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 333Mbps
  ct0: StaticBW 333Mbps, AvailableBW 33Mbps
    MaxAvailableBW 333Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 240Mbps[7]
60Mbps
  Protection: On, Bypass: 1, LSP: 5, Protected LSP: 5, Unprotected LSP: 0
      1 Sep 22 15:25:21 New bypass Bypass->172.30.0.10
    Bypass: Bypass->172.30.0.10, State: Up, Type: LP, LSP: 5, Backup: 0
      3 Sep 22 15:25:23 Record Route:  172.30.0.2 172.30.0.18 172.30.0.46
      2 Sep 22 15:25:23 Up
      1 Sep 22 15:25:22 CSPF: computation result accepted
```

The interface shown is the connection from R1 to R8. A bypass LSP is computed for link protection through R2 (172.30.0.2), R7 (172.30.0.18) and terminates on R8 (172.30.0.46). 5 LSPs are protected by this Bypass LSP (1:n).

The bypass LSPs from other LSRs are listed as Transit and Egress. Verify both tables on R3.

Note: The output below is taken after node-link-protection is enabled as well as link-protection.

```
[edit]
lab@R3# run show mpls lsp
Ingress LSP: 3 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.6      172.30.5.3      Up     0 *     primary-2        R3-to-R6
172.30.5.8      172.30.5.3      Up     0 *     path-1           R3-to-R8-1
172.30.5.8      172.30.5.3      Up     0 *     primary-1        R3-to-R8-2
Total 3 displayed, Up 3, Down 0

Egress LSP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.3      172.30.5.6      Up       0  1 SE       3         - R6-to-R3
172.30.5.3      172.30.5.4      Up       0  1 SE       3         - Bypass->172.30.0.21
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.3  172.30.5.8   Up       0  1 SE       3        - R8-to-R3-1
172.30.5.3  172.30.5.8   Up       0  1 SE       3        - R8-to-R3-2
172.30.5.3  172.30.5.2   Up       0  1 SE       3        - Bypass->172.30.0.14
172.30.5.3  172.30.5.1   Up       0  1 SE       3        - Bypass->172.30.0.2-
>172.30.0.14
Total 6 displayed, Up 6, Down 0

Transit LSP: 17 sessions, 1 detours
To          From         State   Rt Style Labelin Labelout LSPname
172.30.5.1  172.30.5.6   Up       0  1 SE  302336   302576 R6-to-R1
172.30.5.1  172.30.5.6   Up       0  1 SE  302352   302592 R6-to-R1
172.30.5.1  172.30.5.2   Up       0  1 SE  301456   301456 Bypass->172.30.0.1
172.30.5.2  172.30.5.4   Up       0  1 SE  302304        3 Bypass->172.30.0.5-
>172.30.0.2
172.30.5.2  172.30.5.5   Up       0  1 SE  302320        3 R5-to-R2
172.30.5.2  172.30.5.5   Up       0  1 SE  302272        3 R5-to-R2
172.30.5.2  172.30.5.1   Up       0  1 SE  301472        3 Bypass->172.30.0.2
172.30.5.4  172.30.5.7   Up       0  1 SE  302480        3 R7-to-R4-2
172.30.5.4  172.30.5.2   Up       0  1 SE  300736        3 Bypass->172.30.0.1-
>172.30.0.6
172.30.5.4  172.30.5.1   Up       0  1 SE  301920        3 Bypass->172.30.0.6
172.30.5.5  172.30.5.4   Up       0  1 SE  302224   302320 Bypass->172.30.0.30
172.30.5.5  172.30.5.2   Up       0  1 SE  302432   302608 R2-to-R5
172.30.5.5  172.30.5.2   Up       0  1 SE  302288   302560 R2-to-R5
172.30.5.6  172.30.5.5   Up       0  1 SE  302416        3 Bypass->172.30.0.34
172.30.5.6  172.30.5.1   Up       0  1 SE  302240        3 R1-to-R6
172.30.5.6  172.30.5.1   Up       0  1 SE  302400        3 R1-to-R6
172.30.5.7  172.30.5.4   Up       0  1 SE  302448   302672 R4-to-R7-1
Total 17 displayed, Up 17, Down 0
```

Let's list the RSVP sessions.

```
[edit]
lab@R3# run show rsvp session
Ingress RSVP: 8 sessions
To          From         State   Rt Style Labelin Labelout LSPname
172.30.5.2  172.30.5.3   Up       0  1 SE       -   300384 Bypass->172.30.0.13
172.30.5.4  172.30.5.3   Up       0  1 SE       -   299856 Bypass->172.30.0.22
172.30.5.5  172.30.5.3   Up       0  1 SE       -   300096 Bypass->172.30.0.26-
>172.30.0.33
172.30.5.6  172.30.5.3   Up       0  1 SE       -   299872 R3-to-R6
172.30.5.6  172.30.5.3   Up       0  1 SE       -   300080 Bypass->172.30.0.26
172.30.5.8  172.30.5.3   Up       0  1 SE       -   300464 R3-to-R8-2
172.30.5.8  172.30.5.3   Up       0  1 SE       -   300352 R3-to-R8-1
Total 7 displayed, Up 7, Down 0

Egress RSVP: 8 sessions
To          From         State   Rt Style Labelin Labelout LSPname
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.3  172.30.5.2   Up       0  1 SE       3        - Bypass->172.30.0.14
172.30.5.3  172.30.5.1   Up       0  1 SE       3        - Bypass->172.30.0.2-
>172.30.0.14
172.30.5.3  172.30.5.4   Up       0  1 SE       3        - Bypass->172.30.0.21
172.30.5.3  172.30.5.8   Up       0  1 SE       3        - R8-to-R3-1
172.30.5.3  172.30.5.8   Up       0  1 SE       3        - R8-to-R3-1
172.30.5.3  172.30.5.8   Up       0  1 SE       3        - R8-to-R3-2
172.30.5.3  172.30.5.6   Up       0  1 SE       3        - R6-to-R3
172.30.5.3  172.30.5.6   Up       0  1 SE       3        - Bypass->172.30.0.25
Total 8 displayed, Up 8, Down 0

Transit RSVP: 30 sessions, 4 detours
To          From         State   Rt Style Labelin Labelout LSPname
172.30.5.1  172.30.5.2   Up       0  1 SE  300528   300320 Bypass->172.30.0.1
172.30.5.1  172.30.5.4   Up       0  1 SE  300688   300128 Bypass->172.30.0.5
172.30.5.1  172.30.5.8   Up       0  1 SE  300672   300528 R8-to-R1
172.30.5.1  172.30.5.6   Up       0  1 SE  300384   300016 R6-to-R1
172.30.5.1  172.30.5.6   Up       0  1 SE  300416   300032 R6-to-R1
172.30.5.2  172.30.5.1   Up       0  1 SE  300064        3 Bypass->172.30.0.2
172.30.5.2  172.30.5.5   Up       0  1 SE  300480        3 R5-to-R2
172.30.5.2  172.30.5.5   Up       0  1 SE  300656        3 R5-to-R2
172.30.5.2  172.30.5.7   Up       0  1 SE  300544        3 R7-to-R2
172.30.5.5  172.30.5.2   Up       0  1 SE  300640   300320
172.30.5.5:172.30.5.2:32767:mvpn:C2-spoke
172.30.5.4  172.30.5.2   Up       0  1 SE  300080        3 Bypass->172.30.0.1-
>172.30.0.6
172.30.5.4  172.30.5.1   Up       0  1 SE  300048        3 Bypass->172.30.0.6
172.30.5.4  172.30.5.5   Up       0  1 SE  300448        3 R5-to-R4
172.30.5.4  172.30.5.5   Up       0  1 SE  300368        3 Bypass->172.30.0.29
172.30.5.4  172.30.5.7   Up       0  1 SE  300752        3 R7-to-R4-2
172.30.5.5  172.30.5.2   Up       0  1 SE  300096   299824 R2-to-R5
172.30.5.5  172.30.5.2   Up       0  1 SE  300112   299840 R2-to-R5
172.30.5.5  172.30.5.4   Up       0  1 SE  300320   299952 R4-to-R5
172.30.5.5  172.30.5.4   Up       0  1 SE  300160   299856 Bypass->172.30.0.30
172.30.5.5  172.30.5.6   Up       0  1 SE  300352   300048 Bypass->172.30.0.33
172.30.5.6  172.30.5.1   Up       0  1 SE  300288        3 R1-to-R6
172.30.5.6  172.30.5.1   Up       0  1 SE  300304        3 R1-to-R6
172.30.5.6  172.30.5.5   Up       0  1 SE  300176        3 Bypass->172.30.0.34
172.30.5.6  172.30.5.7   Up       0  1 SE  300576        3 Bypass->172.30.0.41
172.30.5.7  172.30.5.2   Up       0  1 SE  300192   299920 R2-to-R7
172.30.5.7  172.30.5.2   Up       0  1 SE  300512   300224 Bypass->172.30.0.18
172.30.5.7  172.30.5.4   Up       0  1 SE  300768   300176 R4-to-R7-1
172.30.5.7  172.30.5.6   Up       0  1 SE  300224   299936 Bypass->172.30.0.42
172.30.5.8  172.30.5.1   Up       0  1 SE  300720   300416 R1-to-R8
172.30.5.8  172.30.5.1   Up       0  1 SE  300704   300400 R1-to-R8
Total 30 displayed, Up 30, Down 0
```

Configure node link protection for LSPs I, J, M and N.

© 2020 Juniper Networks, Inc. All rights reserved.

Node link protection automatically protects the link but also the downstream node.

Configure node link protection on the required devices.

a) R3

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 node-link-protection

[edit]
lab@R3# commit
commit complete
```

b) R4

```
[edit]
lab@R4# set protocols mpls label-switched-path R4-to-R7-1 node-link-protection

[edit]
lab@R4# commit
commit complete
```

c) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R4-1 node-link-protection

[edit]
lab@R7# commit
commit complete
```

d) R8

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3-1 node-link-protection

[edit]
lab@R8# commit
commit complete
```

Verify node-link protection configured on the LSP.

```
[edit]
lab@R3# run show mpls lsp name R3-to-R8-1 extensive
Ingress LSP: 3 sessions
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.8
  From: 172.30.5.3, State: Up, ActiveRoute: 0, LSPname: R3-to-R8-1
  ActivePath: path-1 (primary)
  Node/Link protection desired
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path-1              State: Up, No-decrement-ttl
    Priorities: 6 6
    Bandwidth: 60Mbps
    OptimizeTimer: 28800
    SmartOptimizeTimer: 180
    Reoptimization in 9415 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 172.30.0.13 S 172.30.0.1 S 172.30.0.10 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.5.2(flag=0x29) 172.30.0.13(flag=9 Label=300144)
172.30.5.1(flag=0x21) 172.30.0.1(flag=1 Label=300192) 172.30.5.8(flag=0x20)
172.30.0.10(Label=3)
    OAM state : BFD session up  LSP-ping up
   10 Nov  1 04:25:15.245 Link-protection Up
    9 Nov  1 04:25:12.239 Selected as active path: due to 'primary'
    8 Nov  1 04:24:14.086 Record Route:  172.30.5.2(flag=0x29) 172.30.0.13(flag=9
Label=300144) 172.30.5.1(flag=0x21) 172.30.0.1(flag=1 Label=300192)
172.30.5.8(flag=0x20) 172.30.0.10(Label=3)
    7 Nov  1 04:24:14.068 Record Route:  172.30.5.2(flag=0x29) 172.30.0.13(flag=9
Label=300144) 172.30.5.1(flag=0x20) 172.30.0.1(Label=300192)
172.30.5.8(flag=0x20) 172.30.0.10(Label=3)
    6 Nov  1 04:24:11.134 Record Route:  172.30.5.2(flag=0x20)
172.30.0.13(Label=300144) 172.30.5.1(flag=0x20) 172.30.0.1(Label=300192)
172.30.5.8(flag=0x20) 172.30.0.10(Label=3)
    5 Nov  1 04:24:11.133 Up
    4 Nov  1 04:24:11.054 Originate Call
    3 Nov  1 04:24:11.054 CSPF: computation result accepted  172.30.0.13
172.30.0.1 172.30.0.10
    2 Nov  1 04:23:41.973 CSPF failed: no route toward 172.30.5.8[8 times]
    1 Nov  1 04:19:47.662 CSPF failed: no route toward 172.30.5.1
  Secondary secondary-1      State: Up, No-decrement-ttl
    Priorities: 6 6
    Bandwidth: 60Mbps
    OptimizeTimer: 28800
    SmartOptimizeTimer: 180
    Reoptimization in 28484 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 16777229)
 172.30.0.26 S 172.30.0.33 S 172.30.0.38 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
          172.30.5.6(flag=0x29) 172.30.0.26(flag=9 Label=300352)
172.30.5.5(flag=0x21) 172.30.0.33(flag=1 Label=300176) 172.30.5.8(flag=0x20)
172.30.0.38(Label=3)
    OAM state : BFD session up  LSP-ping up
   14 Nov  1 04:25:12.241 Link-protection Up
   13 Nov  1 04:25:12.239 Deselected as active: due to 'primary'
   12 Nov  1 04:25:12.239 Link-protection Down
   11 Nov  1 04:24:32.628 Record Route:  172.30.5.6(flag=0x29) 172.30.0.26(flag=9
Label=300352) 172.30.5.5(flag=0x21) 172.30.0.33(flag=1 Label=300176)
172.30.5.8(flag=0x20) 172.30.0.38(Label=3)
   10 Nov  1 04:24:03.831 Link-protection Up
    9 Nov  1 04:24:03.829 Link-protection Down
    8 Nov  1 04:23:50.976 Link-protection Up
    7 Nov  1 04:23:44.992 Record Route:  172.30.5.6(flag=0x21) 172.30.0.26(flag=1
Label=300352) 172.30.5.5(flag=0x20) 172.30.0.33(Label=300176)
172.30.5.8(flag=0x20) 172.30.0.38(Label=3)
    6 Nov  1 04:23:42.043 Selected as active path
    5 Nov  1 04:23:42.042 Record Route:  172.30.5.6(flag=0x20)
172.30.0.26(Label=300352) 172.30.5.5(flag=0x20) 172.30.0.33(Label=300176)
172.30.5.8(flag=0x20) 172.30.0.38(Label=3)
    4 Nov  1 04:23:42.042 Up
    3 Nov  1 04:23:41.974 Originate Call
    2 Nov  1 04:23:41.974 CSPF: computation result accepted  172.30.0.26
172.30.0.33 172.30.0.38
    1 Nov  1 04:23:12.604 CSPF failed: no route toward 172.30.5.8[7 times]
  Created: Fri Nov  1 04:18:09 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 7 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 28 sessions, 4 detours
Total 0 displayed, Up 0, Down 0
```

7) Configure per flow load balancing policy on rest of the routers.

To minimize further packet loss during switchover you can apply a per flow load balancing policy. This will place the detour next-hop in the forwarding table.

Note: This is just an example and does not need to be carried out since we already have a load balancing policy applied from earlier in this chapter.

```
[edit]
lab@R1# set policy-options policy-statement load-balance term 1 then load-balance per-
packet

[edit]
lab@R1# set routing-options forwarding-table export load-balance
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

The load balancing policy should already be applied to all routers.

Task 4: IPv6 Tunneling with 6PE

Having a MPLS core build for IPv4 protocols using RSVP and LDP enables a foundation for many services like L3VPN, L2VPN, VPLS and etc. These services are multiplexed over the MPLS core. However, it is also possible for IPv6 traffic to be transported as a service over the MPLS enabled core. Two schemes are available, named 6PE and 6VPE.

The current task requires the configuration of 6PE across the MPLS core. You can avoid enabling IPv6 native in the core and use the IPv4 LSPs that are already present. To copy the LSP routes from `inet.3` to `inet6.3` MPLS must be configured to tunnel IPv6 traffic. This way all routes will be copied as v4 mapped IPv6 address into the `inet6.3` table.

MPLS should know that the tunneled packets inside the LSP are IPv6 and not IPv4. As mentioned earlier the default behavior for the penultimate-hop LSR is to remove the outer label and send the remaining stack/ packet to the egress LSR. In case with IPv6 tunneling, where no native IPv6 is enabled in the core this penultimate router does not know where to send the IPv6 packet to and it will just drop it. 6PE is used to connect IPv6 islands through the core. So, a way to transport the IPv6 routes between the islands should be provided. These IPv6 routes should have next-hops resolved by the LSP in `inet6.3`. Multiprotocol BGP is used to transfer these routes between the PE routers. A new BGP family is used for this purpose. IPv6 tunneled routes carry an additional label. This label is the IPv6 Explicit NULL label (label value 2). This way when a penultimate hop LSR removes the outer label, the IPv6 packet is sent to the egress PE based on this additional inner label. This additional label will be removed by the egress LSR and the IPv6 packet can be send in its native form to the destination.

1) Configure IPv6 MPLS tunneling on all routers.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols mpls ipv6-tunneling

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

This statement instructs Junos to copy the routes from `inet.3` to `inet6.3` as v4 mapped IPv6 routes.

```
[edit]
lab@R1# run show route table inet6.3

inet6.3: 7 destinations, 9 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::ffff:172.30.5.2/128
                   *[LDP/9] 00:00:33, metric 5
                    > to 172.30.0.2 via ae0.0, Push 0
::ffff:172.30.5.3/128
                   *[LDP/9] 00:00:27, metric 15
                    > to 172.30.0.2 via ae0.0, Push 300832
::ffff:172.30.5.4/128
                   *[LDP/9] 00:00:34, metric 10
                    > to 172.30.0.6 via ge-0/0/4.114, Push 0
::ffff:172.30.5.5/128
                   *[LDP/9] 00:00:03, metric 20
                    > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R8
                      to 172.30.0.6 via ge-0/0/4.114, label-switched-path R1-to-
R8
::ffff:172.30.5.6/128
                   *[RSVP/7/1] 00:00:33, metric 25
                    > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R6
                      to 172.30.0.6 via ge-0/0/4.114, label-switched-path R1-to-
R6
                      to 172.30.0.10 via ge-0/0/4.118, label-switched-path R1-to-
R6
                      to 172.30.0.10 via ge-0/0/4.118, label-switched-path R1-to-
R6
                      to 172.30.0.2 via ae0.0, label-switched-path R1-to-R6
                    [LDP/9] 00:00:09, metric 25
                    > to 172.30.0.2 via ae0.0, Push 300848
                      to 172.30.0.6 via ge-0/0/4.114, Push 300336
::ffff:172.30.5.7/128
                   *[LDP/9] 00:00:09, metric 15
                    > to 172.30.0.2 via ae0.0, Push 300816
::ffff:172.30.5.8/128
                   *[RSVP/7/1] 00:00:12, metric 10
                    > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R8
                      to 172.30.0.6 via ge-0/0/4.114, label-switched-path Bypass-
>172.30.0.2
                    [LDP/9] 00:00:03, metric 10
                    > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R8
                      to 172.30.0.6 via ge-0/0/4.114, label-switched-path R1-to-
R8
```


© 2020 Juniper Networks, Inc. All rights reserved.

IPv6 tunneling should be configured on all PE routers.

2) Configure IPv6 labeled-uncast family on all routers.

IPv6 labeled unicast should be enabled on all PE routers. This enables the exchange of IPv6 labeled routes.

Repeat these commands on all eight devices.

```
[edit]
lab@R1# set protocols bgp group ibgp family inet6 labeled-unicast explicit-null
```

Note: The `explicit-null` option is mandatory here. Without it, Junos will not commit the configuration.

When you add another family to BGP the family unicast (which is enabled by default) needs to be explicitly configured also.

```
[edit]
lab@R1# set protocols bgp group ibgp family inet unicast

[edit]
lab@R1# commit
commit complete
```

3) Configure IPv6 labeled unicast on the route reflector.

The IPv6 labeled unicast family should also be configured on the RR as it should reflect the routes to the other PEs.

RR:

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family inet6 labeled-unicast
explicit-null

[edit]
lab@route-reflector# set protocols bgp group cluster-1 family inet unicast

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family inet6 labeled-unicast
explicit-null

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family inet unicast
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# commit
commit complete
```

RR should start to receive routes in table `inet6.0`

```
[edit]
lab@route-reflector# run show route table inet6

inet6.0: 33 destinations, 49 routes (1 active, 0 holddown, 48 hidden)
+ = Active Route, - = Last Active, * = Both

ff02::2/128             *[INET6/0] 1d 05:07:12
                           MultiRecv
```

The routes are present, but they are hidden. The RR needs to resolve the routes so that they can be reflected to the other PEs. You cannot use MPLS IPv6 tunneling to move the routes from `inet.3` to `inet6.3`, because the RR is not MPLS enabled. One solution is to add a static default route in the `inet6.3` table.

The hidden routes are shown below.

```
[edit]
lab@route-reflector# run show route table inet6 hidden

inet6.0: 32 destinations, 48 routes (1  active, 0 holddown, 48 hidden)
+ = Active Route, - = Last Active, * = Both

fd01:aaaa:bbbb::/80 [BGP/170] 00:00:05, localpref 100, from 172.30.5.8
                      AS path: 23456 I
                      Unusable
                    [BGP/170] 00:00:09, localpref 100, from 172.30.5.7
                      AS path: 23456 I
                      Unusable
fd01:aaaa:bbbb:0:1::/80
                    [BGP/170] 00:00:05, localpref 100, from 172.30.5.8
                      AS path: 23456 I
                      Unusable
                    [BGP/170] 00:00:09, localpref 100, from 172.30.5.7
                      AS path: 23456 I
                      Unusable
fd01:aaaa:bbbb:0:2::/80
                    [BGP/170] 00:00:05, localpref 100, from 172.30.5.8
                      AS path: 23456 I
                      Unusable
                    [BGP/170] 00:00:09, localpref 100, from 172.30.5.7
                      AS path: 23456 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                      Unusable
fd01:aaaa:bbbb:0:3::/80
                    [BGP/170] 00:00:05, localpref 100, from 172.30.5.8
                      AS path: 23456 I
                      Unusable
                    [BGP/170] 00:00:09, localpref 100, from 172.30.5.7
                      AS path: 23456 I
                      Unusable
...
```

As you can see the routes are present, but they are unusable.

```
[edit]
lab@route-reflector# set routing-options rib inet6.3 static route 0:0:0:0:0:ffff::/96
receive

[edit]
lab@route-reflector# commit
[edit protocols bgp group cluster-1 cluster]

commit complete
```

Verify the content of the `inet6.0` table, after adding the static route to table `inet6.3`.

```
[edit]
lab@route-reflector# run show route table inet6

inet6.0: 49 destinations, 65 routes (49 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::ffff:0.0.0.0/96  *[Static/5] 00:00:11
                       Receive
fd01:aaaa:bbbb::/80*[BGP/170] 00:18:02, localpref 100, from 172.30.5.7
                      AS path: 23456 I
                      Indirect
                    [BGP/170] 00:17:58, localpref 100, from 172.30.5.8
                      AS path: 23456 I
                      Indirect
fd01:aaaa:bbbb:0:1::/80
                   *[BGP/170] 00:18:02, localpref 100, from 172.30.5.7
                      AS path: 23456 I
                      Indirect
                    [BGP/170] 00:17:58, localpref 100, from 172.30.5.8
                      AS path: 23456 I
                      Indirect
fd01:aaaa:bbbb:0:2::/80
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    *[BGP/170] 00:18:02, localpref 100, from 172.30.5.7
                      AS path: 23456 I
                      Indirect
...
```

```
[edit]
lab@route-reflector# run show route table inet6.0 fd01:aaaa:bbbb::/80 extensive

inet6.0: 49 destinations, 65 routes (49 active, 0 holddown, 0 hidden)
fd01:aaaa:bbbb::/80 (2 entries, 1 announced)
TSI:
KRT in-kernel fd01:aaaa:bbbb::/80 -> {indirect(262146)}
Page 0 idx 0 Type 1 val 8a639d8
    Nexthop: ::ffff:172.30.5.7
    Localpref: 100
    AS path: [54591] 23456 I
    Unrecognized Attributes: 9 bytes
    Attr flags e0 code 11: 02 01 06 8f 30 c3
    Communities:
    Cluster ID: 0.0.0.1
    Originator ID: 172.30.5.7
    Advertise: 0000000b
Page 0 idx 1 Type 1 val 8a63de0
    Nexthop: ::ffff:172.30.5.7
    Localpref: 100
    AS path: [54591] 23456 I
    Unrecognized Attributes: 9 bytes
    Attr flags e0 code 11: 02 01 06 8f 30 c3
    Communities:
    Cluster ID: 0.0.0.1
    Originator ID: 172.30.5.7
Path fd01:aaaa:bbbb:: from 172.30.5.7 Vector len 4.  Val: 0 1
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 48
                Source: 172.30.5.7
                Next hop type: Receive
                Protocol next hop: ::ffff:172.30.5.7
                Label operation: Push 2
                Indirect next hop: 89ee160 262146
                State: <Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 18:32      Metric2: 0
                Task: BGP_54591.172.30.5.7+179
                Announcement bits (3): 0-KRT 3-BGP RT Background 4-Resolve tree 3
                AS path: 23456 I Unrecognized Attributes: 9 bytes
                AS path:  Attr flags e0 code 11: 02 01 06 8f 30 c3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            Route Label: 2
            Localpref: 100
            Router ID: 172.30.5.7
            Indirect next hops: 1
                    Protocol next hop: ::ffff:172.30.5.7
                    Push 2
                    Indirect next hop: 89ee160 262146
                    Indirect path forwarding next hops: 0
                            Next hop type: Receive
                    ::ffff:0.0.0.0/96 Originating RIB: inet6.3
                      Node path count: 1
                      Forwarding nexthops: 0
                            Next hop type: Receive
    BGP     Preference: 170/-101
            Next hop type: Indirect
            Next-hop reference count: 16
            Source: 172.30.5.8
            Next hop type: Receive
            Protocol next hop: ::ffff:172.30.5.8
            Push 2
            Indirect next hop: 89ee0b0 -
            State: <NotBest Int Ext>
            Inactive reason: Not Best in its group - Router ID
            Local AS: 54591 Peer AS: 54591
            Age: 18:28      Metric2: 0
            Task: BGP_54591.172.30.5.8+179
            AS path: 23456 I Unrecognized Attributes: 9 bytes
...
```

The BGP routes are resolved by the static route we created in the `inet6.3` table. You can see that the routes contain a label and the value of the label is 2 (IPv6 Explicit Null).

Verify routes on R3.

```
[edit]
lab@R3# run show route table inet6.0 extensive

inet6.0: 80 destinations, 82 routes (80 active, 0 holddown, 0 hidden)
fd01:aaaa:bbbb::/80 (1 entry, 1 announced)
TSI:
KRT in-kernel fd01:aaaa:bbbb::/80 -> {indirect(262147)}
Page 0 idx 0 Type 1 val 179c460
    Nexthop: Self
    AS path: [54591] 110047427 I
    Communities:
Path fd01:aaaa:bbbb:: from 172.30.5.41 Vector len 4.  Val: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
*BGP    Preference: 170/-101
        Next hop type: Indirect
        Address: 0x155d7d0
        Next-hop reference count: 48
        Source: 172.30.5.41
        Next hop type: Router, Next hop index: 615
        Next hop: 172.30.0.22 via ge-0/0/4.134 weight 0x1, selected
        Label-switched-path R3-to-R6
        Label operation: Push 2, Push 299776, Push 302368(top)
        Label TTL action: no-prop-ttl, no-prop-ttl, no-prop-ttl(top)
        Protocol next hop: ::ffff:172.30.5.7
        Label operation: Push 2
        Indirect next hop: 17b40fc 262147
        State: <Active Int Ext>
        Local AS: 54591 Peer AS: 54591
        Age: 11:51      Metric2: 20
        Task: BGP_54591.172.30.5.41+3334
        Announcement bits (3): 0-KRT 3-BGP_RT_Background 4-Resolve tree 5
        AS path:
        AS2 PA[1]: AS_TRANS
        AS4 PA[1]: 110047427
        Merged[1]: 110047427 I (Originator)
        Cluster list:  0.0.0.1
        Originator ID: 172.30.5.7
        Accepted
        Route Label: 2
        Localpref: 100
        Router ID: 172.30.5.41
        Indirect next hops: 1
                Protocol next hop: ::ffff:172.30.5.7 Metric: 20
                Push 2
                Indirect next hop: 17b40fc 262147
                Indirect path forwarding next hops: 1
                        Next hop type: Router
```

R3 has routes received from the RR, which were reflected from R7.

4) Configure a policy on R5 to change the next hop of the IPv6 routes.

R5 has a single EBGP session, which carries IPv4 and IPv6 routes. R5 will change the next-hop for the IPv6 routes it receives from the RR before it sends them to C3 with the IPv4 next-hop of the direct interface. This is normal EBGP operation and is expected. However, the router on the other end will try to resolve the next-hop in the `inet6.0` table, which causes the route to be unusable.

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.10 table inet6.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet6.0: 87 destinations, 107 routes (87 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop   MED  Lclpref    AS path
* fd01:aaaa:bbbb::/80     Self                                 110047427 I
  fd01:aaaa:bbbb:0:1::/80
*                         Self                                 110047427 I
  fd01:aaaa:bbbb:0:2::/80
*                         Self                                 110047427 I
  fd01:aaaa:bbbb:0:3::/80
*                         Self                                 110047427 I
  fd01:aaaa:bbbb:0:4::/80
*                         Self                                 110047427 I
  fd01:aaaa:bbbb:0:5::/80
*                         Self                                 110047427 I
  fd01:aaaa:bbbb:0:6::/80
*                         Self                                 110047427 I
...
```

R3 is advertising the IPv6 routes reflected by RR with next-hop Self.

You have to create an export policy to change the next-hop with an IPv4-compatible IPv6 address. You have to change the next-hop attribute to the IPv6 address of the interface towards C3.

Configure an export policy for both EBGP sessions.

```
[edit]
lab@R5# set policy-options policy-statement ipv6-next-hop-1 term 1 from family inet6

[edit]
lab@R5# set policy-options policy-statement ipv6-next-hop-1 term 1 from protocol bgp

[edit]
lab@R5# set policy-options policy-statement ipv6-next-hop-1 term 1 then next-hop
::192.168.0.9

[edit]
lab@R5# set policy-options policy-statement ipv6-next-hop-2 term 1 from family inet6

[edit]
lab@R5# set policy-options policy-statement ipv6-next-hop-2 term 1 from protocol bgp

[edit]
lab@R5# set policy-options policy-statement ipv6-next-hop-2 term 1 then next-hop
::192.168.0.13

[edit]
lab@R5# set protocols bgp group C3 neighbor 192.168.0.10 export ipv6-next-hop-1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocol bgp group C3 neighbor 192.168.0.14 export ipv6-next-hop-2

[edit]
lab@R5# commit
commit complete
```

Verify the IPv6 routes advertised from R5.

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.10 table inet6.0

inet6.0: 87 destinations, 107 routes (87 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* fd01:aaaa:bbbb::/80     ::192.168.0.9                           110047427 I
  fd01:aaaa:bbbb:0:1::/80
*                         ::192.168.0.9                           110047427 I
  fd01:aaaa:bbbb:0:2::/80
*                         ::192.168.0.9                           110047427 I
  fd01:aaaa:bbbb:0:3::/80
*                         ::192.168.0.9                           110047427 I
  fd01:aaaa:bbbb:0:4::/80
*                         ::192.168.0.9                           110047427 I
  fd01:aaaa:bbbb:0:5::/80
*                         ::192.168.0.9                           110047427 I
...
```

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.14 table inet6.0

inet6.0: 87 destinations, 107 routes (87 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* fd01:aaaa:bbbb::/80     ::192.168.0.13                          110047427 I
  fd01:aaaa:bbbb:0:1::/80
*                         ::192.168.0.13                          110047427 I
  fd01:aaaa:bbbb:0:2::/80
*                         ::192.168.0.13                          110047427 I
  fd01:aaaa:bbbb:0:3::/80
*                         ::192.168.0.13                          110047427 I
  fd01:aaaa:bbbb:0:4::/80
*                         ::192.168.0.13                          110047427 I
...
```

The IPv6 routes are advertised with different next-hops and they are successfully resolved. This provides reachability for the IPv6 islands over the MPLS core network.

© 2020 Juniper Networks, Inc. All rights reserved.

Appendix - Chapter 5: L3VPN Configuration

Tasks in this chapter will cover the implementation of L3VPNs. Different L3VPNs will be configured with customers running OSPF or BGP as the PE-CE protocol, dual-homed sites, customer internet access, multicast VPNs and IPv6 VPNs (6VPE).

Task 1: L3VPN Configuration

L3VPN Overview

Junos has a rich feature set to provide L3VPN services. With Layer 3 VPNs (L3VPN) services the customers Layer 3 routing is handled by the Service Provider (SP). The customers routing protocol is terminated on the provider router into VPN-specific routing table(s), where the customer prefixes are distributed between customer remote sites.

L3VPN customer sites are called customer edges (CEs) and are connected to provider edge (PE) devices. PE routers can multiplex multiple customers connections over the MPLS network. CE devices are typically not MPLS enabled (however this is possible). Customers can have many customer sites located behind different PE routers. Each customer site is attached to at least one PE router. The PE router maintains a separate VPN routing and forwarding (VRF) table for each customer. This separation provides the capability for customers to use their own (overlapping) IP addressing scheme independently from other customers.

PE routers extend the IPv4 customer addresses with an additional 64-bits route distinguisher (RD). By adding this distinguisher to the IPv4 address, a new address is created which is called a VPNv4 address. To transfer these uniquely identifiable addresses through the network another Multiprotocol BGP address family called `inet-vpn` must be enabled on all PE routers. Multiprotocol BGP carries the VPN label together with VPN prefixes.

Each Multiprotocol BGP route is advertised with an extended `route-target` community. `route-target` communities are usually unique for each VRF; however, they can be the same. Based on these communities each PE router knows if the VPNv4 route should be accepted and to which VRF it must be imported. The attachment of multiple VRFs creates the actual L3VPN.

1) Configure additional interfaces for CE connectivity. These interfaces are shown in Figure 11 and described in Table 12.

© 2020 Juniper Networks, Inc. All rights reserved.

Router	Interface Name	IP Address	IPv6 Address
R1	ge-0/0/5.311	192.168.0.41/30	
	ge-0/0/5.312	192.168.0.45/30	
	ge-0/0/5.313	192.168.0.49/30	
	lo0.1	172.30.5.9/32	
	lo0.2	172.30.5.10/32	
R2	ge-0/0/5.314	192.168.0.53/30	
	ge-0/0/5.315	192.168.0.57/30	
	ge-0/0/5.316	192.168.0.61/30	
	lo0.1	172.30.5.13/32	
	lo0.2	172.30.5.14/32	
R3	ge-0/0/5.317		fc09:c0:ffee::9/126
	ge-0/0/5.318	192.168.0.69/30	
	lo0.1	172.30.5.17/32	
	lo0.2	172.30.5.18/32	fd17:f0f4:f691:5::12/128
R4	ge-0/0/5.319	192.168.0.73/30	
	ge-0/0/5.320	192.168.0.77/30	
	lo0.1	172.30.5.21/32	
	lo0.2	172.30.5.22/32	
R5	ge-0/0/5.321	192.168.0.81/30	
	lo0.1	172.30.5.25/32	
R6	ge-0/0/5.322	192.168.0.85/30	
	lo0.1	172.30.5.29/32	
R7	ge-0/0/5.323	192.168.0.89/30	
	lo0.1	172.30.5.33/32	
R8	ge-0/0/5.324	192.168.0.93/30	
	ge-0/0/5.325		fc09:c0:ffee::d/126
	lo0.1	172.30.5.37/32	
	lo0.2	172.30.5.38/32	fd17:f0f4:f691:5::26/128

Table 12

a) R1

```
[edit]
lab@R1# set interfaces ge-0/0/5.311 vlan-id 311

[edit]
lab@R1# set interfaces ge-0/0/5.311 family inet address 192.168.0.41/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set interfaces ge-0/0/5.311 description "CE2-1 connection 1"

[edit]
lab@R1# set interfaces ge-0/0/5.312 vlan-id 312

[edit]
lab@R1# set interfaces ge-0/0/5.312 family inet address 192.168.0.45/30

[edit]
lab@R1# set interfaces ge-0/0/5.312 description "CE2-1 connection 2"

[edit]
lab@R1# set interfaces ge-0/0/5.313 vlan-id 313

[edit]
lab@R1# set interfaces ge-0/0/5.313 family inet address 192.168.0.49/30

[edit]
lab@R1# set interfaces ge-0/0/5.313 description "CE2-1 connection 3"

[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set interfaces ge-0/0/5.314 vlan-id 314

[edit]
lab@R2# set interfaces ge-0/0/5.314 family inet address 192.168.0.53/30

[edit]
lab@R2# set interfaces ge-0/0/5.314 description "CE2-2 connection 1"

[edit]
lab@R2# set interfaces ge-0/0/5.315 vlan-id 315

[edit]
lab@R2# set interfaces ge-0/0/5.315 family inet address 192.168.0.57/30

[edit]
lab@R2# set interfaces ge-0/0/5.315 description "CE2-2 connection 2"

[edit]
lab@R2# set interfaces ge-0/0/5.316 vlan-id 316

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# set interfaces ge-0/0/5.316 family inet address 192.168.0.61/30

[edit]
lab@R2# set interfaces ge-0/0/5.316 description "CE2-2 connection 3"

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set interfaces ge-0/0/5.317 vlan-id 317

[edit]
lab@R3# set interfaces ge-0/0/5.317 family inet6 address fc09:c0:ffee::9/126

[edit]
lab@R3# set interfaces ge-0/0/5.317 description "CE3-2 connection"

[edit]
lab@R3# set interfaces ge-0/0/5.318 vlan-id 318

[edit]
lab@R3# set interfaces ge-0/0/5.318 family inet address 192.168.0.69/30

[edit]
lab@R3# set interfaces ge-0/0/5.318 description "CE1-2 connection"

[edit]
lab@R3# commit
commit complete
```

d) R4

```
[edit]
lab@R4# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R4# set interfaces ge-0/0/5.319 vlan-id 319

[edit]
lab@R4# set interfaces ge-0/0/5.319 family inet address 192.168.0.73/30

[edit]
lab@R4# set interfaces ge-0/0/5.319 description "CE1-3 connection"
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set interfaces ge-0/0/5.320 vlan-id 320

[edit]
lab@R4# set interfaces ge-0/0/5.320 family inet address 192.168.0.77/30

[edit]
lab@R4# set interfaces ge-0/0/5.320 description "CE2-3 connection"

[edit]
lab@R4# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set interfaces ge-0/0/5.321 vlan-id 321

[edit]
lab@R5# set interfaces ge-0/0/5.321 family inet address 192.168.0.81/30

[edit]
lab@R5# set interfaces ge-0/0/5.321 description "CE2-4 connection"

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set interfaces ge-0/0/5.322 vlan-id 322

[edit]
lab@R6# set interfaces ge-0/0/5.322 family inet address 192.168.0.85/30

[edit]
lab@R6# set interfaces ge-0/0/5.322 description "CE1-4 connection"

[edit]
lab@R6# commit
commit complete
```

g) R7

```
[edit]
lab@R7# set interfaces ge-0/0/5.323 vlan-id 323
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set interfaces ge-0/0/5.323 family inet address 192.168.0.89/30

[edit]
lab@R7# set interfaces ge-0/0/5.323 description "CE2-5 connection"

[edit]
lab@R7# commit
commit complete
```

h) R8

```
[edit]
lab@R8# set interfaces ge-0/0/5.324 vlan-id 324

[edit]
lab@R8# set interfaces ge-0/0/5.324 family inet address 192.168.0.93/30

[edit]
lab@R8# set interfaces ge-0/0/5.324 description "CE1-1 connection"

[edit]
lab@R8# set interfaces ge-0/0/5.325 vlan-id 325

[edit]
lab@R8# set interfaces ge-0/0/5.325 family inet6 address fc09:c0:ffee::d/126

[edit]
lab@R8# set interfaces ge-0/0/5.325 description "CE3-1 connection"

[edit]
lab@R8# commit
commit complete
```

Note: Junos allows only one unit of the loopback interface to participate in a particular routing instance. Therefore, the additional units on the loopback interface will be configured when the routing instances are configured, and these new units are mapped to them.

2) Configure the BGP VPNv4 address family on all routers.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols bgp group ibgp family inet-vpn unicast
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

3) Configure the VPNv4 address family on the route reflector. Make sure you also include the `inet` NLRI to allow the underlying BGP sessions to continue to share routes.

RR

```
[edit]
lab@route-reflector# set protocols bgp group  cluster-1 family inet-vpn unicast

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family inet unicast

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family inet-vpn unicast

[edit]
lab@route-reflector# set protocols bgp group cluster-1 family inet unicast

[edit]
lab@route-reflector# commit
commit complete
```

Use Table 13 to identify PE-CE routing protocol details.

Customer	Site	Router	PE-CE Protocol	Protocol details
C1	S1	CE1-1	OSPF	Area 0
	S2	CE1-2	OSPF	Area 0
		CE1-3	OSPF	Area 0
	S3	CE1-4	OSPF	Area 0
C2	S1	CE2-1	BGP	AS 64600
		CE2-2	BGP	AS 64600
	S2	CE2-3	BGP	AS 64600
		CE2-4	BGP	AS 64600
	S3	CE2-5	BGP	AS 64600

Table 13

4) Configure customer C1 VPN

Customer C1 has sites attached to R3, R4, R7 and R8.

© 2020 Juniper Networks, Inc. All rights reserved.

a) R3

```
[edit]
lab@R3# set interfaces lo0.1 family inet address 172.30.5.17/32

[edit]
lab@R3# set routing-options route-distinguisher-id 172.30.5.3

[edit]
lab@R3# set routing-instances C1 instance-type vrf

[edit]
lab@R3# set routing-instances C1 interface lo0.1

[edit]
lab@R3# set routing-instances C1 interface ge-0/0/5.318

[edit]
lab@R3# set routing-instances C1 vrf-target target:54591:100

[edit]
lab@R3# set routing-instances C1 protocols ospf area 0 interface all

[edit]
lab@R3# commit
commit complete
```

First a new unit is configured for the loopback IP address used in the routing instance of that customer. Secondly a route distinguisher is configured globally. This feature allows the automatically assignment of route distinguisher values for each VRF. Usually the loopback IP address is used because it is unique in the network.

This way the route distinguisher uses the following format 4:2, where 4 is a 4 bytes address and 2 bytes are used for the locally administered field. In the case with automatic assigned route distinguishers the local administered field is unique for each VRF.

A VRF routing-instance with the appropriate PE-CE interfaces is created. An extended route-target community is added in the format **target:ASN:Administrative number**, where ASN is the autonomous system and the administrative number is a number chosen by the administrator. If configured this way all Multiprotocol BGP routes with that specific route target are imported automatically. Also, all routes in the VRF routing table will be exported to other PEs with that route-target extended community.

OSPF is configured as the PE-CE routing protocol.

Verify the OSPF neighbor table.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show ospf neighbor instance C1
Address          Interface              State     ID               Pri  Dead
192.168.0.70     ge-0/0/5.318           Full      172.31.63.2      128    39
```

Note: Keep in mind that you must specify the instance with the command to see the output for a specific routing-instance.

b) R4

```
[edit]
lab@R4# set interfaces lo0.1 family inet address 172.30.5.21/32

[edit]
lab@R4# set routing-options route-distinguisher-id 172.30.5.4

[edit]
lab@R4# set routing-instances C1 instance-type vrf

[edit]
lab@R4# set routing-instances C1 interface lo0.1

[edit]
lab@R4# set routing-instances C1 interface ge-0/0/5.319

[edit]
lab@R4# set routing-instances C1 vrf-target target:54591:100

[edit]
lab@R4# set routing-instances C1 protocols ospf area 0 interface all

[edit]
lab@R4# commit
commit complete
```

The routing instance on R4 is configured following the same principle as on R3. The import and export route target are the same as on R3. However, the route distinguisher is different from the one configured on R3. It can be the same, but for easy troubleshooting on the JNCIE lab it is better to have it set to a unique value for each PE. This way you will know from where a given route is advertised. To be unique, the route distinguisher should be configured under the `routing-instances` hierarchy.

Tip: *Please keep in mind that the route-distinguisher value is used in BGP to distinguish the VPN prefixes. By default, BGP always selects the best route. If for example you want to load balance to a multihomed L3VPN site, then it's recommended to use different route-distinguisher values on the PEs. This way BGP route reflectors see the VPN prefixes as*

© 2020 Juniper Networks, Inc. All rights reserved.

different routes and will advertise multiple VPNv4 routes for the same prefix to the other PEs (thus enabling the possibility for load balancing). If the RD values are the same, then the route reflectors would only send a single best route to the other PEs.

c) R6

```
[edit]
lab@R6# set interfaces lo0.1 family inet address 172.30.5.29/32

[edit]
lab@R6# set routing-options route-distinguisher-id 172.30.5.6

[edit]
lab@R6# set routing-instances C1 interface lo0.1

[edit]
lab@R6# set routing-instances C1 interface ge-0/0/5.322

[edit]
lab@R6# set routing-instances C1 instance-type vrf

[edit]
lab@R6# set routing-instances C1 vrf-target target:54591:100

[edit]
lab@R6# set routing-instances C1 protocols ospf area 0 interface all

[edit]
lab@R6# commit
commit complete
```

```
[edit]
lab@R6# run show ospf neighbor instance C1
Address          Interface              State     ID               Pri  Dead
192.168.0.86     ge-0/0/5.322           Full      172.31.63.4      128    39
```

d) R8

```
[edit]
lab@R8# set interfaces lo0.1 family inet address 172.30.5.37/32

[edit]
lab@R8# set routing-options route-distinguisher-id 172.30.5.8

[edit]
lab@R8# set routing-instances C1 instance-type vrf
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set routing-instances C1 interface lo0.1

[edit]
lab@R8# set routing-instances C1 interface ge-0/0/5.324

[edit]
lab@R8# set routing-instances C1 vrf-target target:54591:100

[edit]
lab@R8# set routing-instances C1 protocols ospf area 0 interface all

[edit]
lab@R8# commit
commit complete
```

```
[edit]
lab@R8# run show ospf neighbor instance C1
Address          Interface              State     ID               Pri  Dead
192.168.0.94     ge-0/0/5.324           Full      172.31.63.1      128    37
```

5) Configure a static route for `inet.3` table on route reflector.

VPNv4 routes are stored in the `bgp.l3vpn.0` table. These next-hop routes (transport LSPs) should be resolved in the `inet.3` table. However, the route reflector is not MPLS enabled so you can configure a static route the same way as you did in the previous IPv6 6PE tunneling task from previous chapter.

RR

```
[edit]
lab@route-reflector# set routing-options rib inet.3 static route 172.30.5.0/24 receive

[edit]
lab@route-reflector# commit
commit complete
```

Verify if the routes specified in the `bgp.l3vpn.0` table are resolved.

```
[edit]
lab@route-reflector# run show bgp summary
Groups: 2 Peers: 8 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 998        598          0          0          0          0
bgp.l3vpn.0             44         44          0          0          0          0
inet6.0                 64         48          0          0          0          0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Peer                 AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
172.30.5.1       54591        378        153       0       2        2:26 Establ
  inet.0: 383/383/0
  bgp.l3vpn.0: 10/10/0
  inet6.0: 0/0/0
172.30.5.2       54591        376        478       0     381          25 Establ
  inet.0: 0/383/0
  bgp.l3vpn.0: 8/8/0
  inet6.0: 0/0/0
172.30.5.3       54591        707      18088       0      12     4:03:07 Establ
  inet.0: 151/151/0
  bgp.l3vpn.0: 9/9/0
  inet6.0: 16/16/0
172.30.5.4       54591        576      17123       0       0     4:02:12 Establ
  inet.0: 0/0/0
  bgp.l3vpn.0: 9/9/0
  inet6.0: 0/0/0
172.30.5.5       54591        553      17186       0       1     4:08:50 Establ
  inet.0: 16/16/0
  bgp.l3vpn.0: 0/0/0
  inet6.0: 16/16/0
172.30.5.6       54591        570      17179       0       1     4:09:31 Establ
  inet.0: 31/32/0
  bgp.l3vpn.0: 4/4/0
  inet6.0: 0/0/0
172.30.5.7       54591        888      27579       0       1     4:09:19 Establ
  inet.0: 17/17/0
  bgp.l3vpn.0: 0/0/0
  inet6.0: 16/16/0
172.30.5.8       54591        781      21040       0       1     4:09:28 Establ
  inet.0: 0/16/0
  bgp.l3vpn.0: 4/4/0
  inet6.0: 0/16/0
...
```

The VPNv4 routes are active.

6) Configure an OSPF sham link.

By default, the MPLS core is represented as a super backbone to the CE routers running OSPF. Routes imported from the other PEs are BGP routes. They should be exported to the OSPF process. This way the routes will enter the OSPF process as an Inter Area (LSA Type 3) or External (LSA Type 5) type. For this task the requirement is to have a contiguous Area 0 without any ABRs. Customer C1 has a backdoor link, but the path through the MPLS network should be preferred. This means that you have to receive Intra Area (LSA Type 1 and 2) routes from the MPLS PEs.

© 2020 Juniper Networks, Inc. All rights reserved.

Each routing-instance has its own routing table where all routes for a specific VPN are stored.

```
[edit]
lab@R3# run show route table C1.inet.0

C1.inet.0: 20 destinations, 74 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.17/32     *[Direct/0] 00:06:59
                    > via lo0.1
                    [BGP/170] 00:02:25, MED 3, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                    > to 172.30.0.22 via ge-0/0/4.134, Push 306912
                    [BGP/170] 00:01:46, MED 7, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-
R8-2
                      to 172.30.0.22 via ge-0/0/4.134, label-switched-path
Bypass->172.30.0.13->172.30.0.1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13
                    [BGP/170] 00:02:24, MED 8, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                    > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-
R6
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.22
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13
172.30.5.21/32     *[OSPF/10] 00:05:45, metric 3
                    > to 192.168.0.70 via ge-0/0/5.318
                    [BGP/170] 00:02:24, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                    > to 172.30.0.22 via ge-0/0/4.134, Push 306896
                    [BGP/170] 00:02:25, MED 7, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                    > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-
R6
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.22
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13
                    [BGP/170] 00:01:46, MED 8, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-
R8-2
                      to 172.30.0.22 via ge-0/0/4.134, label-switched-path
Bypass->172.30.0.13->172.30.0.1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13
172.30.5.29/32     *[OSPF/10] 00:04:20, metric 8
                    > to 192.168.0.70 via ge-0/0/5.318
                    [BGP/170] 00:02:25, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                    > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-
R6
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.22
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13
                    [BGP/170] 00:02:25, MED 7, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                    > to 172.30.0.22 via ge-0/0/4.134, Push 306912
                    [BGP/170] 00:01:46, MED 13, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-
R8-2
                      to 172.30.0.22 via ge-0/0/4.134, label-switched-path
Bypass->172.30.0.13->172.30.0.1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13
```

Verify the OSPF database on R3.

```
[edit]
lab@R3# run show ospf database instance C1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    OSPF database, Area 0.0.0.0
q Type       ID               Adv Rtr          Seq      Age  Opt  Cksum  Len
Router  *172.30.5.17      172.30.5.17      0x8000000a  1636  0x22 0xfe19  48
Router   172.30.5.21      172.30.5.21      0x8000000e   821  0x22 0x6f90  48
Router   172.31.63.1      172.31.63.1      0x8000000e  1642  0x22 0x768c  72
Router   172.31.63.2      172.31.63.2      0x80000013  1637  0x22 0x6f6  120
Router   172.31.63.3      172.31.63.3      0x80000015   859  0x22 0x497c 120
Router   172.31.63.4      172.31.63.4      0x8000000e  1643  0x22 0xc51a  72
Router   172.31.63.5      172.31.63.5      0x80000009  1436  0x22 0x3554  84
Network  192.168.0.70     172.31.63.2      0x80000002  1637  0x22 0x4d44  32
Network  192.168.0.74     172.31.63.3      0x80000001   859  0x22 0x6325  32
...
```

R3 has an LSA Type 1 received from R4 via the backdoor link. There are also others received over the CE interface.

Verify the OSPF database on R4.

```
[edit]
lab@R4# run show ospf database instance C1

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr          Seq      Age  Opt  Cksum  Len
Router   172.30.5.17      172.30.5.17      0x8000000a  1707  0x22 0xfe19  48
Router  *172.30.5.21      172.30.5.21      0x8000000e   885  0x22 0x6f90  48
Router   172.31.63.1      172.31.63.1      0x8000000e  1711  0x22 0x768c  72
Router   172.31.63.2      172.31.63.2      0x80000013  1706  0x22 0x6f6  120
Router   172.31.63.3      172.31.63.3      0x80000015   926  0x22 0x497c 120
Router   172.31.63.4      172.31.63.4      0x8000000e  1710  0x22 0xc51a  72
Router   172.31.63.5      172.31.63.5      0x80000009  1503  0x22 0x3554  84
Network  192.168.0.70     172.31.63.2      0x80000002  1706  0x22 0x4d44  32
Network  192.168.0.74     172.31.63.3      0x80000001   926  0x22 0x6325  32
...
```

On R4 the database is equal, which is as expected as they are Intra Area routers in the same area.

Verify the OSPF database on R6.

```
[edit]
lab@R6# run show ospf database instance C1

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr          Seq      Age  Opt  Cksum  Len
Router  *172.30.5.29      172.30.5.29      0x80000003    21  0x22 0xd604  48
Router   172.31.63.1      172.31.63.1      0x80000010   510  0x22 0x728e  72
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Router    172.31.63.2      172.31.63.2      0x80000015   505  0x22 0x2f8  120
Router    172.31.63.3      172.31.63.3      0x80000016  2725  0x22 0x477d 120
Router    172.31.63.4      172.31.63.4      0x80000012    61  0x22 0xd542  72
Router    172.31.63.5      172.31.63.5      0x8000000b   302  0x22 0x3156  84
Network   192.168.0.70     172.31.63.2      0x80000004   505  0x22 0x4946  32
Network   192.168.0.74     172.31.63.3      0x80000002  2725  0x22 0x6126  32
...
```

Verify the OSPF database on R8.

```
[edit]
lab@R8# run show ospf database instance C1

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.37      172.30.5.37      0x80000003    16  0x22 0xc7ea  48
Router   172.31.63.1      172.31.63.1      0x80000012    18  0x22 0xf63c  72
Router   172.31.63.2      172.31.63.2      0x80000015   716  0x22 0x2f8  120
Router   172.31.63.3      172.31.63.3      0x80000016  2938  0x22 0x477d 120
Router   172.31.63.4      172.31.63.4      0x80000012   276  0x22 0xd542  72
Router   172.31.63.5      172.31.63.5      0x8000000b   514  0x22 0x3156  84
Network  192.168.0.70     172.31.63.2      0x80000004   716  0x22 0x4946  32
Network  192.168.0.74     172.31.63.3      0x80000002  2938  0x22 0x6126  32
...
```

a) Configure a Sham Link on R3.

```
[edit]
lab@R3# set routing-instances C1 protocols ospf sham-link local 172.30.5.17

[edit]
lab@R3# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.21
metric 100

[edit]
lab@R3# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.29

[edit]
lab@R3# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.37

[edit]
lab@R3# commit
commit complete
```

The local sham link address must be the local loopback interface. The local loopback address should be advertised to the other PE routers. The remote sham link address is the loopback address of the remote PE router. The metric of 100 is added to the sham-link.

© 2020 Juniper Networks, Inc. All rights reserved.

b) Configure the sham link on R4.

```
[edit]
lab@R4# set routing-instances C1 protocols ospf sham-link local 172.30.5.21

[edit]
lab@R4# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.17
metric 100

[edit]
lab@R4# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.29

[edit]
lab@R4# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.37

[edit]
lab@R4# commit
commit complete
```

c) Configure the sham link on R6.

```
[edit]
lab@R6# set routing-instances C1 protocols ospf sham-link local 172.30.5.29

[edit]
lab@R6# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.17

[edit]
lab@R6# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.21

[edit]
lab@R6# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.37

[edit]
lab@R6# commit
commit complete
```

On R6 it is not needed to increase the metric.

d) Configure the sham link on R8.

```
[edit]
lab@R8# set routing-instances C1 protocols ospf sham-link local 172.30.5.37

[edit]
lab@R8# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.17
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.21

[edit]
lab@R8# set routing-instances C1 protocols ospf area 0 sham-link-remote 172.30.5.29

[edit]
lab@R8# commit
commit complete
```

Verify the operation of the sham link.

```
[edit]
lab@R3# run show ospf neighbor instance C1
Address          Interface              State     ID               Pri  Dead
192.168.0.70     ge-0/0/5.318           Full      172.31.63.2      128    34
172.30.5.21      shamlink.0             Full      172.30.5.21        0    35
172.30.5.29      shamlink.1             Full      172.30.5.29        0    33
172.30.5.37      shamlink.2             Full      172.30.5.37        0    36
```

R3 has three new neighbors established on virtual sham link interface.

```
[edit]
lab@R3# run show ospf data instance C1

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.17      172.30.5.17     0x80000012    16  0x22 0xbec6  72
Router   172.30.5.21      172.30.5.21     0x80000019     8  0x22 0x78f9  72
Router   172.30.5.29      172.30.5.29     0x8000000c    13  0x22 0x223d  72
Router   172.30.5.37      172.30.5.37     0x80000009     8  0x22 0xf653  72
Router   172.31.63.1      172.31.63.1     0x80000012   375  0x22 0xf63c  72
Router   172.31.63.2      172.31.63.2     0x80000015  1071  0x22 0x2f8  120
Router   172.31.63.3      172.31.63.3     0x80000017   293  0x22 0x457e 120
Router   172.31.63.4      172.31.63.4     0x80000012   631  0x22 0xd542  72
Router   172.31.63.5      172.31.63.5     0x8000000b   870  0x22 0x3156  84
Network  192.168.0.70     172.31.63.2     0x80000004  1071  0x22 0x4946  32
Network  192.168.0.74     172.31.63.3     0x80000003   293  0x22 0x5f27  32
Network  192.168.0.86     172.31.63.4     0x80000001   631  0x22 0x5f13  32
Network  192.168.0.94     172.31.63.1     0x80000001   375  0x22 0x73f4  32
```

The LSA Type 1 from other PEs are now present in the OSPF database via the sham links.

7) Configure customer C2 VPN.

© 2020 Juniper Networks, Inc. All rights reserved.

Customer C2 has three sites, one of them is the hub router, the others are the spokes. The traffic between the spokes should pass through the hub site. The hub site is connected to both R1 and R2. Site 2 is also multihomed to R4 and R5.

To simulate a logical hub-and-spoke design over an MPLS network involves playing with route-target communities.

At the hub site each of the two customers CE routers are connected to one of the PEs using three different separate physical interfaces. One is used for the traffic to the hub, the other for traffic from the hub and the third for Internet access.

a) Configure C2 VPN on R1.

Configure the new units on the loopback interface and add them to the both VRFs.

```
[edit]
lab@R1# set interfaces lo0.1 family inet address 172.30.5.9/32

[edit]
lab@R1# set interfaces lo0.2 family inet address 172.30.5.10/32
```

Configure the automatic route distinguisher.

```
[edit]
lab@R1# set routing-options route-distinguisher-id 172.30.5.1
```

Next you have to configure the router to accept the AS paths with its own Autonomous system. This is because of the fact that the hub should advertise the routes from one spoke to the other spokes. The routes will be exchanged twice in Multiprotocol BGP. This command must be applied on the route reflector, we will do this after configuring R1.

```
[edit]
lab@R1# set routing-options autonomous-system 54591 loops 3
```

Now you can start with the configuration of the routing-instance which will advertise all routes to the spokes.

```
[edit]
lab@R1# set routing-instances C2-hub instance-type vrf

[edit]
lab@R1# set routing-instances C2-hub interface lo0.1

[edit]
lab@R1# set routing-instances C2-hub interface ge-0/0/5.311
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set routing-instances C2-hub vrf-import C2-hub-import

[edit]
lab@R1# set routing-instances C2-hub vrf-export C2-hub-export
```

Different policies are used for importing and exporting the route-targets.

Configure the EBGP session to the CE on the hub interface.

```
lab@R1# set routing-instances C2-hub protocols bgp group ce peer-as 64600

[edit]
lab@R1# set routing-instances C2-hub protocols bgp group ce neighbor 192.168.0.42
```

Configure the other routing-instance which will accept the routes from the spokes.

```
[edit]
lab@R1# set routing-instances C2-spoke instance-type vrf

[edit]
lab@R1# set routing-instances C2-spoke interface lo0.2

[edit]
lab@R1# set routing-instances C2-spoke interface ge-0/0/5.312

[edit]
lab@R1# set routing-instances C2-spoke vrf-import C2-spoke-import

[edit]
lab@R1# set routing-instances C2-spoke vrf-export C2-spoke-export
```

Again, different policies are used for importing and exporting route-targets.

```
[edit]
lab@R1# set routing-instances C2-spoke protocols bgp group ce peer-as 64600

[edit]
lab@R1# set routing-instances C2-spoke protocols bgp group ce as-override

[edit]
lab@R1# set routing-instances C2-spoke protocols bgp group ce neighbor 192.168.0.46
```

Configure the EBGP session on the spoke physical interface. Here you have to use the `as-override` option. This ensures that the CE router on the other end will accept the routes and not drop them because of the AS path loop.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set routing-instances C2-spoke protocols bgp group ce peer-as 64600

[edit]
lab@R1# set routing-instances C2-spoke protocols bgp group ce as-override

[edit]
lab@R1# set routing-instances C2-spoke protocols bgp group ce neighbor 192.168.0.46
```

Configure the import and export policies for the two routing instances.

```
[edit]
lab@R1# set policy-options policy-statement C2-hub-export term 1 from protocol [
bgp direct ]

[edit]
lab@R1# set policy-options policy-statement C2-hub-export term 1 then community
set CE2-hub

[edit]
lab@R1# set policy-options policy-statement C2-hub-export term 1 then accept
```

An export policy on the hub routing instance exports the routes learned from the CE and the direct routes. CE2-hub community is set for these routes. If you do not include the direct routes they will be rejected.

```
[edit]
lab@R1# set policy-options policy-statement C2-hub-import term 1 then reject
```

The import policy rejects everything.

```
[edit]
lab@R1# set policy-options policy-statement C2-spoke-import term 1 from protocol bgp

[edit]
lab@R1# set policy-options policy-statement C2-spoke-import term 1 from community CE2-
spoke

[edit]
lab@R1# set policy-options policy-statement C2-spoke-import term 1 then accept
```

The spoke routing instance is configured the opposite way. The import policy imports routes from the spokes with the CE2-spoke community attribute.

```
[edit]
lab@R1# set policy-options policy-statement C2-spoke-export then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

Nothing is exported from the spoke routing instance.

At last the communities must be defined.

```
[edit]
lab@R1# set policy-options community CE2-hub members target:54591:200

[edit]
lab@R1# set policy-options community CE2-spoke members target:54591:201

[edit]
lab@R1# commit
commit complete
```

b) Configure AS looping on the route reflector

```
[edit]
lab@route-reflector# set routing-options autonomous-system 54591 loops 3

[edit]
lab@route-reflector# commit
commit complete
```

c) Configure the C2 VPN on R2.

R2 is also connected to the hub CE of customer C2. The configuration of the routing instances is the same.

```
[edit]
lab@R2# set interfaces lo0.1 family inet address 172.30.5.13/32

[edit]
lab@R2# set interfaces lo0.2 family inet address 172.30.5.14/32

[edit]
lab@R2# set routing-options route-distinguisher-id 172.30.5.2

[edit]
lab@R2# set routing-options autonomous-system 54591 loops 3
```

Configure the hub routing instance.

```
[edit]
lab@R2# set routing-instances C2-hub instance-type vrf

[edit]
lab@R2# set routing-instances C2-hub interface lo0.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set routing-instances C2-hub interface ge-0/0/5.314

[edit]
lab@R2# set routing-instances C2-hub vrf-import C2-hub-import

[edit]
lab@R2# set routing-instances C2-hub vrf-export C2-hub-export

[edit]
lab@R2# set routing-instances C2-hub protocols bgp group ce peer-as 64600

[edit]
lab@R2# set routing-instances C2-hub protocols bgp group ce neighbor 192.168.0.54
```

Configure the spoke routing instance.

```
[edit]
lab@R2# set routing-instances C2-spoke instance-type vrf

[edit]
lab@R2# set routing-instances C2-spoke interface lo0.2

[edit]
lab@R2# set routing-instances C2-spoke interface ge-0/0/5.315

[edit]
lab@R2# set routing-instances C2-spoke vrf-import C2-spoke-import

[edit]
lab@R2# set routing-instances C2-spoke vrf-export C2-spoke-export

[edit]
lab@R2# set routing-instances C2-spoke protocols bgp group ce peer-as 64600

[edit]
lab@R2# set routing-instances C2-spoke protocols bgp group ce neighbor 192.168.0.58

[edit]
lab@R2# set routing-instances C2-spoke protocols bgp group ce as-override
```

Configure the import and export route target policies.

```
[edit]
lab@R2# set policy-options policy-statement C2-hub-export term 1 from protocol [ bgp
direct ]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set policy-options policy-statement C2-hub-export term 1 then community set
CE2-hub

[edit]
lab@R2# set policy-options policy-statement C2-hub-export term 1 then accept

[edit]
lab@R2# set policy-options policy-statement C2-hub-import term 1 then reject

[edit]
lab@R2# set policy-options policy-statement C2-spoke-export term 1 then reject

[edit]
lab@R2# set policy-options policy-statement C2-spoke-import term 1 from protocol bgp

[edit]
lab@R2# set policy-options policy-statement C2-spoke-import term 1 from community CE2-
spoke

[edit]
lab@R2# set policy-options policy-statement C2-spoke-import term 1 then accept

[edit]
lab@R2# set policy-options community CE2-hub members target:54591:200

[edit]
lab@R2# set policy-options community CE2-spoke members target:54591:201

[edit]
lab@R2# commit
commit complete
```

d) Configure C2 VPN on R4.

R4 is a spoke site, which is multi-homed on R5 also.

```
[edit]
lab@R4# set interfaces lo0.2 family inet address 172.30.5.22/32

[edit]
lab@R4# set routing-options autonomous-system 54591 loops 3
```

As a spoke site R4 has just one routing instance with specific route target policies.

```
[edit]
lab@R4# set routing-instances C2-spoke instance-type vrf
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set routing-instances C2-spoke interface lo0.2

[edit]
lab@R4# set routing-instances C2-spoke interface ge-0/0/5.320

[edit]
lab@R4# set routing-instances C2-spoke vrf-import C2-spoke-import

[edit]
lab@R4# set routing-instances C2-spoke vrf-export C2-spoke-export

[edit]
lab@R4# set routing-instances C2-spoke protocols bgp group ce peer-as 64600

[edit]
lab@R4# set routing-instances C2-spoke protocols bgp group ce as-override

[edit]
lab@R4# set routing-instances C2-spoke protocols bgp group ce neighbor 192.168.0.78
```

Configure the route target policies for import and export from the routing instance.

```
[edit]
lab@R4# set policy-options policy-statement C2-spoke-import term 1 from protocol bgp

[edit]
lab@R4# set policy-options policy-statement C2-spoke-import term 1 from community CE2-
hub

[edit]
lab@R4# set policy-options policy-statement C2-spoke-import term 1 then accept
```

The hub routes are imported. These are the routes advertised from the hub routing instance on R1 and R2.

```
[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 1 from protocol [ bgp
direct ]

[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 1 then community set
CE2-spoke

[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 1 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 1 then accept
```

The export route-target policy sets the spoke community to the direct and BGP routes.

```
[edit]
lab@R4# set policy-options community CE2-hub members target:54591:200

[edit]
lab@R4# set policy-options community CE2-spoke members target:54591:201

[edit]
lab@R4# commit
commit complete
```

e) Configure the C2 VPN on R5.

The configuration on R5 is basically the same as on R4. R5 is the second multi-homed site for the spoke site.

```
[edit]
lab@R5# set interfaces lo0.1 family inet address 172.30.5.25/32

[edit]
lab@R5# set routing-options route-distinguisher-id 172.30.5.5

[edit]
lab@R5# set routing-options autonomous-system 54591 loops 3

[edit]
lab@R5# set routing-instances C2-spoke instance-type vrf

[edit]
lab@R5# set routing-instances C2-spoke interface lo0.1

[edit]
lab@R5# set routing-instances C2-spoke interface ge-0/0/5.321

[edit]
lab@R5# set routing-instances C2-spoke vrf-import C2-spoke-import

[edit]
lab@R5# set routing-instances C2-spoke vrf-export C2-spoke-export

[edit]
lab@R5# set routing-instances C2-spoke protocols bgp group ce peer-as 64600
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set routing-instances C2-spoke protocols bgp group ce as-override

[edit]
lab@R5# set routing-instances C2-spoke protocols bgp group ce neighbor 192.168.0.82
```

Configure the export and import policies.

```
[edit]
lab@R5# set policy-options policy-statement C2-spoke-import term 1 from protocol bgp

[edit]
lab@R5# set policy-options policy-statement C2-spoke-import term 1 from community CE2-
hub

[edit]
lab@R5# set policy-options policy-statement C2-spoke-import term 1 then accept

[edit]
lab@R5# set policy-options policy-statement C2-spoke-export term 1 from protocol [ bgp
direct ]

[edit]
lab@R5# set policy-options policy-statement C2-spoke-export term 1 then community set
CE2-spoke

[edit]
lab@R5# set policy-options policy-statement C2-spoke-export term 1 then accept

[edit]
lab@R5# set policy-options community CE2-hub members target:54591:200

[edit]
lab@R5# set policy-options community CE2-spoke members target:54591:201

[edit]
lab@R5# commit
commit complete
```

f) Configure C2 VPN on R7.

Customer C2 has one or more spoke sites connected to R7.

```
[edit]
lab@R7# set interfaces lo0.1 family inet address 172.30.5.33/32

[edit]
lab@R7# set routing-options autonomous-system 54591 loops 3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set routing-options route-distinguisher-id 172.30.5.7

[edit]
lab@R7# set routing-instances C2-spoke instance-type vrf

[edit]
lab@R7# set routing-instances C2-spoke interface lo0.1

[edit]
lab@R7# set routing-instances C2-spoke interface ge-0/0/5.323

[edit]
lab@R7# set routing-instances C2-spoke vrf-import C2-spoke-import

[edit]
lab@R7# set routing-instances C2-spoke vrf-export C2-spoke-export

[edit]
lab@R7# set routing-instances C2-spoke protocols bgp group ce peer-as 64600

[edit]
lab@R7# set routing-instances C2-spoke protocols bgp group ce as-override

[edit]
lab@R7# set routing-instances C2-spoke protocols bgp group ce neighbor 192.168.0.90
```

Configure the export and import policies.

```
[edit]
lab@R7# set policy-options policy-statement C2-spoke-import term 1 from protocol bgp

[edit]
lab@R7# set policy-options policy-statement C2-spoke-import term 1 from community CE2-
hub

[edit]
lab@R7# set policy-options policy-statement C2-spoke-import term 1 then accept

[edit]
lab@R7# set policy-options policy-statement C2-spoke-export term 1 from protocol [ bgp
direct ]

[edit]
lab@R7# set policy-options policy-statement C2-spoke-export term 1 then community set
CE2-spoke
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set policy-options policy-statement C2-spoke-export term 1 then accept

[edit]
lab@R7# set policy-options community CE2-hub members target:54591:200

[edit]
lab@R7# set policy-options community CE2-spoke members target:54591:201

[edit]
lab@R7# commit
commit complete
```

Verify the hub and spoke L3VPN configuration.

```
[edit]
lab@R1# run show bgp summary
Groups: 4 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    1343        598          0          0          0          0
bgp.l3vpn.0
                      16         16          0          0          0          0
inet6.0
                      48         48          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591       2580        863       0       0     6:24:19
Establ
  inet.0: 215/215/215/0
  bgp.l3vpn.0: 16/16/16/0
  inet6.0: 48/48/48/0
  C2-spoke.inet.0: 12/16/16/0
192.168.0.42          64600        832        846       0       0     6:24:42
Establ
  C2-hub.inet.0: 20/20/20/0
192.168.0.46          64600        824        853       0       0     6:24:38
Establ
  C2-spoke.inet.0: 0/0/0/0
192.168.1.3            1620       1332       1095       0       0     6:24:46
Establ
  inet.0: 0/564/402/0
192.168.1.4            1620       1321       1096       0       0     6:24:50
Establ
  inet.0: 383/564/402/0
```

Both EBGP session are established, one from hub routing instance and one from the spoke routing instance. Routes are received only in the hub instance. Those routes are

© 2020 Juniper Networks, Inc. All rights reserved.

learned from the spokes. The spoke routing instance is not receiving routes from the CE. R1 should only advertise routes received from the other spokes through multiprotocol BGP.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.42
```

Nothing is advertised from R1 to the hub CE EBGP peer.

```
[edit]
lab@R1# run show route receive-protocol bgp 192.168.0.42

inet.0: 840 destinations, 1423 routes (672 active, 0 holddown, 330 hidden)

inet.3: 11 destinations, 16 routes (7 active, 0 holddown, 7 hidden)

C2-hub.inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
  Prefix             Nexthop          MED     Lclpref    AS path
* 172.30.5.22/32     192.168.0.42                        64600 54591 I
* 172.30.5.25/32     192.168.0.42                        64600 54591 I
* 172.30.5.33/32     192.168.0.42                        64600 54591 I
* 172.31.64.0/24     192.168.0.42                        64600 I
* 172.31.65.0/24     192.168.0.42                        64600 I
* 172.31.66.0/24     192.168.0.42                        64600 I
* 172.31.67.0/24     192.168.0.42                        64600 I
* 172.31.68.0/24     192.168.0.42                        64600 I
* 172.31.69.0/24     192.168.0.42                        64600 I
* 172.31.70.0/24     192.168.0.42                        64600 I
* 172.31.71.0/24     192.168.0.42                        64600 I
* 172.31.72.0/24     192.168.0.42                        64600 54591 54591 I
* 172.31.73.0/24     192.168.0.42                        64600 54591 54591 I
* 172.31.74.0/24     192.168.0.42                        64600 54591 54591 I
* 172.31.75.0/24     192.168.0.42                        64600 54591 54591 I
* 172.31.76.0/24     192.168.0.42                        64600 54591 54591 I
* 172.31.77.0/24     192.168.0.42                        64600 54591 54591 I
* 192.168.0.76/30    192.168.0.42                        64600 54591 I
* 192.168.0.80/30    192.168.0.42                        64600 54591 I
* 192.168.0.88/30    192.168.0.42                        64600 54591 I
...
```

All spoke routes are received from the CE in the hub routing instance. Routes that are received from the spokes on R4, R5 and R7 are having an AS path that contains twice the AS 54591. This is because we used the `as-override` option. This is also the reason why you need to enable each router to accept an AS path with its own AS.

© 2020 Juniper Networks, Inc. All rights reserved.

The spoke routing instance advertises the routes learned from the other spokes to the CE router.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.46

C2-spoke.inet.0: 15 destinations, 19 routes (15 active, 0 holddown, 0 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 172.30.5.22/32     Self                                    I
* 172.30.5.25/32     Self                                    I
* 172.30.5.33/32     Self                                    I
* 172.31.72.0/24     Self                                    54591 I
* 172.31.73.0/24     Self                                    54591 I
* 172.31.74.0/24     Self                                    54591 I
* 172.31.75.0/24     Self                                    54591 I
* 172.31.76.0/24     Self                                    54591 I
* 172.31.77.0/24     Self                                    54591 I
* 192.168.0.76/30    Self                                    I
* 192.168.0.80/30    Self                                    I
* 192.168.0.88/30    Self                                    I
```

No routes are received from that CE peer.

```
[edit]
lab@R1# run show route receive-protocol bgp 192.168.0.46

inet.0: 840 destinations, 1423 routes (672 active, 0 holddown, 330 hidden)

inet.3: 11 destinations, 16 routes (7 active, 0 holddown, 7 hidden)

C2-hub.inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)

C2-spoke.inet.0: 15 destinations, 19 routes (15 active, 0 holddown, 0 hidden)

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

mpls.0: 50 destinations, 50 routes (50 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)

inet6.0: 79 destinations, 80 routes (79 active, 0 holddown, 0 hidden)

inet6.3: 7 destinations, 9 routes (7 active, 0 holddown, 0 hidden)
```

Routes learned from remote spokes are reachable via MPLS LSP.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show route table C2-spoke.inet.0 172.31.72.0/24

C2-spoke.inet.0: 15 destinations, 19 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.72.0/24     *[BGP/170] 02:40:10, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.6 via ge-0/0/4.114, Push 301168
                    [BGP/170] 00:01:31, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R8
                      to 172.30.0.2 via ae0.0, label-switched-path R1-to-R8
```

R2 is also connected to the hub site for redundancy.

R2

```
[edit]
lab@R2# run show bgp summary
Groups: 4 Peers: 5 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    1128        402          0          0          0          0
bgp.l3vpn.0
                       0          0          0          0          0          0
inet6.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591        608        444       0       1       40:52
Establ
  inet.0: 217/600/600/0
  inet6.0: 48/48/48/0
  bgp.l3vpn.0: 16/16/16/0
  C2-spoke.inet.0: 12/16/16/0
192.168.0.54          64600        635        635       0       0     4:48:31
Establ
  C2-hub.inet.0: 20/20/20/0
192.168.0.58          64600        620        649       0       0     4:48:25
Establ
  C2-spoke.inet.0: 0/0/0/0
192.168.1.3            1620       1089        871       0       0     4:48:32
Establ
  inet.0: 383/564/402/0
192.168.1.4            1620       1089        870       0       0     4:48:30
Establ
  inet.0: 0/564/402/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Both EBGP sessions are established and as on R1 only the hub session is receiving routes.

```
[edit]
lab@R2# run show route receive-protocol bgp 192.168.0.54

inet.0: 645 destinations, 1209 routes (476 active, 0 holddown, 331 hidden)

inet.3: 11 destinations, 16 routes (7 active, 0 holddown, 7 hidden)

C2-hub.inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
  Prefix             Nexthop         MED   Lclpref  AS path
* 172.30.5.22/32     192.168.0.54                   64600 54591 I
* 172.30.5.25/32     192.168.0.54                   64600 54591 I
* 172.30.5.33/32     192.168.0.54                   64600 54591 I
* 172.31.64.0/24     192.168.0.54                   64600 I
* 172.31.65.0/24     192.168.0.54                   64600 I
* 172.31.66.0/24     192.168.0.54                   64600 I
* 172.31.67.0/24     192.168.0.54                   64600 I
* 172.31.68.0/24     192.168.0.54                   64600 I
* 172.31.69.0/24     192.168.0.54                   64600 I
* 172.31.70.0/24     192.168.0.54                   64600 I
* 172.31.71.0/24     192.168.0.54                   64600 I
* 172.31.72.0/24     192.168.0.54                   64600 54591 54591 I
* 172.31.73.0/24     192.168.0.54                   64600 54591 54591 I
* 172.31.74.0/24     192.168.0.54                   64600 54591 54591 I
* 172.31.75.0/24     192.168.0.54                   64600 54591 54591 I
* 172.31.76.0/24     192.168.0.54                   64600 54591 54591 I
* 172.31.77.0/24     192.168.0.54                   64600 54591 54591 I
* 192.168.0.76/30    192.168.0.54                   64600 54591 I
* 192.168.0.80/30    192.168.0.54                   64600 54591 I
* 192.168.0.88/30    192.168.0.54                   64600 54591 I
...
```

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.0.54
```

No routes are advertised to that peer from hub routing instance.

```
[edit]
lab@R2# run show route receive-protocol bgp 192.168.0.58

inet.0: 645 destinations, 1209 routes (476 active, 0 holddown, 331 hidden)

inet.3: 11 destinations, 16 routes (7 active, 0 holddown, 7 hidden)

C2-hub.inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
C2-spoke.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

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

mpls.0: 46 destinations, 46 routes (46 active, 0 holddown, 0 hidden)

inet6.0: 31 destinations, 32 routes (31 active, 0 holddown, 0 hidden)

inet6.3: 7 destinations, 9 routes (7 active, 0 holddown, 0 hidden)
```

With spoke session as the opposite. No routes are received from the spoke peer.

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.0.58

C2-spoke.inet.0: 15 destinations, 19 routes (15 active, 0 holddown, 0 hidden)
  Prefix             Nexthop              MED     Lclpref    AS path
* 172.30.5.22/32     Self                                    I
* 172.30.5.25/32     Self                                    I
* 172.30.5.33/32     Self                                    I
* 172.31.72.0/24     Self                                    54591 I
* 172.31.73.0/24     Self                                    54591 I
* 172.31.74.0/24     Self                                    54591 I
* 172.31.75.0/24     Self                                    54591 I
* 172.31.76.0/24     Self                                    54591 I
* 172.31.77.0/24     Self                                    54591 I
* 192.168.0.76/30    Self                                    I
* 192.168.0.80/30    Self                                    I
* 192.168.0.88/30    Self                                    I
```

Routes received from the Multiprotocol BGP session from other peers are advertised via the spoke EBGP session.

R4

```
[edit]
lab@R4# run show route advertising-protocol bgp 192.168.0.78

C2-spoke.inet.0: 23 destinations, 49 routes (23 active, 0 holddown, 0 hidden)
  Prefix             Nexthop          MED     Lclpref    AS path
* 172.30.5.9/32      Self                                I
* 172.30.5.25/32     Self                                54591 54591 I
* 172.30.5.33/32     Self                                54591 54591 I
* 172.31.64.0/24     Self                                54591 I
* 172.31.65.0/24     Self                                54591 I
* 172.31.66.0/24     Self                                54591 I
* 172.31.67.0/24     Self                                54591 I
* 172.31.68.0/24     Self                                54591 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.69.0/24    Self                                  54591 I
* 172.31.70.0/24    Self                                  54591 I
* 172.31.71.0/24    Self                                  54591 I
* 172.31.72.0/24    192.168.0.78                          54591 I
* 172.31.73.0/24    192.168.0.78                          54591 I
* 172.31.74.0/24    192.168.0.78                          54591 I
* 172.31.75.0/24    192.168.0.78                          54591 I
* 172.31.76.0/24    Self                                  54591 54591 54591 I
* 172.31.77.0/24    Self                                  54591 54591 54591 I
* 192.168.0.40/30   Self                                  I
* 192.168.0.80/30   Self                                  54591 54591 I
* 192.168.0.88/30   Self                                  54591 54591 I
```

Routes received from the hub site are advertised to the spoke. Four routes have a next-hop of the IP address of the local spoke peer. Those routes are advertised back to the spoke CE because the `as-override` command overrides the 64600 AS number from the as-path.

```
[edit]
lab@R4# run show route receive-protocol bgp 192.168.0.78

inet.0: 678 destinations, 678 routes (673 active, 0 holddown, 5 hidden)

inet.3: 13 destinations, 19 routes (8 active, 0 holddown, 9 hidden)

C1.inet.0: 20 destinations, 34 routes (20 active, 0 holddown, 6 hidden)

C2-spoke.inet.0: 23 destinations, 29 routes (23 active, 0 holddown, 0 hidden)
  Prefix                Nexthop              MED     Lclpref    AS path
* 172.31.72.0/24        192.168.0.78                            64600 I
* 172.31.73.0/24        192.168.0.78                            64600 I
* 172.31.74.0/24        192.168.0.78                            64600 I
* 172.31.75.0/24        192.168.0.78                            64600 I

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

mpls.0: 55 destinations, 55 routes (55 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 39 destinations, 39 routes (39 active, 0 holddown, 0 hidden)

inet6.0: 81 destinations, 83 routes (81 active, 0 holddown, 0 hidden)

inet6.3: 8 destinations, 10 routes (8 active, 0 holddown, 0 hidden)
```

The four routes are received from the spoke CE attached to R4.

© 2020 Juniper Networks, Inc. All rights reserved.

On R7 the two prefixes are received from the spoke. All other prefixes from the hub are send to the CE.

R7

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.90

C2-spoke.inet.0: 23 destinations, 47 routes (23 active, 0 holddown, 0 hidden)
  Prefix             Nexthop      MED     Lclpref    AS path
* 172.30.5.9/32      Self                            I
* 172.30.5.22/32     Self                            54591 54591 I
* 172.30.5.25/32     Self                            54591 54591 I
* 172.31.64.0/24     Self                            54591 I
* 172.31.65.0/24     Self                            54591 I
* 172.31.66.0/24     Self                            54591 I
* 172.31.67.0/24     Self                            54591 I
* 172.31.68.0/24     Self                            54591 I
* 172.31.69.0/24     Self                            54591 I
* 172.31.70.0/24     Self                            54591 I
* 172.31.71.0/24     Self                            54591 I
* 172.31.72.0/24     Self                            54591 54591 54591 I
* 172.31.73.0/24     Self                            54591 54591 54591 I
* 172.31.74.0/24     Self                            54591 54591 54591 I
* 172.31.75.0/24     Self                            54591 54591 54591 I
* 172.31.76.0/24     192.168.0.90                    54591 I
* 172.31.77.0/24     192.168.0.90                    54591 I
* 192.168.0.40/30    Self                            I
* 192.168.0.76/30    Self                            54591 54591 I
* 192.168.0.80/30    Self                            54591 54591 I
```

Verify the VPN connectivity from R4 to R7.

R4

```
[edit]
lab@R4# run ping routing-instance C2-spoke 172.30.5.33
PING 172.30.5.33 (172.30.5.33): 56 data bytes
64 bytes from 172.30.5.33: icmp_seq=0 ttl=61 time=4.608 ms
64 bytes from 172.30.5.33: icmp_seq=1 ttl=61 time=4.472 ms
64 bytes from 172.30.5.33: icmp_seq=2 ttl=61 time=9.751 ms
^C
--- 172.30.5.33 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.472/6.277/9.751/2.457 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run traceroute routing-instance C2-spoke 172.30.5.33 no-resolve
traceroute to 172.30.5.33 (172.30.5.33), 30 hops max, 52 byte packets
 1  172.30.0.5  6.556 ms  5.385 ms  5.674 ms
     MPLS Label=354208 CoS=0 TTL=1 S=1
 2  192.168.0.42  152.793 ms  98.433 ms  10.521 ms
 3  192.168.0.45  5.279 ms  4.621 ms  4.692 ms
 4  * * *
 5  * * *
 6  * * *
 7  172.30.0.45  48.077 ms  79.763 ms  735.806 ms
     MPLS Label=351824 CoS=0 TTL=1 S=1
 8  172.30.5.33  327.947 ms  408.870 ms  396.012 ms
```

Ping and traceroute from R4 to R5.

```
[edit]
lab@R4# run ping routing-instance C2-spoke 172.30.5.25
PING 172.30.5.25 (172.30.5.25): 56 data bytes
64 bytes from 172.30.5.25: icmp_seq=0 ttl=61 time=9.578 ms
64 bytes from 172.30.5.25: icmp_seq=1 ttl=61 time=4.588 ms
^C
--- 172.30.5.25 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.588/7.083/9.578/2.495 ms

[edit]
lab@R4# run traceroute routing-instance C2-spoke 172.30.5.25
traceroute to 172.30.5.25 (172.30.5.25), 30 hops max, 40 byte packets
 1  172.30.5.9 (172.30.5.9)  3.854 ms  2.514 ms  2.478 ms
 2  192.168.0.42 (192.168.0.42)  4.153 ms  4.864 ms  4.951 ms
 3  192.168.0.45 (192.168.0.45)  8.611 ms  3.502 ms  4.086 ms
 4  172.30.5.25 (172.30.5.25)  9.497 ms *  3.861 ms
```

From the output of the traceroute command is can be seen that the traffic is passing through the hub site.

8) Make sure that the PE-CE link subnets in the customer C2 VPN are advertised to the customer remote VPN sites.

This requirement is accomplished with matching the direct subnets in export route target policies. However, to provide for data plane connectivity for these subnets you also have to enable the `vrf-table-label` option.

By default, Junos assigns a different VPN label for each CE interface. When the penultimate router removes the transport label, the packet that is received on the egress

© 2020 Juniper Networks, Inc. All rights reserved.

LSR only has the VPN label applied. Based on that label the Egress LSR forwards the traffic to the CE interface. To have connectivity to the directly connected subnets, an active route pointing to the CE routes should be present in the table. The `vrf-table-label` enables a virtual LSI infrastructure. A VRF label is advertised for all directly connected routes for a given VRF. This way the egress PE router performs IP lookup to find the CE interface. Configure this on all PE-CE routing instances on R1, R2, R4, R5 and R7.

R1

```
[edit]
lab@R1# set routing-instances C2-hub vrf-table-label

[edit]
lab@R1# set routing-instances C2-spoke vrf-table-label

[edit]
lab@R1# commit
commit complete
```

R2

```
[edit]
lab@R2# set routing-instances C2-hub vrf-table-label

[edit]
lab@R2# set routing-instances C2-spoke vrf-table-label

[edit]
lab@R2# commit
commit complete
```

R4

```
[edit]
lab@R4# set routing-instances C2-spoke vrf-table-label

[edit]
lab@R4# commit
commit complete
```

R5

```
[edit]
lab@R5# set routing-instances C2-spoke vrf-table-label
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# commit
commit complete
```

R7

```
[edit]
lab@R7# set routing-instances C2-spoke vrf-table-label

[edit]
lab@R7# commit
commit complete
```

Verify the MPLS label advertisement after the change.

R1

```
[edit]
lab@R1# run show route advertising-protocol bgp 172.30.5.41 table C2-hub.inet.0
extensive

C2-hub.inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
* 172.30.5.9/32 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.30.5.1:4
     VPN Label: 16
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [54591] I
     Communities: target:54591:200 origin:54591:1
...
```

R4

```
[edit]
lab@R4# run show route table mpls.0 label 16 extensive

mpls.0: 77 destinations, 77 routes (77 active, 0 holddown, 0 hidden)
16  (1 entry, 0 announced)
        *VPN    Preference: 0
                Next hop type: Router, Next hop index: 0
                Address: 0xd9d7a30
                Next-hop reference count: 1
                Next hop: via lsi.256 (C2-spoke), selected
                Label operation: Pop
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Load balance label: None;
Label element ptr: 0xd9da720
Label parent element ptr: 0x0
Label element references: 1
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
State: <Active NotInstall Int Ext LsiL3>
Age: 2:35
Validation State: unverified
Task: RT
AS path: I
```

An IP route lookup is performed in the C2-hub.inet.0 table.

9) Make sure that if a route is originated from customer C2 site S1 and S2, it is never advertised back to the same site.

To accomplish this requirement, you need to not accept routes advertised from the same site. This is needed in cases where you have multi-homed sites with `as-override` feature, the situation that you have with the HUB site S1 and spoke site S2.

You cannot use route-target community for this purpose because of two reasons. First that community is used for attaching the routes to the proper routing instance. The second reason is the inability to distinguish the routes that are advertised by the local Spoke site.

A new extended community is used to mark the routes and filter them on the other multi-homed site. That community is called the Sight of Origin (SoO).

Configure R1, R2, R3 and R4 with SoO community.

```
[edit]
lab@R1# set policy-options community C2-S1-origin members origin:54591:1

[edit]
lab@R1# set policy-options policy-statement C2-hub-export term 1 then community add C2-
S1-origin
```

Create the `origin` extended community type. The structure is the same as the target community but instead of target, the `origin` keyword is used when creating the community.

© 2020 Juniper Networks, Inc. All rights reserved.

To filter the routes that are advertised from R2 and coming to R1, you should reject the routes from Multiprotocol BGP with that specific origin community.

```
[edit]
lab@R1# set policy-options policy-statement C2-spoke-import term reject-origin from
protocol bgp

[edit]
lab@R1# set policy-options policy-statement C2-spoke-import term reject-origin from
community C2-S1-origin

[edit]
lab@R1# set policy-options policy-statement C2-spoke-import term reject-origin then
reject
```

That term should be moved before the term accepting the routes with appropriate route target community.

```
[edit]
lab@R1# insert policy-options policy-statement C2-spoke-import term reject-origin
before term 1

[edit]
lab@R1# commit
commit complete
```

After the commit, the routes from R1 are sent with both a route target and origin community.

```
[edit]
lab@R1# run show route advertising-protocol bgp 172.30.5.41 table C2-hub.inet.0
extensive

C2-hub.inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
* 172.30.5.9/32 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.30.5.1:4
     VPN Label: 16
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [54591] I
     Communities: target:54591:200 origin:54591:1

* 172.30.5.22/32 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.30.5.1:4
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
VPN Label: 16
Nexthop: Self
Flags: Nexthop Change
Localpref: 100
AS path: [54591] 64600 54591 I
Communities: target:54591:200 origin:54591:1
```

The configuration on R2 should be the same as on R1. You need the same because it connects to the same customer CE site.

On R4 and R5 a different origin community should be used to represent the routes advertised from that side. This way they will accepted from the other site because of the route-target applied, however they will be limited by R4 and R5 in the import direction.

```
[edit]
lab@R4# set policy-options community C2-S2-origin members origin:54591:2

[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 1 then community add
C2-S2-origin

[edit]
lab@R4# set policy-options policy-statement C2-spoke-import term reject-origin from
protocol bgp

[edit]
lab@R4# set policy-options policy-statement C2-spoke-import term reject-origin from
community C2-S2-origin

[edit]
lab@R4# set policy-options policy-statement C2-spoke-import term reject-origin then
reject

[edit]
lab@R4# insert policy-options policy-statement C2-spoke-import term reject-origin
before term 1

[edit]
lab@R4# commit
commit complete
```

On R5 the same configuration should be applied.

10) Configure route target BGP family on all routes.

The route target family provides scalability improvement for large L3VPN networks. The PE routers send a list with route targets that they are interested in based on their locally

© 2020 Juniper Networks, Inc. All rights reserved.

configured VRFs. The route reflector only reflects the routes with the required route target to the PEs.

Repeat this on all eight devices.

```
[edit]
lab@R1# set protocols bgp group ibgp family route-target

[edit]
lab@R1# commit
commit complete
```

Configure the route target family on the route reflector.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family route-target

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family route-target

[edit]
lab@route-reflector# commit
commit complete
```

Verify the content of `bgp.rtarget.0` table, which is used for storing route target family information.

```
[edit]
lab@route-reflector# run show route table bgp.rtarget.0

bgp.rtarget.0: 3 destinations, 9 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

54591:54591:100/96
                   *[BGP/170] 00:00:37, localpref 100, from 172.30.5.3
                      AS path: I, validation-state: unverified
                    to Receive
                    [BGP/170] 00:00:58, localpref 100, from 172.30.5.4
                      AS path: I, validation-state: unverified
                    to Receive
                    [BGP/170] 00:00:52, localpref 100, from 172.30.5.6
                      AS path: I, validation-state: unverified
                    to Receive
                    [BGP/170] 00:00:46, localpref 100, from 172.30.5.8
                      AS path: I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    to Receive
54591:54591:200/96
                  *[BGP/170] 00:00:58, localpref 100, from 172.30.5.4
                    AS path: I, validation-state: unverified
                    to Receive
                   [BGP/170] 00:00:53, localpref 100, from 172.30.5.5
                    AS path: I, validation-state: unverified
                    to Receive
                   [BGP/170] 00:00:49, localpref 100, from 172.30.5.7
                    AS path: I, validation-state: unverified
                    to Receive
54591:54591:201/96
                  *[BGP/170] 00:00:59, localpref 100, from 172.30.5.1
                    AS path: I, validation-state: unverified
                   to Receive
                   [BGP/170] 00:00:41, localpref 100, from 172.30.5.2
                    AS path: I, validation-state: unverified
                    to Receive
```

On the route reflector the routes are present. You can see the request received from all the PEs which have a particular VRF configured.

```
[edit]
lab@R1# run show route table bgp.rtarget.0

bgp.rtarget.0: 3 destinations, 4 routes (1 active, 0 holddown, 3 hidden)
+ = Active Route, - = Last Active, * = Both

54591:54591:201/96
                  *[RTarget/5] 00:03:15
                     Type Default
                     Local
```

But on R1 only the local route is present in the table. Besides that, three more hidden routes are present.

```
[edit]
lab@R1# run show route table bgp.rtarget.0 hidden extensive

bgp.rtarget.0: 3 destinations, 4 routes (1 active, 0 holddown, 3 hidden)
54591:54591:100/96 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Next hop type: Unusable, Next hop index: 0
                Address: 0xbe73e50
                Next-hop reference count: 3
                State: <Hidden Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 2:37
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                Validation State: unverified
                Task: BGP_54591.172.30.5.41
                AS path: I  (Originator)
                Cluster list:  0.0.0.2
                Originator ID: 172.30.5.41
                Accepted
                Localpref: 100
                Router ID: 172.30.5.41
                Indirect next hops: 1
                        Protocol next hop: 172.30.5.41
                        Indirect next hop: 0x0 - INH Session ID: 0x0
...
```

The route reflector is changing the protocol next-hop to self. The routes are hidden because the route reflector is not enabled for MPLS and its loopback is not present in the `inet.3` table.

To make route reflector loopback IP address present in the `inet.3` table you can use `rib-groups` to copy the RR loopback IP address from `inet.0` to `inet.3`.

Repeat this on all eight devices.

```
[edit]
lab@R1# set policy-options policy-statement rr-loopback-to-inet3 term 1 from route-
filter 172.30.5.41/32 exact

[edit]
lab@R1# set policy-options policy-statement rr-loopback-to-inet3 term 1 to rib inet.3

[edit]
lab@R1# set policy-options policy-statement rr-loopback-to-inet3 term 1 then accept

[edit]
lab@R1# set policy-options policy-statement rr-loopback-to-inet3 term 2 to rib inet.3

[edit]
lab@R1# set policy-options policy-statement rr-loopback-to-inet3 term 2 then reject
```

Routing policy is created so that only the loopback IP address of the route reflector is copied from `inet.0` to `inet.3`.

Then the rib-group is created. The first routing table specified with `rib-groups` is the source of the routes, others are where the routes are transferred to. The `import-policy` is used to filter which routes are copied.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set routing-options rib-groups rr-inet0-inet3 import-rib [ inet.0 inet.3 ]

[edit]
lab@R1# set routing-options rib-groups rr-inet0-inet3 import-policy rr-loopback-to-
inet3

[edit]
lab@R1# set protocols isis rib-group rr-inet0-inet3

[edit]
lab@R1# commit
commit complete
```

When defined the rib-group should be applied to the protocol from where the route is inserted in the primary routing table. In that case this is IS-IS.

```
[edit]
lab@R1# run show route table inet.3 172.30.5.41

inet.3: 12 destinations, 17 routes (8 active, 0 holddown, 7 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.41/32     *[IS-IS/15] 00:00:40, metric 10
                     > to 172.30.0.66 via ge-0/0/4.206
```

After the commit the route is present in the `inet.3` table. The route is learned via the IS-IS protocol.

```
[edit]
lab@R1# run show route table bgp.rtarget.0

bgp.rtarget.0: 3 destinations, 4 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

54591:54591:100/96
                   *[BGP/170] 00:02:23, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                     > to 172.30.0.66 via ge-0/0/4.206
54591:54591:200/96
                   *[BGP/170] 00:02:23, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
                     > to 172.30.0.66 via ge-0/0/4.206
54591:54591:201/96
                   *[RTarget/5] 00:02:23
                      Type Default
                      Local
                    [BGP/170] 00:02:23, localpref 100, from 172.30.5.41
                      AS path: I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
> to 172.30.0.66 via ge-0/0/4.206
```

Routes received from the route reflector in the route target routing table are resolved and active.

11) Configure route exchange between customer C1 site 2 and customer C2 site 2.

Customers' C1 and C2 second sites should communicate locally on R4. This could be accomplished again with rib-groups.

One rib-group is needed for copying the routes from the C1 routing instance to C2 routing instance and one more in the opposite direction.

```
[edit]
lab@R4# set routing-options rib-groups C1-C2-vpn import-rib [ C1.inet.0 C2-spoke.inet.0
]

[edit]
lab@R4# set routing-options rib-groups C2-C1-vpn import-rib [ C2-spoke.inet.0 C1.inet.0
]
```

Then both policies should be applied to the routing protocols in the appropriate routing instances.

```
[edit]
lab@R4# set routing-instances C1 routing-options interface-routes rib-group inet C1-C2-
vpn
```

The rib-groups should copy the direct interface route between the tables. This is needed for resolving the routes received from the routing protocols.

Note: rib-groups applied under the `interfaces-routes` hierarchy copies routes in both directions. The order of the ribs in rib-groups definition does not matter.

```
[edit]
lab@R4# set routing-instances C1 protocols ospf rib-group C1-C2-vpn

[edit]
lab@R4# set routing-instances C2-spoke routing-options interface-routes rib-group inet
C2-C1-vpn
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set routing-instances C2-spoke protocols bgp family inet unicast rib-group C2-
C1-vpn

[edit]
lab@R4# commit
commit complete
```

Verify the presence of the copied routes.

```
[edit]
lab@R4# run show route table C1.inet.0 172.31.72.0/24 extensive

C1.inet.0: 27 destinations, 42 routes (27 active, 0 holddown, 6 hidden)
172.31.72.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.31.72.0/24 -> {192.168.0.78}
RTF { 54591:100/64 }
Page 0 idx 0 Type 1 val 154600c
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [54591] 64600 I
    Communities: target:54591:100
Path 172.31.72.0 from 192.168.0.78 Vector len 4.  Val: 0
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 602
                Address: 0x155e104
                Next-hop reference count: 17
                Source: 192.168.0.78
                Next hop: 192.168.0.78 via ge-0/0/5.320, selected
                State: <Secondary Active Ext>
                Peer AS: 64600
                Age: 5:48
                Task: BGP_64600.192.168.0.78+50073
                Announcement bits (3): 1-KRT 2-BGP Route Target 3-
BGP_RT_Background
                AS path: 64600 I
                Accepted
                Localpref: 100
                Router ID: 172.31.75.1
                Primary Routing Table C2-spoke.inet.0
```

The primary routing table for the route is table C2-spoke.inet.0.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# run show route table C2-spoke.inet.0 172.31.48.4/30 extensive

C2-spoke.inet.0: 39 destinations, 69 routes (33 active, 0 holddown, 6 hidden)
172.31.48.4/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.31.48.4/30 -> {192.168.0.74}
        *OSPF   Preference: 10
                Next hop type: Router, Next hop index: 656
                Address: 0x155e0b8
                Next-hop reference count: 29
                Next hop: 192.168.0.74 via ge-0/0/5.319, selected
                State: <Secondary Active Int>
                Age: 6:29       Metric: 2
                Area: 0.0.0.0
                Task: C1-OSPF
                Announcement bits (1): 0-KRT
                AS path: I
                Communities: rte-type:0.0.0.0:1:0
                Primary Routing Table C1.inet.0
```

The routes from C1 are copied to the C2-spoke routing table.

But copying the routes between the routing instances makes them eligible for export by the automatic route target export policies on both routing instances. That behavior is because no explicit `vrf-export` policy is used. Instead we used the `vrf-target` option.

```
[edit]
lab@R4# run show route advertising-protocol bgp 172.30.5.41 table C1.inet.0

C1.inet.0: 27 destinations, 52 routes (27 active, 0 holddown, 6 hidden)
  Prefix                Nexthop              MED     Lclpref    AS path
* 172.30.5.21/32        Self                         100        I
* 172.30.5.22/32        Self                         100        I
* 172.31.48.4/30        Self                 2       100        I
* 172.31.48.8/30        Self                 3       100        I
* 172.31.48.12/30       Self                 2       100        I
* 172.31.48.16/30       Self                 6       100        I
* 172.31.63.2/32        Self                 2       100        I
* 172.31.63.3/32        Self                 1       100        I
* 172.31.63.5/32        Self                 2       100        I
* 172.31.72.0/24        Self                         100        64600 I
* 172.31.73.0/24        Self                         100        64600 I
* 172.31.74.0/24        Self                         100        64600 I
* 172.31.75.0/24        Self                         100        64600 I
* 192.168.0.72/30       Self                         100        I
* 192.168.0.76/30       Self                         100        I
```

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# run show route advertising-protocol bgp 172.30.5.41 table C2-spoke.inet.0

C2-spoke.inet.0: 39 destinations, 47 routes (33 active, 0 holddown, 6 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.5.21/32          Self                         100        I
* 172.30.5.22/32          Self                         100        I
* 172.31.72.0/24          Self                         100        64600 I
* 172.31.73.0/24          Self                         100        64600 I
* 172.31.74.0/24          Self                         100        64600 I
* 172.31.75.0/24          Self                         100        64600 I
* 192.168.0.72/30         Self                         100        I
* 192.168.0.76/30         Self                         100        I
```

The outputs above prove that behavior. But the requirement is to limit the interconnected traffic only to the local site. Fix that by creating explicit export route target policies for both routing instances.

Configure C1 routing instance with a custom route target import and export policies.

```
[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 1 from protocol direct

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 1 from route-filter
192.168.0.72/30 exact

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 1 from route-filter
172.30.5.21/32 exact

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 1 then community set CE1

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 1 then accept

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 2 from protocol ospf

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 2 then community set CE1

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 2 then accept

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 3 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: The term 3 with a reject action is needed, because the routes copied from the C2 customer are BGP routes and will automatically advertised.

```
[edit]
lab@R4# set policy-options policy-statement C1-vpn-import term 1 from protocol bgp

[edit]
lab@R4# set policy-options policy-statement C1-vpn-import term 1 from community CE1

[edit]
lab@R4# set policy-options policy-statement C1-vpn-import term 1 then accept

[edit]
lab@R4# set policy-options community CE1 members target:54591:100

[edit]
lab@R4# delete routing-instances C1 vrf-target

[edit]
lab@R4# set routing-instances C1 vrf-export C1-vpn-export

[edit]
lab@R4# set routing-instances C1 vrf-import C1-vpn-import
```

Modifying C2-spoke export policy.

```
[edit]
lab@R4# delete policy-options policy-statement C2-spoke-export term 1 from protocol bgp

[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 1 from route-filter
192.168.0.76/30 exact

[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 1 from route-filter
172.30.5.22/32 exact

[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 2 from protocol bgp

[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 2 then community set
CE2-spoke

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# set policy-options policy-statement C2-spoke-export term 2 then community add
C2-S2-origin

[edit]
lab@R4# set policy-options policy-statement C2-spoke-export term 2 then accept

[edit]
lab@R4# commit
commit complete
```

The result is that only the desired routes are advertised to the other PE routers from both routing instances.

Verify connectivity from the C2-spoke routing instance to C1 loopback IP address.

```
[edit]
lab@R4# run ping routing-instance C2-spoke 172.30.5.21
PING 172.30.5.21 (172.30.5.21): 56 data bytes
64 bytes from 172.30.5.21: icmp_seq=0 ttl=64 time=4.350 ms
64 bytes from 172.30.5.21: icmp_seq=1 ttl=64 time=0.405 ms
64 bytes from 172.30.5.21: icmp_seq=2 ttl=64 time=0.240 ms
64 bytes from 172.30.5.21: icmp_seq=3 ttl=64 time=0.235 ms
^C
--- 172.30.5.21 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.235/1.308/4.350/1.758 ms
```

12) Configure Customer C1 Internet access.

Customer C1 should be provided with Internet access at site S2. Internet access is available in the global routing table `inet.0` on R3 and R4. You have to ensure that a failure of R3 or R4 will not cause Internet service disruption.

To provide reachability of C1 routing instance routes from the global routing table again rib-groups can be used.

a) Configure R3 for Internet access.

```
[edit]
lab@R3# set routing-options rib-groups C1-vpn-inet import-rib [ C1.inet.0 inet.0 ]

[edit]
lab@R3# set routing-instances C1 routing-options interface-routes rib-group inet C1-
vpn-inet

[edit]
lab@R3# set routing-instances C1 protocols ospf rib-group C1-vpn-inet
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# commit
commit complete
```

The rib-group copies routing instance C1 routes to `inet.0` routing table.

```
[edit]
lab@R3# run show route protocol ospf extensive

inet.0: 758 destinations, 760 routes (685 active, 0 holddown, 75 hidden)
172.31.48.0/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.31.48.0/30 -> {192.168.0.70}
Aggregated into 172.31.48.0/20
        *OSPF   Preference: 10
                Next hop type: Router, Next hop index: 635
                Address: 0x155c8f8
                Next-hop reference count: 35
                Next hop: 192.168.0.70 via ge-0/0/5.318, selected
                State: <Secondary Active Int>
                Age: 26:00      Metric: 6
                Area: 0.0.0.0
                Task: C1-OSPF
                Announcement bits (4): 0-KRT 2-LDP 5-Aggregate 7-Resolve tree 4
                AS path: I
                Communities: rte-type:0.0.0.0:1:0
                Primary Routing Table C1.inet.0
...
```

OSPF routes from `C1.inet.0` are copied to the global routing table `inet.0`.

To direct traffic from the C1 routing instance to the global routing table without copying all routes in the instance routing table, you can use default route with a next-hop lookup in the `inet.0` table.

```
[edit]
lab@R3# set routing-instances C1 routing-options static route 0/0 next-table inet.0
```

This static route should be announced to all CE routers including the locally attached CE. OSPF export policy is applied to advertise the static route to the local CE and for export towards the remote CEs the route-target policy should be modified.

```
[edit]
lab@R3# set policy-options policy-statement C1-default-to-ospf term 1 from protocol
static
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set policy-options policy-statement C1-default-to-ospf term 1 from route-filter
0/0 exact


[edit]
lab@R3# set policy-options policy-statement C1-default-to-ospf term 1 then accept

[edit]
lab@R3# set routing-instances C1 protocols ospf export C1-default-to-ospf

[edit]
lab@R3# commit
commit complete
```

On R3 everything in the routing instance routing table is advertised to Multiprotocol BGP because the `vrf-target` option is used.

Verify the default route advertisement to other PE routers.

```
[edit]
lab@R8# run show route table C1.inet.0

C1.inet.0: 21 destinations, 40 routes (20 active, 0 holddown, 12 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.17/32     *[BGP/170] 00:00:57, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-1
                      to 172.30.0.45 via ge-0/0/4.178, label-switched-path
Bypass->172.30.0.9->172.30.0.2
                      to 172.30.0.45 via ge-0/0/4.178, label-switched-path
Bypass->172.30.0.37->172.30.0.34
172.30.5.21/32     *[BGP/170] 00:00:57, localpref 100, from 172.30.5.41
                      AS path: I
...
```

On R8 default route is not present. R3 is not advertising the static route. The reason for that is the next-hop of the static route. Next-hop is not pointing to a local CE interface. As with advertising direct routes on C2 customers' routing instance Configuring the `vrf-table-label` will advertise a common VRF label for all routes per routing instance routing

© 2020 Juniper Networks, Inc. All rights reserved.

table. This way R3 will perform a route lookup matching the static default route pointing to `inet.0`.

R3

```
[edit]
lab@R3# set routing-instances C1 vrf-table-label

[edit]
lab@R3# commit
commit complete
```

```
[edit]
lab@R8# run show route table C1.inet.0

C1.inet.0: 21 destinations, 48 routes (21 active, 0 holddown, 11 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[BGP/170] 00:00:05, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
                      to 172.30.0.45 via ge-0/0/4.178, label-switched-path
Bypass->172.30.0.9->172.30.0.2
172.30.5.17/32      *[BGP/170] 00:00:05, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                      to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
                      to 172.30.0.45 via ge-0/0/4.178, label-switched-path
Bypass->172.30.0.9->172.30.0.2
...
```

After adding `vrf-table-label`, the static default route is advertised to the other PE routers.

The last step is to advertise C1 prefixes to the Internet. You can create an aggregate route and advertise it to IBGP. The route reflector will advertise them to all external peers and customers.

```
[edit]
lab@R3# set routing-options aggregate route 172.31.48.0/20

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set policy-options policy-statement C1-inet-routes term 1 from protocol
aggregate

[edit]
lab@R3# set policy-options policy-statement C1-inet-routes term 1 from route-filter
172.31.48.0/20 exact

[edit]
lab@R3# set policy-options policy-statement C1-inet-routes term 1 then accept

[edit]
lab@R3# set protocols bgp group ibgp export C1-inet-routes

[edit]
lab@R3# commit
commit complete
```

Verify the advertisement of the aggregate route.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.41 172.31.48.0/20 exact

inet.0: 758 destinations, 760 routes (685 active, 0 holddown, 75 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.48.0/20          Self                         100        I
```

b) Configure R4 for Internet access.

The requirement is to provide redundancy to the Internet service. The configuration of R4 should be the same as of R3 to provide the desired redundancy.

There is already a rib-group applied, which copies the routes from the C1 routing instance from previous step. Adding rib `inet.0` to it, will copy the routes not only to the C2-spoke routing instance but also to the global routing table.

```
[edit]
lab@R4# delete routing-options rib-groups C1-C2-vpn

[edit]
lab@R4# set routing-options rib-groups C1-C2-vpn import-rib [ C1.inet.0 C2-spoke.inet.0
inet.0 ]

[edit]
lab@R4# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Modifying the policies and advertising the default and aggregate routes.

```
[edit]
lab@R4# set routing-instances C1 routing-options static route 0/0 next-table inet.0

[edit]
lab@R4# set policy-options policy-statement C1-default-to-ospf term 1 from protocol
static

[edit]
lab@R4# set policy-options policy-statement C1-default-to-ospf term 1 from route-filter
0/0 exact

[edit]
lab@R4# set policy-options policy-statement C1-default-to-ospf term 1 then accept

[edit]
lab@R4# set routing-instances C1 protocols ospf export C1-default-to-ospf
```

Modifying the route target export policy.

```
[edit]
lab@R4# delete policy-options policy-statement C1-vpn-export term 3

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 3 from protocol static

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 3 from route-filter 0/0
exact

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 3 then community set CE1

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 3 then accept

[edit]
lab@R4# set policy-options policy-statement C1-vpn-export term 4 then reject

[edit]
lab@R4# set routing-instances C1 vrf-table-label

[edit]
lab@R4# set routing-options aggregate route 172.31.48.0/20

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# set policy-options policy-statement C1-inet-routes term 1 from protocol
aggregate

[edit]
lab@R4# set policy-options policy-statement C1-inet-routes term 1 from route-filter
172.31.48.0/20 exact

[edit]
lab@R4# set policy-options policy-statement C1-inet-routes term 1 then accept

[edit]
lab@R4# set protocols bgp group ibgp export C1-inet-routes

[edit]
lab@R4# commit
commit complete
```

13) Configure Customer C2 internet access.

Customer C2 should have Internet access on the dedicated physical interface on R1 (192.168.0.50) and R2 (192.168.0.62). You have to configure an EBGP session with the CE router and advertise the default route.

```
[edit]
lab@R1# set protocols bgp group CE2 peer-as 64600 neighbor 192.168.0.50
```

Configure and apply the export policy to advertise the default route.

```
[edit]
lab@R1# set routing-options aggregate route 0/0

[edit]
lab@R1# set routing-options aggregate defaults brief

[edit]
lab@R1# set policy-options policy-statement default-only term 1 from protocol aggregate

[edit]
lab@R1# set policy-options policy-statement default-only term 1 from route-filter 0/0
exact

[edit]
lab@R1# set policy-options policy-statement default-only term 1 then accept

[edit]
lab@R1# set policy-options policy-statement default-only term 2 then reject

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set protocols bgp group CE2 type external export default-only

[edit]
lab@R1# commit
commit complete
```

Verify the advertisement of the default route.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.50

inet.0: 645 destinations, 1209 routes (477 active, 0 holddown, 330 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                                    ?
```

The same configuration steps have to be applied also to R2.

Task 2: Multicast in L3VPNs

The second task requires configuring general multicast and NG-MVPN multicast for L3VPN.

1) Configure PIM on all routers.

Protocol Independent Multicast (PIM) protocol should be configured on all core interfaces.

R1

```
[edit]
lab@R1# set protocols pim interface ge-0/0/4.114

[edit]
lab@R1# set protocols pim interface ge-0/0/4.118

[edit]
lab@R1# set protocols pim interface ae0.0

[edit]
lab@R1# set protocols pim interface lo0.0

[edit]
lab@R1# commit
commit complete
```

R2

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# set protocols pim interface ge-0/0/4.123

[edit]
lab@R2# set protocols pim interface ge-0/0/4.127

[edit]
lab@R2# set protocols pim interface ae0.0

[edit]
lab@R2# set protocols pim interface lo0.0

[edit]
lab@R2# commit
commit complete
```

R3

```
[edit]
lab@R3# set protocols pim interface ge-0/0/4.123

[edit]
lab@R3# set protocols pim interface ge-0/0/4.134

[edit]
lab@R3# set protocols pim interface ge-0/0/4.136

[edit]
lab@R3# set protocols pim interface lo0.0

[edit]
lab@R3# commit
commit complete
```

R4

```
[edit]
lab@R4# set protocols pim interface ge-0/0/4.114

[edit]
lab@R4# set protocols pim interface ge-0/0/4.134

[edit]
lab@R4# set protocols pim interface ge-0/0/4.145

[edit]
lab@R4# set protocols pim interface lo0.0

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# commit
commit complete
```

R5

```
[edit]
lab@R5# set protocols pim interface ge-0/0/4.145

[edit]
lab@R5# set protocols pim interface ge-0/0/4.158

[edit]
lab@R5# set protocols pim interface ae0.0

[edit]
lab@R5# set protocols pim interface lo0.0

[edit]
lab@R5# commit
commit complete
```

R6

```
[edit]
lab@R6# set protocols pim interface ge-0/0/4.136

[edit]
lab@R6# set protocols pim interface ge-0/0/4.167

[edit]
lab@R6# set protocols pim interface ae0.0

[edit]
lab@R6# set protocols pim interface lo0.0

[edit]
lab@R6# commit
commit complete
```

R7

```
[edit]
lab@R7# set protocols pim interface ge-0/0/4.127

[edit]
lab@R7# set protocols pim interface ge-0/0/4.167

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# set protocols pim interface ge-0/0/4.178

[edit]
lab@R7# set protocols pim interface lo0.0

[edit]
lab@R7# commit
commit complete
```

R8

```
[edit]
lab@R8# set protocols pim interface ge-0/0/4.118

[edit]
lab@R8# set protocols pim interface ge-0/0/4.158

[edit]
lab@R8# set protocols pim interface ge-0/0/4.178

[edit]
lab@R8# set protocols pim interface lo0.0

[edit]
lab@R8# commit
commit complete
```

By default, PIM is started in sparse mode. Verify the PIM neighbor table after PIM is enabled on all routers.

```
[edit]
lab@R1# run show pim interfaces
Instance: PIM.master

Name          Stat Mode    IP V State NbrCnt JoinCnt(sg) JoinCnt(*g) DR address
ae0.0         Up   Sparse  4 2 NotDR      1            0            0 172.30.0.2
ge-0/0/4.114  Up   Sparse  4 2 NotDR      1            0            0 172.30.0.6
ge-0/0/4.118  Up   Sparse  4 2 NotDR      1            0            0 172.30.0.10
lo0.0         Up   Sparse  4 2 DR         0            0            0 172.30.5.1
ppd0.32769    Up   Sparse  4 2 P2P        0            0            0
ae0.0         Up   Sparse  6 2 NotDR      1            0            0
fe80::fac0:1ff:fedd:280
ge-0/0/4.118  Up   Sparse  6 2 NotDR      1            0            0
fe80::fac0:100:76dc:2e84
lo0.0         Up   Sparse  6 2 DR         0            0            0
fe80::226:880f:fcef:7580
```

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# run show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Interface          IP V Mode        Option      Uptime Neighbor addr
ae0.0              4 2              HPLGT       04:37:52 172.30.0.2
ge-0/0/4.114       4 2              HPLGT       04:37:40 172.30.0.6
ge-0/0/4.118       4 2              HPLGT       04:37:18 172.30.0.10
ae0.0              6 2              HPLGT       04:37:52 fe80::fac0:1ff:fedd:280
ge-0/0/4.118       6 2              HPLGT       04:37:18 fe80::fac0:100:76dc:2e84
```

2) Configure PIM to use the `inet.2` table.

PIM sparse mode performs Reverse Path Forwarding (RPF) checks of the source of the stream. By default, `inet.0` is used for that.

```
[edit]
lab@R1# run show multicast rpf summary
Multicast RPF table: inet.0 , 844 entries

Multicast RPF table: inet6.0 , 82 entries
```

If you want to have different topology for multicast RFP checks, you can use the `inet.2` table. BGP and IS-IS can directly place routes into the `inet.2` table. For other routing protocols rib-groups must be used. Rib-groups are used also to instruct multicast routing protocols to use the `inet.2` for RFP checks.

Let's modify the existing rib-group applied to IS-IS and apply interface routes also.

Repeat this on all eight devices.

```
[edit]
lab@R1# set routing-options rib-groups rr-inet0-inet3 import-rib [ inet.0 inet.3 inet.2
]

[edit]
lab@R1# set routing-options interface-routes rib-group inet rr-inet0-inet3
```

Configure a new rib-group, which imports `inet.2` and ensure the PIM protocol uses this rib-group.

```
[edit]
lab@R1# set routing-options rib-groups mcast-rib import-rib inet.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols pim rib-group inet mcast-rib

[edit]
lab@R1# commit
commit complete
```

The rib-group for PIM is configured with single group `inet.2` as PIM only uses the first import rib for RFP checks.

Verify if PIM uses the `inet.2` table for RFP checks.

```
[edit]
lab@R1# run show multicast rpf summary
Multicast RPF table: inet.2 , 57 entries

Multicast RPF table: inet6.0 , 82 entries
```

`Inet.2` is populated with copied routes.

```
[edit]
lab@R1# run show route table inet.2

inet.2: 73 destinations, 73 routes (72 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.1.0/24       *[Direct/0] 00:06:09
                    > via fxp0.0
10.10.1.1/32       *[Local/0] 00:06:09
                       Local via fxp0.0
172.30.0.0/30      *[Direct/0] 00:00:40
                    > via ae0.0
172.30.0.1/32      *[Local/0] 00:00:40
                       Local via ae0.0
172.30.0.4/30      *[Direct/0] 00:06:09
                    > via ge-0/0/4.114
172.30.0.5/32      *[Local/0] 00:06:09
                       Local via ge-0/0/4.114
172.30.0.8/30      *[Direct/0] 00:06:09
                    > via ge-0/0/4.118
172.30.0.9/32      *[Local/0] 00:06:09
                       Local via ge-0/0/4.118
172.30.0.12/30     *[IS-IS/15] 00:00:20, metric 15
                    > to 172.30.0.2 via ae0.0
172.30.0.16/30     *[IS-IS/15] 00:00:20, metric 15
                    > to 172.30.0.2 via ae0.0
172.30.0.20/30     *[IS-IS/15] 00:06:09, metric 20
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.6 via ge-0/0/4.114
```

This should be configured on all routers.

3) Configure a PIM RP on R1 and R2.

R1 and R2 are dedicated Rendezvous Points (RPs) for the provider PIM domain. They should be configured for anycast-RP. This means that both PEs share a common loopback IP address, which is advertised as the RP. The mechanism used for discovery of the RP address is Bootstrap router (BSR). R1 and R2 are both configured as Bootstrap routers, where R1 has a higher priority configured.

a) Configure RP on R1.

```
[edit]
lab@R1# set interfaces lo0.0 family inet address 172.30.5.254/32

[edit]
lab@R1# set protocols pim rp bootstrap family inet priority 200

[edit]
lab@R1# set protocols pim rp local family inet address 172.30.5.254

[edit]
lab@R1# set protocols pim rp local family inet anycast-pim rp-set address 172.30.5.2

[edit]
lab@R1# set protocols pim rp local family inet anycast-pim local-address 172.30.5.1

[edit]
lab@R1# commit
commit complete
```

The IP address 172.30.5.254 is used as the anycast-RP. The local IP address is R1's loopback IP address. R2's loopback IP address is configured as RP set address announced by the Bootstrap router.

Verify the RP table and bootstrap table.

```
[edit]
lab@R1# run show pim rps
Instance: PIM.master
Address family INET
RP address             Type       Holdtime Timeout Groups Group prefixes
172.30.5.254           bootstrap       150     117      0 224.0.0.0/4
172.30.5.254           static            0    None      0 224.0.0.0/4
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show pim bootstrap
Instance: PIM.master

BSR                  Pri Local address          Pri State      Timeout
172.30.5.1           200 172.30.5.1             200 Elected          3
None                   0 fd17:f0f4:f691:5::1       0 InEligible Infinity
```

Local router is elected as the Bootstrap router.

b) Configure RP on R2.

```
[edit]
lab@R2# set interfaces lo0.0 family inet address 172.30.5.254/32

[edit]
lab@R2# set protocols pim rp bootstrap family inet priority 100

[edit]
lab@R2# set protocols pim rp local family inet address 172.30.5.254

[edit]
lab@R2# set protocols pim rp local family inet anycast-pim rp-set address 172.30.5.1

[edit]
lab@R2# set protocols pim rp local family inet anycast-pim local-address 172.30.5.2

[edit]
lab@R2# commit
commit complete
```

The same loopback IP address is configured on R2 also and is used for anycast-RP. The bootstrap priority is lower than R1.

```
[edit]
lab@R2# run show pim rps
Instance: PIM.master
Address family INET
RP address           Type       Holdtime Timeout Groups Group prefixes
172.30.5.254         bootstrap       150      97      0 224.0.0.0/4
172.30.5.254         static          150    None      0 224.0.0.0/4
```

```
[edit]
lab@R2# run show pim bootstrap
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Instance: PIM.master

BSR                     Pri Local address           Pri State       Timeout
172.30.5.1              200 172.30.5.2              100 Candidate        46
None                      0 fd17:f0f4:f691:5::2       0 InEligible Infinity
```

The elected Bootstrap router is R1, whereas R2 is a candidate to become the Bootstrap router.

4) Configure multicast in customer C2 VPN.

Customer C2 multicast is implemented using NG MVPN. This type of Multicast VPN implementation does not use Provider PIM for control plane information exchange, but instead uses BGP. The MCAST-VPN NLRI (AFI1/ SAFI 5) carries new attributes to signal Customer PIM information and data plane setup operations across the core. The provider multicast service interface (PMSI) is an BGP optional transitive attribute, which identifies the tunnel that carries the multicast traffic in a particular customer MVPN. NG MVPN routes uses the following BGP routing tables: L3VPN multicast routes are placed in the `bgp.mvpn.0` table and specific `<VPN>.mvpn.0 table`.

Seven MVPN NLRIs are defined to support NG MVPNs.

- *Type 1: Intra AS Inclusive MVPN Membership Discovery – send from all PE routers. These routes carry PMSI Tunnel information.*
- *Type 2: Inter AS Inclusive MVPN Membership Discovery – used for membership discovery between different ASs.*
- *Type 3: Selective MVPN Autodiscovery route – send by PE that initiates a selective PMSI.*
- *Type 4: Selective MVPN Autodiscovery route for Leaf – send by receiver PE in request of Type 3 (with leaf information bit set).*
- *Type 5: Source active autodiscovery route –send by PE router discovers an active multicast source. The actual source can be discovered by MSDP, locally attached source or through PIM RP register message.*
- *Type 6: Shared tree join route – send by receiver PE that receives a PIM join message on its VRF interface.*
- *Type 7: Source tree join route – send by the receiver PE that received a source specific PIM join.*

a) Configure NG MVPN on R1 and R2.

First configure PIM in the customer spoke instance.

```
[edit]
lab@R1# set routing-instances C2-spoke protocols pim interface all
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set routing-instances C2-spoke protocols pim rp local address 172.30.5.253

[edit]
lab@R1# set routing-instances C2-spoke protocols pim rp local group-ranges 239.0.0.0/24

[edit]
lab@R1# commit
commit complete
```

A static local RP is configured on R1. The address for the RP is shared with R2 and is advertised to the CE with a different MED value via EBGP session.

```
[edit]
lab@R1# set interfaces lo0.2 family inet address 172.30.5.253/32

[edit]
lab@R1# set interfaces lo0.2 family inet address 172.30.5.10/32 primary
```

Enable tunnel services interface to allow the RP to de-encapsulate PIM packets.

```
[edit]
lab@R1# set chassis fpc 0 pic 0 tunnel-services
```

The shared IP address is configured on the loopback interface as secondary.

```
[edit]
lab@R1# set policy-options policy-statement C2-direct-routes term 1 from protocol
direct

[edit]
lab@R1# set policy-options policy-statement C2-direct-routes term 1 from route-filter
172.30.5.253/32 exact

[edit]
lab@R1# set policy-options policy-statement C2-direct-routes term 1 then metric 10

[edit]
lab@R1# set policy-options policy-statement C2-direct-routes term 1 then accept

[edit]
lab@R1# set policy-options policy-statement C2-direct-routes term 2 from protocol
direct

[edit]
lab@R1# set policy-options policy-statement C2-direct-routes term 2 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set routing-instances C2-spoke protocols bgp group ce export C2-direct-routes

[edit]
lab@R1# commit
commit complete
```

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.46

C2-spoke.inet.0: 19 destinations, 23 routes (19 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.5.10/32          Self                                    I
* 172.30.5.22/32          Self                                    I
* 172.30.5.25/32          Self                                    I
* 172.30.5.33/32          Self                                    I
* 172.30.5.253/32         Self                 10                 I
* 172.31.72.0/24          Self                                    54591 I
* 172.31.73.0/24          Self                                    54591 I
* 172.31.74.0/24          Self                                    54591 I
* 172.31.75.0/24          Self                                    54591 I
* 172.31.76.0/24          Self                                    54591 I
* 172.31.77.0/24          Self                                    54591 I
* 192.168.0.44/30         Self                                    I
* 192.168.0.76/30         Self                                    I
* 192.168.0.80/30         Self                                    I
* 192.168.0.88/30         Self                                    I
```

On R1 configure `inet-mvpn signaling` on the IBGP session with the Route.

```
[edit]
lab@R1# set protocols bgp group ibgp family inet-mvpn signaling

[edit]
lab@R1# commit
commit complete
```

Configure `inet-mvpn signaling` on the route reflectors IBGP groups.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family inet-mvpn signaling

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family inet-mvpn signaling

[edit]
lab@route-reflector# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verifying the rt-import community that should be advertised by the C2-hub routing-instance.

```
[edit]
lab@R2# run show policy __vrf-mvpn-import-cmcast-C2-spoke-internal__
Policy __vrf-mvpn-import-cmcast-C2-spoke-internal__:
    Term unnamed:
        from community __vrf-mvpn-community-rt_import-target-C2-spoke-internal__
[target: 172.30.5.2:8]
        then accept
    Term unnamed:
        then reject
```

```
[edit]
lab@R2# set policy-options community src-as members src-as:54591:0

[edit]
lab@R2# set policy-options community rt-import-spoke members rt-import:172.30.5.2:8

[edit]
lab@R2# set policy-options policy-statement C2-hub-export term 1 then community add rt-
import-spoke

[edit]
lab@R2# set policy-options policy-statement C2-hub-export term 1 then community add
src-as
```

NG MVPN supports different provider tunnels. Customer C2 must use a point-to-multipoint LSPs build with RSVP. The advantage of using a point-to-multipoint LSP is that there will be not traffic replication in the MPLS network for multicast traffic. Remember that the main advantage of multicast traffic is that traffic to multiple destinations is only send once on each link. The problem with MPLS is that each receiving PE in the network can be reached on a different LSP. The LSPs might follow the same physical link. So, if a source PE wants to send multicast traffic to 10 different receiving PE's, the multicast traffic is still replicated 10 times if the LSPs towards these 10 PEs transit the same link. When you use a point-to-multipoint LSP for L3 MVPNs or VPLS, then Broadcast, Unknown Unicast and Multicast traffic can be sent without replication. Please note that unicast traffic still uses a regular point-to-point LSP.

Configure P2MP RSVP LSP.

```
[edit]
lab@R1# set protocols mpls label-switched-path mcast-p2mp-template template

[edit]
lab@R1# set protocols mpls label-switched-path mcast-p2mp-template bandwidth 30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols mpls label-switched-path mcast-p2mp-template hop-limit 5

[edit]
lab@R1# set protocols mpls label-switched-path mcast-p2mp-template priority 5 5

[edit]
lab@R1# set protocols mpls label-switched-path mcast-p2mp-template link-protection

[edit]
lab@R1# set protocols mpls label-switched-path mcast-p2mp-template p2mp
```

The RSVP P2MP LSP is configured with link-protection, bandwidth and priority. It is defined as a dynamic template.

Configure the MVPN protocol under the routing instance.

```
[edit]
lab@R1# set routing-instances C2-spoke protocols mvpn mvpn-mode spt-only

[edit]
lab@R1# set routing-instances C2-spoke protocols mvpn sender-site

[edit]
lab@R1# set routing-instances C2-spoke protocols mvpn route-target import-target target
target:54591:202

[edit]
lab@R1# set routing-instances C2-spoke protocols mvpn route-target export-target target
target:54591:202
```

MVPN mode `spt-only` is the default, so it does not have to be explicitly configured. The `spt-only` mode ensures that shortest path tree multicast must be used in the MPLS core. Sources are learned from multicast VPN Source-Active routes. The site is designated as sender-site only. Sender-only sites have PEs that do not join the P-tunnels advertised by the remote PEs. Different route targets are configured for the MVPN routes.

```
[edit]
lab@R1# set routing-instances C2-spoke provider-tunnel rsvp-te label-switched-path-
template mcast-p2mp-template

[edit]
lab@R1# commit
commit complete
```

The dynamic template (P2MP RSVP LSP) is configured as Inclusive P-tunnel.

© 2020 Juniper Networks, Inc. All rights reserved.

b) Configure NG MVPN on R2.

R2 configuration is the same as on R1 (redundancy).

```
[edit]
lab@R2# set interfaces lo0.2 family inet address 172.30.5.253/32

[edit]
lab@R2# set interfaces lo0.2 family inet address 172.30.5.14/32 primary

[edit]
lab@R2# set chassis fpc 0 pic 0 tunnel-services

[edit]
lab@R2# set policy-options policy-statement C2-direct-routes term 1 from protocol
direct

[edit]
lab@R2# set policy-options policy-statement C2-direct-routes term 1 from route-filter
172.30.5.253/32 exact

[edit]
lab@R2# set policy-options policy-statement C2-direct-routes term 1 then metric 100

[edit]
lab@R2# set policy-options policy-statement C2-direct-routes term 1 then accept

[edit]
lab@R2# set policy-options policy-statement C2-direct-routes term 2 from protocol
direct

[edit]
lab@R2# set policy-options policy-statement C2-direct-routes term 2 then accept

[edit]
lab@R2# set routing-instances C2-spoke protocols bgp group ce export C2-direct-routes

[edit]
lab@R2# commit
commit complete
```

The shared IP address used for the RP is configured as a secondary loopback interface address. It is advertised to the CE router with higher metric of 100.

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.0.58 table C2-
spoke.inet.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
C2-spoke.inet.0: 19 destinations, 23 routes (19 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.5.14/32          Self                                    I
* 172.30.5.22/32          Self                                    I
* 172.30.5.25/32          Self                                    I
* 172.30.5.33/32          Self                                    I
* 172.30.5.253/32         Self                 100                I
* 172.31.72.0/24          Self                                    54591 I
* 172.31.73.0/24          Self                                    54591 I
* 172.31.74.0/24          Self                                    54591 I
* 172.31.75.0/24          Self                                    54591 I
* 172.31.76.0/24          Self                                    54591 I
* 172.31.77.0/24          Self                                    54591 I
* 192.168.0.56/30         Self                                    I
* 192.168.0.76/30         Self                                    I
* 192.168.0.80/30         Self                                    I
* 192.168.0.88/30         Self                                    I
```

The route of shared IP address is advertised with metric of 100 to the CE.

Configure PIM for routing-instance C2-spoke.

```
[edit]
lab@R2# set routing-instances C2-spoke protocols pim interface all

[edit]
lab@R2# set routing-instances C2-spoke protocols pim rp local address 172.30.5.253

[edit]
lab@R2# set routing-instances C2-spoke protocols pim rp local group-ranges 239.0.0.0/24

[edit]
lab@R2# commit
commit complete
```

R2 is also statically configured to act as a local RP using the shared address. The CEs have the shared RP address statically configured. BGP announces the shared RP address, which is more or less emulating Anycast RP functionality for the customer CEs.

```
[edit]
lab@R2# set protocols mpls label-switched-path mcast-p2mp-template template

[edit]
lab@R2# set protocols mpls label-switched-path mcast-p2mp-template bandwidth 30

[edit]
lab@R2# set protocols mpls label-switched-path mcast-p2mp-template hop-limit 5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set protocols mpls label-switched-path mcast-p2mp-template priority 5 5

[edit]
lab@R2# set protocols mpls label-switched-path mcast-p2mp-template link-protection

[edit]
lab@R2# set protocols mpls label-switched-path mcast-p2mp-template p2mp
```

Configure protocols mvpn for routing instance C2-spoke.

```
[edit]
lab@R2# set routing-instances C2-spoke protocols mvpn sender-site

[edit]
lab@R2# set routing-instances C2-spoke protocols mvpn mvpn-mode spt-only

[edit]
lab@R2# set routing-instances C2-spoke protocols mvpn route-target import-target target
target:54591:202

[edit]
lab@R2# set routing-instances C2-spoke protocols mvpn route-target export-target target
target:54591:202

[edit]
lab@R2# set routing-instances C2-spoke provider-tunnel rsvp-te label-switched-path-
template mcast-p2mp-template

[edit]
lab@R2# commit
commit complete
```

c) Configure NG MVPN support on R4.

```
[edit]
lab@R4# set protocols bgp group ibgp family inet-mvpn signaling

[edit]
lab@R4# set routing-instances C2-spoke protocols pim interface all

[edit]
lab@R4# set routing-instances C2-spoke protocols mvpn receiver-site

[edit]
lab@R4# set routing-instances C2-spoke protocols mvpn mvpn-mode spt-only

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# set routing-instances C2-spoke protocols mvpn route-target import-target target
target:54591:202

[edit]
lab@R4# set routing-instances C2-spoke protocols mvpn route-target export-target target
target:54591:202

[edit]
lab@R4# commit
commit complete
```

R4 is configured as a receiver-site only. PEs designated as a receiver-site only, do not send the PMSI Tunnels attribute. Therefore, no configuration for a provider tunnel under the routing instance is needed.

d) Configure NG MVPN on R5.

Configuration on R5 is the same as R4.

```
[edit]
lab@R5# set protocols bgp group ibgp family inet-mvpn signaling

[edit]
lab@R5# set routing-instances C2-spoke protocols pim interface all

[edit]
lab@R5# set routing-instances C2-spoke protocols mvpn mvpn-mode spt-only

[edit]
lab@R5# set routing-instances C2-spoke protocols mvpn receiver-site

[edit]
lab@R5# set routing-instances C2-spoke protocols mvpn route-target import-target target
target:54591:202

[edit]
lab@R5# set routing-instances C2-spoke protocols mvpn route-target export-target target
target:54591:202

[edit]
lab@R5# commit
commit complete
```

e) Configure NG MVPN on R7.

R7 is also configured as a receiver site only. The configuration is the same as on R4 and R5.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols bgp group ibgp family inet-mvpn signaling

[edit]
lab@R7# set routing-instances C2-spoke protocols pim interface all

[edit]
lab@R7# set routing-instances C2-spoke protocols mvpn mvpn-mode spt-only

[edit]
lab@R7# set routing-instances C2-spoke protocols mvpn receiver-site

[edit]
lab@R7# set routing-instances C2-spoke protocols mvpn route-target import-target target
target:54591:202

[edit]
lab@R7# set routing-instances C2-spoke protocols mvpn route-target export-target target
target:54591:202

[edit]
lab@R7# commit
commit complete
```

Verify the operation of NG MVPN for customer C2 on R1.

```
[edit]
lab@R1# run show pim neighbors instance C2-spoke
Instance: PIM.C2-spoke
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Interface           IP V Mode        Option      Uptime Neighbor addr
ge-0/0/5.312         4 2             HPLGT      05:20:43 192.168.0.46
```

One PIM neighbor is established on the CE interface.

```
[edit]
lab@R1# run show pim interfaces instance C2-spoke
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, DDR = Dual DR, DistDR = Distributed DR,
P2P = Point-to-point link, P2MP = Point-to-Multipoint,
Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name                Stat Mode IP V State        NbrCnt JoinCnt(sg/*g) DR address
ge-0/0/5.312        Up   S     4 2 NotDR,NotCap      1 0/0             192.168.0.46
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lo0.2                Up    S      4 2 DR,NotCap          0 0/0              172.30.5.10
lsi.257              Up    SD     4 2 P2P,NotCap         0 0/0
pd-0/0/0.32770       Up    S      4 2 P2P,NotCap         0 0/0
lsi.257              Up    SD     6 2 P2P,NotCap         0 0/0
```

A virtual LSI interface is created because of the `vrf-table-label` option under the routing instance. The `ppd0` or `pd-0/0/0` interface is a PIM de-encapsulation interface. The interface created depends on the platform being used. The local router is the RP and needs to de-encapsulate the register messages from the sources.

```
[edit]
lab@R1# run show pim rps instance C2-spoke
Instance: PIM.C2-spoke
Address family INET
RP address          Type        Holdtime Timeout Groups Group prefixes
172.30.5.253        static           150    None      0 239.0.0.0/24

Address family INET6
```

Listing the BGP neighbors with all the NLRI/ families enabled.

```
[edit]
lab@R1# run show bgp summary
Groups: 5 Peers: 6 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    1344        599          0          0          0          0
bgp.l3vpn.0
                      16         16          0          0          0          0
inet6.0
                      48         48          0          0          0          0
bgp.mvpn.0
                       3          3          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591       4958       7003       0       0     5:20:16
Establ
  inet.0: 216/216/216/0
  bgp.l3vpn.0: 16/16/16/0
  inet6.0: 48/48/48/0
  bgp.rtarget.0: 8/12/12/0
  bgp.mvpn.0: 3/3/3/0
  C2-spoke.inet.0: 12/16/16/0
  C2-spoke.mvpn.0: 3/3/3/0
182.168.0.50          64600          0          0       0       0     5:22:27
Idle
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
192.168.0.42            64600       4437        708        0        0     5:21:14
Establ
  C2-hub.inet.0: 24/24/24/0
192.168.0.46            64600        686       1872        0        0     5:21:13
Establ
  C2-spoke.inet.0: 0/0/0/0
192.168.1.3              1620       1154        838        0        0     5:21:05
Establ
  inet.0: 0/564/402/0
192.168.1.4              1620       1186        840        0        0     5:21:15
Establ
  inet.0: 383/564/402/0
```

As mentioned, `bgp.mvpn.0` is where the `family mvpn` routes received from other PEs are stored. It has the same role as the `bgp.l3vpn.0` for L3VPN services. From there the routes are inserted into the correct routing instances tables. In this case table `C2-spoke.mvpn.0`.

```
[edit]
lab@R1# run show route table C2-spoke.mvpn.0

C2-spoke.mvpn.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:172.30.5.1:32767:172.30.5.1/240
                   *[MVPN/70] 05:54:39, metric2 1
                      Indirect
1:172.30.5.2:5:172.30.5.2/240
                   *[BGP/170] 00:00:23, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.2 via ae0.0, Push 0
1:172.30.5.4:5:172.30.5.4/240
                   *[BGP/170] 00:00:22, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.6 via ge-0/0/4.114, Push 0
1:172.30.5.5:4:172.30.5.5/240
                   *[BGP/170] 01:02:43, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.6 via ge-0/0/4.114, label-switched-path R1-to-
R8
                      to 172.30.0.10 via ge-0/0/4.118, label-switched-path R1-to-
R8
1:172.30.5.7:4:172.30.5.7/240
                   *[BGP/170] 00:01:23, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.2 via ae0.0, Push 303536
```


© 2020 Juniper Networks, Inc. All rights reserved.

Extensive output of the local originated Type 1 route shows the signaled PMSI Tunnel attribute.

```
[edit]
lab@R1# run show route table C2-spoke.mvpn.0 extensive

C2-spoke.mvpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
1:172.30.5.1:32767:172.30.5.1/240 (1 entry, 1 announced)
TSI:
RTF { 54591:202/64 }
Page 0 idx 0 Type 1 val 1853b9c
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [54591] I
    Communities: target:54591:202
    PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[172.30.5.1:0:2498:172.30.5.1]Session_13[172.30.5.1:0:2498:172.30.5.1]
Path 1:172.30.5.1:32767:172.30.5.1 Vector len 4.  Val: 0
        *MVPN   Preference: 70
                Next hop type: Indirect
                Address: 0x155d064
                Next-hop reference count: 3
                Protocol next hop: 172.30.5.1
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 5:23:34    Metric2: 1
                Task: mvpn global task
                Announcement bits (4): 0-PIM.C2-spoke 1-mvpn global task 2-BGP
Route Target 3-BGP_RT_Background
                AS path: I
...
```

Verify the P2MP RSVP sessions.

```
[edit]
lab@R1# run show rsvp session p2mp
Ingress RSVP: 15 sessions
P2MP name: 172.30.5.1:32767:mvpn:C2-spoke, P2MP branch count: 4
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.1      Up       0  1 SE       -   304368
172.30.5.7:172.30.5.1:32767:mvpn:C2-spoke
172.30.5.5      172.30.5.1      Up       0  1 SE       -   324896
172.30.5.5:172.30.5.1:32767:mvpn:C2-spoke
172.30.5.2      172.30.5.1      Dn       0  0  -       -        -
172.30.5.2:172.30.5.1:32767:mvpn:C2-spoke
172.30.5.4      172.30.5.1      Up       0  1 SE       -       17
172.30.5.4:172.30.5.1:32767:mvpn:C2-spoke
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Total 4 displayed, Up 3, Down 1

Egress RSVP: 13 sessions, 6 detours
P2MP name: 172.30.5.2:5:mvpn:C2-spoke, P2MP branch count: 3
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.2      Up       0  1 SE       3         -
172.30.5.1:172.30.5.2:5:mvpn:C2-spoke
172.30.5.1      172.30.5.2      Up       0  1 SE       3         -
172.30.5.1:172.30.5.2:5:mvpn:C2-spoke
Total 2 displayed, Up 2, Down 0

Transit RSVP: 24 sessions
P2MP name: 172.30.5.2:5:mvpn:C2-spoke, P2MP branch count: 3
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.5      172.30.5.2      Up       0  1 SE  306032   325152
172.30.5.5:172.30.5.2:5:mvpn:C2-spoke
Total 1 displayed, Up 1, Down 0
```

The LSP to R2 is down, because R1 is configured as sender-only site and does not join P-tunnels advertised to other PE routers.

```
[edit]
lab@R1# run show mvpn instance C2-spoke

MVPN instance:
Legend for provider tunnel
S-    Selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Family : INET

Instance : C2-spoke
  MVPN Mode : SPT-ONLY
  Provider tunnel: I-P-tnl:RSVP-TE P2MP:172.30.5.1, 2498,172.30.5.1
  Neighbor                                Inclusive Provider Tunnel
  172.30.5.2
  172.30.5.4
  172.30.5.5
  172.30.5.7
...
```

The output shown above displays the MVPN instances, Inclusive P-tunnel and the neighbors.

Verify the operation of the NG MVPN on the receiver site R4.

© 2020 Juniper Networks, Inc. All rights reserved.

R4

```
[edit]
lab@R4# run show pim neighbors instance C2-spoke
B = Bidirectional Capable, G = Generation Identifier
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit,
A = Hello Option Join Attribute

Instance: PIM.C2-spoke
Interface            IP V Mode        Option       Uptime Neighbor addr
ge-0/0/5.320          4 2             HPLGTA      00:18:38 192.168.0.78
```

One PIM neighbor is established on the CE interface.

```
[edit]
lab@R4# run show pim rps instance C2-spoke
Instance: PIM.C2-spoke
Address family INET

Address family INET6
```

No RPs are known (and not configured statically) because NG MVPN is configured in spt-only mode.

```
[edit]
lab@R4# run show route table C2-spoke.mvpn.0

C2-spoke.mvpn.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
1:172.30.5.1:32767:172.30.5.1/240
                   *[BGP/170] 01:02:39, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.5 via ge-0/0/4.114, Push 0
1:172.30.5.2:5:172.30.5.2/240
                   *[BGP/170] 00:00:19, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.5 via ge-0/0/4.114, Push 304608
1:172.30.5.4:5:172.30.5.4/240
                   *[MVPN/70] 01:19:44, metric2 1
                      Indirect
1:172.30.5.5:4:172.30.5.5/240
                   *[BGP/170] 01:17:38, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.21 via ge-0/0/4.134, label-switched-path R4-to-
R5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                        to 172.30.0.30 via ge-0/0/4.145, label-switched-path
Bypass->172.30.0.21->172.30.0.26
1:172.30.5.7:4:172.30.5.7/240
                    *[BGP/170] 00:01:19, localpref 100, from 172.30.5.41
                       AS path: I
                       to 172.30.0.21 via ge-0/0/4.134, label-switched-path R4-to-
R7-1
                     > to 172.30.0.30 via ge-0/0/4.145, label-switched-path R4-to-
R7-2
                       to 172.30.0.30 via ge-0/0/4.145, label-switched-path
Bypass->172.30.0.21
```

Type 1 routes are received from all PE having routing instances in that VPN.

```
[edit]
lab@R4# run show rsvp session p2mp
Ingress RSVP: 8 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 12 sessions
P2MP name: 172.30.5.1:32767:mvpn:C2-spoke, P2MP branch count: 2
To          From            State   Rt Style Labelin Labelout LSPname
172.30.5.4  172.30.5.1      Up       0  1 SE      17        -
172.30.5.4:172.30.5.1:32767:mvpn:C2-spoke
P2MP name: 172.30.5.2:5:mvpn:C2-spoke, P2MP branch count: 2
To          From            State   Rt Style Labelin Labelout LSPname
172.30.5.4  172.30.5.2      Up       0  1 SE      17        -
172.30.5.4:172.30.5.2:5:mvpn:C2-spoke
Total 2 displayed, Up 2, Down 0

Transit RSVP: 26 sessions
P2MP name: 172.30.5.1:32767:mvpn:C2-spoke, P2MP branch count: 2
To          From            State   Rt Style Labelin Labelout LSPname
172.30.5.5  172.30.5.1      Up       0  1 SE  324896       16
172.30.5.5:172.30.5.1:32767:mvpn:C2-spoke
P2MP name: 172.30.5.2:5:mvpn:C2-spoke, P2MP branch count: 2
To          From            State   Rt Style Labelin Labelout LSPname
172.30.5.5  172.30.5.2      Up       0  1 SE  325168       16
172.30.5.5:172.30.5.2:5:mvpn:C2-spoke
Total 2 displayed, Up 2, Down 0
```

Two P2MP RSVP Egress sessions are created, one to R1 and one to R2.

```
[edit]
lab@R4# run show mvpn instance C2-spoke

MVPN instance:
Legend for provider tunnel
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
S-    Selective provider tunnel
F-    Flood NH forwarding NH
M-    Multicast Composite NH
C-    Cloned NH

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
I -- Inactive
Family : INET

Instance : C2-spoke
  MVPN Mode : SPT-ONLY
  Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Provider tunnel: I-P-tnl:invalid
  Neighbor                                Inclusive Provider Tunnel
  172.30.5.1                              RSVP-TE P2MP:172.30.5.1, 44971,172.30.5.1
  172.30.5.2                              RSVP-TE P2MP:172.30.5.2, 63509,172.30.5.2
  172.30.5.5
  172.30.5.7

MVPN instance:
Legend for provider tunnel
S-    Selective provider tunnel
F-    Flood NH forwarding NH
M-    Multicast Composite NH
C-    Cloned NH

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
I -- Inactive
Family : INET6

Instance : C2-spoke
  MVPN Mode : SPT-ONLY
  Sender-Based RPF: Disabled. Reason: Not enabled by configuration.
  Hot Root Standby: Disabled. Reason: Not enabled by configuration.
  Provider tunnel: I-P-tnl:invalid
```

The output above displays the MVPN instances, Inclusive P-tunnel and the neighbors.

5) Configure Selective provider tunnels on R1 and R2.

The selective P-tunnel should be created for groups 239.0.0.1 and 239.0.0.1 when the traffic reaches 100Mbps. The P2MP LSP parameters are defined in table 20.

```
[edit]
lab@R1# set protocols mpls label-switched-path mcast-selective-template template
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols mpls label-switched-path mcast-selective-template bandwidth 60m

[edit]
lab@R1# set protocols mpls label-switched-path mcast-selective-template hop-limit 5

[edit]
lab@R1# set protocols mpls label-switched-path mcast-selective-template priority 5 5

[edit]
lab@R1# set protocols mpls label-switched-path mcast-selective-template link-protection

[edit]
lab@R1# set protocols mpls label-switched-path mcast-selective-template p2mp

[edit]
lab@R1# set routing-instances C2-spoke provider-tunnel selective tunnel-limit 5

[edit]
lab@R1# set routing-instances C2-spoke provider-tunnel selective group 239.0.0.1/32
source 172.31.64.0/21 threshold-rate 100000

[edit]
lab@R1# set routing-instances C2-spoke provider-tunnel selective group 239.0.0.1/32
source 172.31.64.0/21 rsvp-te label-switched-path-template mcast-selective-template

[edit]
lab@R1# set routing-instances C2-spoke provider-tunnel selective group 239.0.0.2/32
source 172.31.64.0/21 threshold-rate 100000

[edit]
lab@R1# set routing-instances C2-spoke provider-tunnel selective group 239.0.0.2/32
source 172.31.64.0/21 rsvp-te label-switched-path-template mcast-selective-template

[edit]
lab@R1# commit
commit complete
```

R2 should be configured with same configuration for redundancy.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 3: IPv6 Tunneling with 6VPE

The 6VPE model for tunneling IPv6 traffic over an MPLS network uses a similar mechanism as used for IPv4 L3VPNs. Like 6PE, the LSPs are signaled with an IPv4 signaling protocol. The IPv6 CE interfaces are configured inside VRFs. Multiprotocol BGP is used to advertise the IPv6 routes to other PE routers. You can have IPv4 and IPv6 support within a single VRF.

The `ipv6-tunneling` option under `protocols mpls` hierarchy is also needed for 6VPE to copy the LSP routes from `inet.3` to `inet6.3`

In the previous chapter `ipv6-tunneling` was already enabled on all the routers.

```
[edit]
lab@R3# run show route table inet6.3

inet6.3: 8 destinations, 10 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::ffff:172.30.5.1/128
                   *[LDP/9] 00:10:13, metric 15
                    > to 172.30.0.13 via ge-0/0/4.123, Push 303712
::ffff:172.30.5.2/128
                   *[LDP/9] 00:10:07, metric 10
                    > to 172.30.0.13 via ge-0/0/4.123, Push 0
::ffff:172.30.5.4/128
                   *[LDP/9] 00:10:13, metric 10
                    > to 172.30.0.22 via ge-0/0/4.134, Push 0
::ffff:172.30.5.5/128
                   *[LDP/9] 00:08:55, metric 15
                    > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-
R6
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
::ffff:172.30.5.6/128
                   *[RSVP/7/1] 00:10:08, metric 10
                    > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-
R6
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path
Bypass->172.30.0.22
                    [LDP/9] 00:08:55, metric 10
                    > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-
R6
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
::ffff:172.30.5.7/128
                   *[LDP/9] 00:08:55, metric 20
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.22 via ge-0/0/4.134, label-switched-path R3-to-
R6
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
::ffff:172.30.5.8/128
                    *[RSVP/7/1] 00:10:06, metric 25
                     > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-
R8-2
                      to 172.30.0.22 via ge-0/0/4.134, label-switched-path
Bypass->172.30.0.13->172.30.0.1
                    [LDP/9] 00:07:32, metric 25
                     > to 172.30.0.13 via ge-0/0/4.123, Push 304112
::ffff:192.168.1.0/120
                    *[LDP/9] 00:10:07, metric 20
                     > to 172.30.0.13 via ge-0/0/4.123, Push 0
```

The `inet6.3` routing table is populated with v4 mapped addresses.

1) Configure the IPv6 VPN BGP family on R3 and R8.

```
[edit]
lab@R3# set protocols bgp group ibgp family inet6-vpn unicast

[edit]
lab@R3# commit
commit complete
```

On R8 the same configuration should be applied.

2) Configure the IPv6 BGP family on the route reflector.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family inet6-vpn unicast

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family inet6-vpn unicast

[edit]
lab@route-reflector# commit
commit complete
```

3) Configure customer C3 VPN.

The configuration of the routing instance is basically the same as with IPv4 L3VPN. The difference is the IPv6 enabled CE interfaces.

© 2020 Juniper Networks, Inc. All rights reserved.

a) Configure C3 VPN on R3.

```
[edit]
lab@R3# set interfaces lo0.2 family inet address 172.30.5.18/32

[edit]
lab@R3# set interfaces lo0.2 family inet6 address fd17:f0f4:f691:5::12/128

[edit]
lab@R3# set routing-instances C3 instance-type vrf

[edit]
lab@R3# set routing-instances C3 interface ge-0/0/5.317

[edit]
lab@R3# set routing-instances C3 interface lo0.2

[edit]
lab@R3# set routing-instances C3 vrf-target target:54591:300

[edit]
lab@R3# set routing-instances C3 protocols bgp group ce peer-as 64601

[edit]
lab@R3# set routing-instances C3 protocols bgp group ce neighbor fc09:c0:ffee::a

[edit]
lab@R3# set routing-instances C3 protocols bgp group ce as-override

[edit]
lab@R3# commit
commit complete
```

As with IPv4 , the `as-override` knob is needed to override the customer AS with the core AS number, allowing the CE routers to accept the routes from remote sites.

b) Configure C3 VPN on R8.

```
[edit]
lab@R8# set interfaces lo0.2 family inet address 172.30.5.38/32

[edit]
lab@R8# set interfaces lo0.2 family inet6 address fd17:f0f4:f691:5::26/128

[edit]
lab@R8# set routing-instances C3 instance-type vrf
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set routing-instances C3 interface lo0.2

[edit]
lab@R8# set routing-instances C3 interface ge-0/0/5.325

[edit]
lab@R8# set routing-instances C3 vrf-target target:54591:300

[edit]
lab@R8# set routing-instances C3 protocols bgp group ce peer-as 64601

[edit]
lab@R8# set routing-instances C3 protocols bgp group ce neighbor fc09:c0:ffee::e

[edit]
lab@R8# set routing-instances C3 protocols bgp group ce as-override

[edit]
lab@R8# commit
commit complete
```

Verify the operation of 6VPE sites between R3 and R8.

```
[edit]
lab@R3# run show bgp summary
Groups: 5 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     660        598          0          0          0          0
bgp.l3vpn.0
                      19         19          0          0          0          0
inet6.0
                      48         48          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591        440        220       0       0        4:46
Establ
  inet.0: 447/447/447/0
  bgp.l3vpn.0: 19/19/19/0
  inet6.0: 32/32/32/0
  bgp.l3vpn-inet6.0: 10/10/10/0
  C1.inet.0: 9/18/18/0
  bgp.rtarget.0: 13/15/15/0
  C3.inet6.0: 9/10/10/0
  C3.inet.0: 1/1/1/0
192.168.0.2      3521382357        952       9630       0       0     6:49:33
Establ
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  inet.0: 69/97/69/0
192.168.0.6       2831679853         962         9631         0         0      6:49:49
Establ
  inet.0: 82/116/82/0
fc09:c0:ffee::a         64601          46           49         0         0        19:44
Establ
  C3.inet6.0: 8/8/8/0
fe80::223:9c01:2d8b:6c81  3521382357         876         931         0         0
6:49:41 Establ
  inet6.0: 16/16/16/0
```

R3 receives 10 routes from the inet6-vpn family/ NLRI. These routes are placed in the `bgp.l3vpn-inet6.0` table. This is analogous to the `bgp.l3vpn.0` table for IPv4 L3VPN.

Based on the route target the routes are moved to the corresponding VRF `inet6.0` table.

```
[edit]
lab@R3# run show route table C3.inet6.0 | no-more

C3.inet6.0: 23 destinations, 24 routes (23 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fc09:c0:ffee::8/126*[Direct/0] 00:22:02
                    > via ge-0/0/5.317
fc09:c0:ffee::9/128*[Local/0] 00:22:02
                      Local via ge-0/0/5.317
fd17:f0f4:f691:5::12/128*[Direct/0] 00:22:02
                    > via lo0.2
fc09:c0:ffee::c/126*[BGP/170] 00:00:50, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
fd18:cccc:dddd:5::/80
                    *[BGP/170] 00:00:50, localpref 100, from 172.30.5.41
                      AS path: 64601 I
                     > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
fd18:cccc:dddd:5:1::/80
                    *[BGP/170] 00:00:50, localpref 100, from 172.30.5.41
                      AS path: 64601 I
                     > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
fd18:cccc:dddd:5:2::/80
                   *[BGP/170] 00:00:50, localpref 100, from 172.30.5.41
                      AS path: 64601 I
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
fd18:cccc:dddd:5:3::/80
                   *[BGP/170] 00:00:50, localpref 100, from 172.30.5.41
                      AS path: 64601 I
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
fd18:cccc:dddd:5:4::/80
                   *[BGP/170] 00:00:50, localpref 100, from 172.30.5.41
                      AS path: 64601 I
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
fd18:cccc:dddd:5:5::/80
                   *[BGP/170] 00:00:50, localpref 100, from 172.30.5.41
                      AS path: 64601 I
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
fd18:cccc:dddd:5:6::/80
                   *[BGP/170] 00:00:50, localpref 100, from 172.30.5.41
                      AS path: 64601 I
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
fd18:cccc:dddd:5:7::/80
                   *[BGP/170] 00:00:50, localpref 100, from 172.30.5.41
                      AS path: 64601 I
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
fd18:cccc:dddd:5:8::/80
                   *[BGP/170] 00:21:40, localpref 100
                      AS path: 64601 I
                    > to fc09:c0:ffee::a via ge-0/0/5.317
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

The routes are advertised to the CE routers via the native IPv6 EBGP session.

```
[edit]
lab@R3# run show route advertising-protocol bgp fc09:c0:ffee::a

C3.inet6.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* fc09:c0:ffee::c/126     Self                                    I
  fd17:f0f4:f691:5::26/128
*                         Self                                    I
  fd18:cccc:dddd:5::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:1::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:2::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:3::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:4::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:5::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:6::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:7::/80
*                         Self                                    54591 I
```

Verify if the routes are present in the `C3.inet6.0` table at R8.

```
[edit]
lab@R8# run show route table C3.inet6.0

C3.inet6.0: 23 destinations, 24 routes (23 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fc09:c0:ffee::8/126*[BGP/170] 00:00:13, localpref 100, from 172.30.5.41
                      AS path: I
                      to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                    > to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
                      to 172.30.0.45 via ge-0/0/4.178, label-switched-path
Bypass->172.30.0.9->172.30.0.2
fc09:c0:ffee::c/126*[Direct/0] 00:18:39
                    > via ge-0/0/5.325
fc09:c0:ffee::d/128*[Local/0] 00:18:39
                      Local via ge-0/0/5.325
fd17:f0f4:f691:5::12/128*[Direct/0] 00:18:39
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                     > via lo0.2
fd18:cccc:dddd:5::/80
                    *[BGP/170] 00:18:03, localpref 100
                       AS path: 64601 I
                     > to fc09:c0:ffee::e via ge-0/0/5.325
fd18:cccc:dddd:5:1::/80
                    *[BGP/170] 00:18:03, localpref 100
                       AS path: 64601 I
                     > to fc09:c0:ffee::e via ge-0/0/5.325
fd18:cccc:dddd:5:2::/80
                    *[BGP/170] 00:18:03, localpref 100
                       AS path: 64601 I
                     > to fc09:c0:ffee::e via ge-0/0/5.325
fd18:cccc:dddd:5:3::/80
                    *[BGP/170] 00:18:03, localpref 100
                       AS path: 64601 I
                     > to fc09:c0:ffee::e via ge-0/0/5.325
fd18:cccc:dddd:5:4::/80
                    *[BGP/170] 00:18:03, localpref 100
                       AS path: 64601 I
                     > to fc09:c0:ffee::e via ge-0/0/5.325
fd18:cccc:dddd:5:5::/80
                    *[BGP/170] 00:18:03, localpref 100
                       AS path: 64601 I
                     > to fc09:c0:ffee::e via ge-0/0/5.325
fd18:cccc:dddd:5:6::/80
                    *[BGP/170] 00:18:03, localpref 100
                       AS path: 64601 I
                     > to fc09:c0:ffee::e via ge-0/0/5.325
fd18:cccc:dddd:5:7::/80
                    *[BGP/170] 00:18:03, localpref 100
                       AS path: 64601 I
                     > to fc09:c0:ffee::e via ge-0/0/5.325
fd18:cccc:dddd:5:8::/80
                    *[BGP/170] 00:00:13, localpref 100, from 172.30.5.41
                       AS path: 64601 I
                       to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R3-1
                     > to 172.30.0.37 via ge-0/0/4.158, label-switched-path R8-to-
R3-2
                       to 172.30.0.45 via ge-0/0/4.178, label-switched-path
Bypass->172.30.0.9->172.30.0.2
...
```

Verify route advertisement to the CE router on R8.

```
[edit]
lab@R8# run show route advertising-protocol bgp fc09:c0:ffee::e
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
C3.inet6.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* fc09:c0:ffee::8/126     Self                                    I
  fd17:f0f4:f691:5::12/128
*                         Self                                    I
  fd18:cccc:dddd:5:8::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:9::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:a::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:b::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:c::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:d::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:e::/80
*                         Self                                    54591 I
  fd18:cccc:dddd:5:f::/80
*                         Self                                    54591 I
```

Customer C3 IPv6 routes are exchanged between the sites. The routes are advertised to the CE routers.

This concludes Chapter Five.

© 2020 Juniper Networks, Inc. All rights reserved.

Appendix - Chapter 6: Layer 2 VPN Services

The tasks included in this chapter introduce L2VPN, VPLS and EVPN services to the MPLS network. LDP and BGP signaled L2VPN, VPLS and EVPN will be configured for the customer networks as mentioned in the tasks.

Task 1: L2VPN Configuration

L2VPN Overview

BGP (Kompella) and LDP (Martini) signaled L2VPNs are supported in Junos. Both L2VPN technologies are used for connectivity of customer attachment circuits across the provider MPLS network. A L2VPN is a point-to-point connection between two PE routers. This type of VPN emulates transparent L2 services for the customers. Normally the L2 attachment circuits on each end are the same but may also be different. For example, it is possible to have a tagged interface on one end and an untagged interface on the other end of the L2VPN.

Kompella and Martini L2VPN's both use Martini data plane encapsulation.

BGP based L2VPNs use an additional BGP family/ NLRI (`family l2-vpn`) to signal VPN label information. This L2VPN type also supports auto-provisioning, which reduces the amount of work that needs to be performed when adding or moving (new) sites.

LDP based L2VPN use LDP for signaling extended/ targeted LDP sessions, which is the same as used for LDP tunneling over RSVP LSPs. A Virtual Circuit Label is distributed between PEs, which identifies the specific L2VPN. The Local and remote sites are configured manually, no auto-discovery is supported.

1) Configure additional interfaces on the routers with L2VPN sites.

```
[edit]
lab@R1# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R1# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services
```

Using `flexible-ethernet-services` allows you to configure a different encapsulation type per unit. If you use `encapsulation vlan-ccc` every unit must be configured with `encapsulation vlan-ccc`

```
[edit]
lab@R1# set interfaces ge-0/0/3.512 encapsulation vlan-ccc
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set interfaces ge-0/0/3.512 vlan-id 512

[edit]
lab@R1# set interfaces ge-0/0/3.513 encapsulation vlan-ccc

[edit]
lab@R1# set interfaces ge-0/0/3.513 vlan-id 513

[edit]
lab@R1# set interfaces ge-0/0/3.514 encapsulation vlan-ccc

[edit]
lab@R1# set interfaces ge-0/0/3.514 vlan-id 514

[edit]
lab@R1# commit
commit complete
```

The `vlan-ccc` encapsulation configures the port for VLAN circuit cross connect.

R3

```
[edit]
lab@R3# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R3# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R3# set interfaces ge-0/0/3.512 encapsulation vlan-ccc

[edit]
lab@R3# set interfaces ge-0/0/3.512 vlan-id 512

[edit]
lab@R3# set interfaces ge-0/0/3.513 encapsulation vlan-ccc

[edit]
lab@R3# set interfaces ge-0/0/3.513 vlan-id 513

[edit]
lab@R3# set interfaces ge-0/0/3.514 encapsulation vlan-ccc


[edit]
lab@R3# set interfaces ge-0/0/3.514 vlan-id 514
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# commit
commit complete
```

R5

```
[edit]
lab@R5# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R5# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R5# set interfaces ge-0/0/3.512 encapsulation vlan-ccc

[edit]
lab@R5# set interfaces ge-0/0/3.512 vlan-id 512

[edit]
lab@R5# set interfaces ge-0/0/3.513 encapsulation vlan-ccc

[edit]
lab@R5# set interfaces ge-0/0/3.513 vlan-id 513

[edit]
lab@R5# set interfaces ge-0/0/3.514 encapsulation vlan-ccc

[edit]
lab@R5# set interfaces ge-0/0/3.514 vlan-id 514

[edit]
lab@R5# commit
commit complete
```

R6

```
[edit]
lab@R6# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R6# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R6# set interfaces ge-0/0/3.512 encapsulation vlan-ccc

[edit]
lab@R6# set interfaces ge-0/0/3.512 vlan-id 512
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set interfaces ge-0/0/3.513 encapsulation vlan-ccc

[edit]
lab@R6# set interfaces ge-0/0/3.513 vlan-id 513

[edit]
lab@R6# set interfaces ge-0/0/3.514 encapsulation vlan-ccc

[edit]
lab@R6# set interfaces ge-0/0/3.514 vlan-id 514

[edit]
lab@R6# commit
commit complete
```

R7

```
[edit]
lab@R7# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R7# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R7# set interfaces ge-0/0/3.512 encapsulation vlan-ccc

[edit]
lab@R7# set interfaces ge-0/0/3.512 vlan-id 512

[edit]
lab@R7# set interfaces ge-0/0/3.513 encapsulation vlan-ccc

[edit]
lab@R7# set interfaces ge-0/0/3.513 vlan-id 513

[edit]
lab@R7# set interfaces ge-0/0/3.514 encapsulation vlan-ccc

[edit]
lab@R7# set interfaces ge-0/0/3.514 vlan-id 514

[edit]
lab@R7# commit
commit complete
```

R8

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R8# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R8# set interfaces ge-0/0/3.512 encapsulation vlan-ccc

[edit]
lab@R8# set interfaces ge-0/0/3.512 vlan-id 512

[edit]
lab@R8# set interfaces ge-0/0/3.513 encapsulation vlan-ccc

[edit]
lab@R8# set interfaces ge-0/0/3.513 vlan-id 513

[edit]
lab@R8# set interfaces ge-0/0/3.514 encapsulation vlan-ccc

[edit]
lab@R8# set interfaces ge-0/0/3.514 vlan-id 514

[edit]
lab@R8# commit
commit complete
```

Note: Interface encapsulation `vlan-ccc` is used for both BGP and LDP L2VPNs.

2) Configure BGP family L2VPN signaling on R3, R5, R7 and on the route reflector.

BGP L2VPN uses BGP for signaling the labels.

```
[edit]
lab@R3# set protocols bgp group ibgp family l2vpn signaling

[edit]
lab@R3# commit
commit complete
```

Configure both IBGP groups on the route reflector also.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family l2vpn signaling
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 family l2vpn signaling

[edit]
lab@route-reflector# commit
commit complete
```

3) Configure customer C5's L2VPN

For the L2VPN the routing-instance type of `l2vpn` must be chosen. Route distinguishers and route-targets are needed as with L3VPN configuration.

a) R7.

According to figure 12 Site 1 for customer C5 is connected to R7.

```
[edit]
lab@R7# set routing-instances C5-l2vpn instance-type l2vpn

[edit]
lab@R7# set routing-instances C5-l2vpn interface ge-0/0/3.512

[edit]
lab@R7# set routing-instances C5-l2vpn interface ge-0/0/3.513

[edit]
lab@R7# set routing-instances C5-l2vpn vrf-target target:54591:500
```

The route distinguisher is automatically generated because of the `route-distinguisher-id` option is configured under the `routing-options` hierarchy.

The actual L2VPN configuration is configured under the `protocols l2vpn` hierarchy.

```
[edit]
lab@R7# set routing-instances C5-l2vpn protocols l2vpn encapsulation-type ethernet-vlan

[edit]
lab@R7# set routing-instances C5-l2vpn protocols l2vpn site site-1 site-identifier 1

[edit]
lab@R7# set routing-instances C5-l2vpn protocols l2vpn site site-1 interface ge-
0/0/3.512


[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# set routing-instances C5-l2vpn protocols l2vpn site site-1 interface ge-
0/0/3.513

[edit]
lab@R7# commit
commit complete
```

The encapsulation of the L2VPN is set to `ethernet-vlan`. This site is identified with number 1. Next, both local interfaces are configured. The site-identifier and the order of the local interfaces are important for the correct remote site and interface associations. To determine the remote site ids: the interfaces of R7 start from 1, which in that case is the local site. The next value is 2, so interface `ge-0/0/3.512` is assigned to remote-site 2. The next value is 3, interface `ge-0/0/3.513` is assigned to remote-site 3.

Note: Remote-site id can be manually configured with remote-site-id knob under each interface. Site-identifiers should be unique for each PE in L2VPN.

b) R3.

```
[edit]
lab@R3# set routing-instances C5-l2vpn instance-type l2vpn

[edit]
lab@R3# set routing-instances C5-l2vpn interface ge-0/0/3.512

[edit]
lab@R3# set routing-instances C5-l2vpn interface ge-0/0/3.514

[edit]
lab@R3# set routing-instances C5-l2vpn vrf-target target:54591:500

[edit]
lab@R3# set routing-instances C5-l2vpn protocols l2vpn encapsulation-type ethernet-vlan

[edit]
lab@R3# set routing-instances C5-l2vpn protocols l2vpn site site-2 site-identifier 2

[edit]
lab@R3# set routing-instances C5-l2vpn protocols l2vpn site site-2 interface ge-
0/0/3.512

[edit]
lab@R3# set routing-instances C5-l2vpn protocols l2vpn site site-2 interface ge-
0/0/3.514


[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# commit
commit complete
```

R3s local site-identifier is 2. Starting from a value of 1 and assign it to interface ge-0/0/3.512 as remote-site id. This corresponds to R7. Value 2 is used to identify the local site. Next value 3 is assigned as the remote-site id for interface ge-0/0/3.514.

c) R5.

```
[edit]
lab@R5# set routing-instances C5-l2vpn instance-type l2vpn

[edit]
lab@R5# set routing-instances C5-l2vpn interface ge-0/0/3.513

[edit]
lab@R5# set routing-instances C5-l2vpn interface ge-0/0/3.514

[edit]
lab@R5# set routing-instances C5-l2vpn vrf-target target:54591:500

[edit]
lab@R5# set routing-instances C5-l2vpn protocols l2vpn encapsulation-type ethernet-vlan

[edit]
lab@R5# set routing-instances C5-l2vpn protocols l2vpn site site-3 site-identifier 3

[edit]
lab@R5# set routing-instances C5-l2vpn protocols l2vpn site site-3 interface ge-
0/0/3.513

[edit]
lab@R5# set routing-instances C5-l2vpn protocols l2vpn site site-3 interface ge-
0/0/3.514

[edit]
lab@R5# commit
commit complete
```

The remote-site id for interface ge-0/0/3.513 is 1 on R7. For interface ge-0/0/3.514 the remote-side id is 2 and the local site id is 3.

Verify the operation of the BGP L2VPN.

```
lab@R7# run show l2vpn connections extensive
Layer-2 VPN connections:
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
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           HS -- Hot-standby Connection

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

Instance: C5-l2vpn
Edge protection: Not-Primary
  Local site: site-1 (1)
    Number of local interfaces: 2
    Number of local interfaces up: 2
    ge-0/0/3.512        2
    ge-0/0/3.513        3
    Label-base        Offset     Size  Range     Preference
    800000            1          4      3         100
      status-vector:  0
    connection-site           Type  St     Time last up          # Up trans
    2                         rmt   Up     Apr  4 17:06:17 2019           1
      Remote PE: 172.30.5.3, Negotiated control-word: Yes (Null)
      Incoming label: 800001, Outgoing label: 800000
      Local interface: ge-0/0/3.512, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
  Connection History:
      Apr  4 17:06:17 2019  status update timer
      Apr  4 17:06:17 2019  PE route changed
      Apr  4 17:06:17 2019  Out lbl Update                    800000
      Apr  4 17:06:17 2019  In lbl Update                     800001
      Apr  4 17:06:17 2019  loc intf up                ge-0/0/3.512
  3                         rmt   Up     Apr  4 17:06:17 2019           1
      Remote PE: 172.30.5.5, Negotiated control-word: Yes (Null)
      Incoming label: 800002, Outgoing label: 800000
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Local interface: ge-0/0/3.513, Status: Up, Encapsulation: VLAN
    Flow Label Transmit: No, Flow Label Receive: No
  Connection History:
      Apr  4 17:06:17 2019  status update timer
      Apr  4 17:06:17 2019  PE route changed
      Apr  4 17:06:17 2019  Out lbl Update                  800000
      Apr  4 17:06:17 2019  In lbl Update                   800002
      Apr  4 17:06:17 2019  loc intf up                ge-0/0/3.513
```

Both connections to the remote sites are in the **UP** state. The remote-site id automatic assignment is also shown for each interface.

How are the Incoming and Outgoing label determined? The following formulas are used for VPN label calculation.

L2VPN label (Outgoing) = Label-Base (remote) + Site-Id(local) – Label-Offset (remote)

L2VPN label (Incoming) = Label-Base (local) + Site-Id(remote) – Label-Offset (local)

© 2020 Juniper Networks, Inc. All rights reserved.

Let's look at R3's L2VPN connection table for the calculated labels.

```
[edit]
lab@R3# run show l2vpn connections
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           HS -- Hot-standby Connection

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

Instance: C5-l2vpn
Edge protection: Not-Primary
  Local site: site-2 (2)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Apr  4 17:06:18 2019           1
      Remote PE: 172.30.5.7, Negotiated control-word: Yes (Null)
      Incoming label: 800000, Outgoing label: 800001
      Local interface: ge-0/0/3.512, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
    3                         rmt   Up     Apr  4 17:06:16 2019           1
      Remote PE: 172.30.5.5, Negotiated control-word: Yes (Null)
      Incoming label: 800002, Outgoing label: 800001
      Local interface: ge-0/0/3.514, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
```

To find the Label-Base and the Label-offset you have to look at the L2VPN prefixes in `C5-l2vpn.l2vpn.0` routing table.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show route table C5-l2vpn.l2vpn.0 extensive

C5-l2vpn.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
 172.30.5.3:65534:2:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 18d4b28
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [54591] I
    Communities: target:54591:500 Layer2-info: encaps:VLAN, control
flags:Control-Word, mtu: 0, site preference: 100
Path 172.30.5.3:65534:2:1 Vector len 4.  Val: 0
RTF { 54591:500/64 }
        *L2VPN  Preference: 170/-101
                Next hop type: Indirect
                Address: 0x18ec8f8
                Next-hop reference count: 2
                Protocol next hop: 172.30.5.3
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 20:23      Metric2: 1
                Task: C5-l2vpn-l2vpn
                Announcement bits (2): 1-BGP_RT_Background 2-BGP Route Target
                AS path: I
                Communities: Layer2-info: encaps:VLAN, control flags:Control-
Word, mtu: 0, site preference: 100
                Label-base: 800000, range: 4, status-vector: 0x0

 172.30.5.5:65534:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 172.30.5.5:65534
                Next hop type: Indirect
                Address: 0x18edb60
                Next-hop reference count: 5
                Source: 172.30.5.41
                Protocol next hop: 172.30.5.5
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 3:29       Metric2: 1
                Task: BGP_54591.172.30.5.41+1513
                Announcement bits (1): 0-C5-l2vpn-l2vpn
                AS path: I (Originator)
                Cluster list:  0.0.0.2
                Originator ID: 172.30.5.5
                Communities: target:54591:500 Layer2-info: encaps:VLAN, control
flags:Control-Word, mtu: 0, site preference: 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                Import Accepted
                Label-base: 800000, range: 2, status-vector: 0x0
                Localpref: 100
                Router ID: 172.30.5.41
                Primary Routing Table bgp.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 172.30.5.5 Metric: 15
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 172.30.0.22 via ge-0/0/4.134 weight 0x1
                                Next hop: 172.30.0.13 via ge-0/0/4.123 weight
0x8001
                        172.30.5.5/32 Originating RIB: inet.3
                          Metric: 15                        Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 172.30.0.22 via ge-0/0/4.134
                                Nexthop: 172.30.0.13 via ge-0/0/4.123

 172.30.5.7:65533:1:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 172.30.5.7:65533
                Next hop type: Indirect
                Address: 0x18ed738
                Next-hop reference count: 5
                Source: 172.30.5.41
                Protocol next hop: 172.30.5.7
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 3:22       Metric2: 1
                Task: BGP_54591.172.30.5.41+1513
                Announcement bits (1): 0-C5-l2vpn-l2vpn
                AS path: I (Originator)
                Cluster list:  0.0.0.1
                Originator ID: 172.30.5.7
                Communities: target:54591:500 Layer2-info: encaps:VLAN, control
flags:Control-Word, mtu: 0, site preference: 100
                Import Accepted
                Label-base: 800000, range: 4, status-vector: 0x0
                Localpref: 100
                Router ID: 172.30.5.41
                Primary Routing Table bgp.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 172.30.5.7 Metric: 20
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 172.30.0.22 via ge-0/0/4.134 weight 0x1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                                Next hop: 172.30.0.13 via ge-0/0/4.123 weight
0x8001
                            172.30.5.7/32 Originating RIB: inet.3
                              Metric: 20                        Node path count: 1
                              Forwarding nexthops: 2
                                    Nexthop: 172.30.0.22 via ge-0/0/4.134
                                    Nexthop: 172.30.0.13 via ge-0/0/4.123
```

L2VPN route is constructed from the RD:Remote-site-id:Remote-label-offset. The L2VPN NLRI also contains a remote label base, encapsulation, control word, MTU and site preference.

For example, the route from R7 `172.30.5.7:65533:1:1/96` can be interpreted as:

RD – 172.30.5.7:65533
Remote-site-Id – 1
Remote-label-offset – 1

Now you have all data to calculate the outgoing label.

800000 (Label-Base) + 2 (Local-site-Id) – 1 (Remote-label-offset) = 800001

4) Configure loopback in LDP on R6, R1 and R8.

LDP L2VPN is using LDP for signaling the tunnels. Extended LDP sessions are used for exchanging the control information. To enable extended LDP you have to configure interface `lo0.0` under the `protocols ldp` hierarchy.

Because LDP tunneling over RSVP is also using extended LDP sessions, loopback interface is configured as LDP interface in Chapter Four.

```
[edit]
lab@R1# show protocols ldp
track-igp-metric;
egress-policy ldp-routes;
deaggregate;
explicit-null;
interface ge-0/0/4.114;
interface ae0.0;
interface lo0.0;
session 172.30.5.2 {
    authentication-key "$9$X/f-Yo5T3/A0k.fz6A1IX7-w2a"; ## SECRET-DATA
}
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
session 172.30.5.4 {
    authentication-key "$9$AVqP0EyVb2oZjX7ds4ZkqAp0Ihr"; ## SECRET-DATA
```

5) Configure the customer C4 L2VPN.

Unlike BGP L2VPN, LDP L2VPNs are not configured under routing instances. For each L2 connection a L2circuit must be configured including all required parameters. The remote PE routes are configured manually, and auto-discovery is not supported with LDP L2VPNs.

a) R1

To accomplish the full mesh requirement, R1 should have a L2 connection to R6 and R8 using the VLANs.

```
[edit]
lab@R1# set protocols l2circuit neighbor 172.30.5.8 interface ge-0/0/3.512 virtual-
circuit-id 512

[edit]
lab@R1# set protocols l2circuit neighbor 172.30.5.6 interface ge-0/0/3.513 virtual-
circuit-id 513

[edit]
lab@R1# commit
commit complete
```

The neighbor is the remote PE loopback IP address enabled for LDP. An extended LDP session is establish to the neighbors configured for l2circuits. The `virtual-circuit-id` number configured under the local interface should match on both ends of the L2 connection.

b) R6

R6 should have a l2circuit to R1 and R8.

```
[edit]
lab@R6# set protocols l2circuit neighbor 172.30.5.1 interface ge-0/0/3.513 virtual-
circuit-id 513

[edit]
lab@R6# set protocols l2circuit neighbor 172.30.5.8 interface ge-0/0/3.514 virtual-
circuit-id 514
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# commit
commit complete
```

c) R8

R8 should have a l2circuit to R1 and R6.

```
[edit]
lab@R8# set protocols l2circuit neighbor 172.30.5.1 interface ge-0/0/3.512 virtual-
circuit-id 512

[edit]
lab@R8# set protocols l2circuit neighbor 172.30.5.6 interface ge-0/0/3.514 virtual-
circuit-id 514

[edit]
lab@R8# commit
commit complete
```

Verify the operation of the LDP L2VPN on R1.

```
[edit]
lab@R1# run 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  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- Hot-standby Connection
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 172.30.5.6
    Interface                 Type  St     Time last up          # Up trans
    ge-0/0/3.513(vc 513)      rmt   Up     Apr  4 17:22:23 2019           1
      Remote PE: 172.30.5.6, Negotiated control-word: Yes (Null)
      Incoming label: 378880, Outgoing label: 363872
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
      Negotiated PW status TLV: No
      Local interface: ge-0/0/3.513, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
Neighbor: 172.30.5.8
    Interface                 Type  St     Time last up          # Up trans
    ge-0/0/3.512(vc 512)      rmt   Up     Apr  4 17:22:53 2019           1
      Remote PE: 172.30.5.8, Negotiated control-word: Yes (Null)
      Incoming label: 378864, Outgoing label: 371904
      Negotiated PW status TLV: No
      Local interface: ge-0/0/3.512, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
```

Both l2circuit connections are UP. Also, the Incoming and Outgoing labels are shown.

```
[edit]
lab@R1# run show ldp database session 172.30.5.6 l2circuit extensive
Input label database, 172.30.5.1:0--172.30.5.6:0
Labels received: 10
  Label     Prefix
 363872      L2CKT CtrlWord VLAN VC 513
            MTU: 1500 Requested VLAN ID: 513 Flow Label T Bit: 0 Flow Label R
Bit: 0
            State: Active
            Age: 2:20
            VCCV Control Channel types:
                PWE3 control word
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
                BFD with PW-ACH-encapsulation for Fault Detection
                BFD with IP/UDP-encapsulation for Fault Detection

Output label database, 172.30.5.1:0--172.30.5.6:0
Labels advertised: 9
  Label     Prefix
 378880      L2CKT CtrlWord VLAN VC 513
            MTU: 1500 Requested VLAN ID: 513 Flow Label T Bit: 0 Flow Label R
Bit: 0
            State: Active
            Age: 2:50
            VCCV Control Channel types:
                PWE3 control word
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
                BFD with PW-ACH-encapsulation for Fault Detection
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                BFD with IP/UDP-encapsulation for Fault Detection
```

```
[edit]
lab@R1# run show ldp database session 172.30.5.8 l2circuit extensive
Input label database, 172.30.5.1:0--172.30.5.8:0
Labels received: 10
  Label     Prefix
 371904      L2CKT CtrlWord VLAN VC 512
            MTU: 1500 Requested VLAN ID: 512 Flow Label T Bit: 0 Flow Label R
Bit: 0
            State: Active
            Age: 4:01
            VCCV Control Channel types:
                PWE3 control word
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
                BFD with PW-ACH-encapsulation for Fault Detection
                BFD with IP/UDP-encapsulation for Fault Detection

Output label database, 172.30.5.1:0--172.30.5.8:0
Labels advertised: 9
  Label     Prefix
 378864      L2CKT CtrlWord VLAN VC 512
            MTU: 1500 Requested VLAN ID: 512 Flow Label T Bit: 0 Flow Label R
Bit: 0
            State: Active
            Age: 5:01
            VCCV Control Channel types:
                PWE3 control word
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
                BFD with PW-ACH-encapsulation for Fault Detection
                BFD with IP/UDP-encapsulation for Fault Detection
```

The extensive output of the LDP database displays the Incoming and Outgoing labels signaled with R6 and R8. The MTU of the circuit and the VLANs are exchanged and should match.

Verify the operation of the LDP L2VPN on R6.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# run 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  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- Hot-standby Connection
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 172.30.5.1
    Interface                 Type  St     Time last up          # Up trans
    ge-0/0/3.513(vc 513)      rmt   Up     Apr  4 17:22:18 2019           1
      Remote PE: 172.30.5.1, Negotiated control-word: Yes (Null)
      Incoming label: 363872, Outgoing label: 378880
      Negotiated PW status TLV: No
      Local interface: ge-0/0/3.513, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
Neighbor: 172.30.5.8
    Interface                 Type  St     Time last up          # Up trans
    ge-0/0/3.514(vc 514)      rmt   Up     Apr  4 17:22:56 2019           1
      Remote PE: 172.30.5.8, Negotiated control-word: Yes (Null)
      Incoming label: 363888, Outgoing label: 371920
      Negotiated PW status TLV: No
      Local interface: ge-0/0/3.514, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
```

```
[edit]
lab@R6# run show ldp database session 172.30.5.1 l2circuit extensive | no-more
Input label database, 172.30.5.6:0--172.30.5.1:0
Labels received: 9
  Label     Prefix
 378880      L2CKT CtrlWord VLAN VC 513
            MTU: 1500 Requested VLAN ID: 513 Flow Label T Bit: 0 Flow Label R
Bit: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
            State: Active
            Age: 16:07
            VCCV Control Channel types:
                PWE3 control word
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
                BFD with PW-ACH-encapsulation for Fault Detection
                BFD with IP/UDP-encapsulation for Fault Detection

Output label database, 172.30.5.6:0--172.30.5.1:0
Labels advertised: 10
  Label      Prefix
 363872       L2CKT CtrlWord VLAN VC 513
            MTU: 1500 Requested VLAN ID: 513 Flow Label T Bit: 0 Flow Label R
Bit: 0
            State: Active
            Age: 15:37
            VCCV Control Channel types:
                PWE3 control word
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
                BFD with PW-ACH-encapsulation for Fault Detection
                BFD with IP/UDP-encapsulation for Fault Detection
```

Verify the operation of the LDP L2VPN on R8.

```
[edit]
lab@R6# run show ldp database session 172.30.5.8 l2circuit extensive | no-more
Input label database, 172.30.5.6:0--172.30.5.8:0
Labels received: 10
  Label      Prefix
 371920       L2CKT CtrlWord VLAN VC 514
            MTU: 1500 Requested VLAN ID: 514 Flow Label T Bit: 0 Flow Label R
Bit: 0
            State: Active
            Age: 16:01
            VCCV Control Channel types:
                PWE3 control word
                MPLS router alert label
                MPLS PW label with TTL=1
            VCCV Control Verification types:
                LSP ping
                BFD with PW-ACH-encapsulation for Fault Detection
                BFD with IP/UDP-encapsulation for Fault Detection
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Output label database, 172.30.5.6:0--172.30.5.8:0
Labels advertised: 10
   Label      Prefix
  363888       L2CKT CtrlWord VLAN VC 514
               MTU: 1500 Requested VLAN ID: 514 Flow Label T Bit: 0 Flow Label R
Bit: 0
               State: Active
               Age: 16:35
               VCCV Control Channel types:
                   PWE3 control word
                   MPLS router alert label
                   MPLS PW label with TTL=1
              VCCV Control Verification types:
                   LSP ping
                   BFD with PW-ACH-encapsulation for Fault Detection
                   BFD with IP/UDP-encapsulation for Fault Detection
```

Verify the operation of the LDP L2VPN on R8.

```
[edit]
lab@R8# run 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  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- Hot-standby Connection
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 172.30.5.1
    Interface                 Type  St     Time last up          # Up trans
    ge-0/0/3.512(vc 512)      rmt   Up     Apr  4 17:22:48 2019           1
      Remote PE: 172.30.5.1, Negotiated control-word: Yes (Null)
      Incoming label: 371904, Outgoing label: 378864
      Negotiated PW status TLV: No
      Local interface: ge-0/0/3.512, Status: Up, Encapsulation: VLAN
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
      Flow Label Transmit: No, Flow Label Receive: No
Neighbor: 172.30.5.6
    Interface                 Type  St     Time last up          # Up trans
    ge-0/0/3.514(vc 514)      rmt   Up     Apr  4 17:22:48 2019           1
      Remote PE: 172.30.5.6, Negotiated control-word: Yes (Null)
      Incoming label: 371920, Outgoing label: 363888
      Negotiated PW status TLV: No
      Local interface: ge-0/0/3.514, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
```

```
[edit]
lab@R8# run show ldp database session 172.30.5.1 l2circuit
Input label database, 172.30.5.8:0--172.30.5.1:0
Labels received: 9
  Label     Prefix
 378864      L2CKT CtrlWord VLAN VC 512

Output label database, 172.30.5.8:0--172.30.5.1:0
Labels advertised: 10
  Label     Prefix
 371904      L2CKT CtrlWord VLAN VC 512
```

```
[edit]
lab@R8# run show ldp database session 172.30.5.6 l2circuit
Input label database, 172.30.5.8:0--172.30.5.6:0
Labels received: 10
  Label     Prefix
 363888      L2CKT CtrlWord VLAN VC 514

Output label database, 172.30.5.8:0--172.30.5.6:0
Labels advertised: 10
  Label     Prefix
 371920      L2CKT CtrlWord VLAN VC 514
```

Task 2: VPLS Configuration

VPLS Overview

Virtual Private LAN Service (VPLS) emulates an Ethernet LAN segment over an MPLS network. VPLS acts as a bridge and MAC learning, similar to L2 switching, is enabled on the PEs. PE devices learn the MAC addresses of the frames that they receive from the CE interfaces and build a forwarding table. Based on this table frames are forwarded to locally attached PE-CE interfaces or over MPLS LSPs to other sites. A full mesh of LSPs is required between the PEs that participates in the same VPLS.

© 2020 Juniper Networks, Inc. All rights reserved.

As with L2VPN two standards for VPLS exist. The primary difference is that one uses BGP for signaling, while the other uses LDP. Junos supports both standards. BGP based VPLS service has the following benefits over LDP VPLS: auto-discovery, scalability using route reflectors and inter-AS operation.

BGP VPLS uses the same L2VPN NLRI as used with L2VPNs. So BGP `family l2vpn signaling` must be enabled on the PE and RR. BGP VPLS uses two tables, VRF routing and MAC table. Remote sites are learned dynamically. The Incoming and Outgoing labels are also calculated as they are for BGP L2VPNs.

LDP VPLS requires extended LDP sessions to be established between the PE loopback IP addresses. For every VPLS instance all neighbors should be configured manually. Junos also supporting BGP for auto-discovery of LDP VPLS neighbors.

1) Configure additional interfaces on **R2, R3, R4** and **R5**.

```
[edit]
lab@R2# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R2# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R2# set interface ge-0/0/3 unit 600 vlan-id 600

[edit]
lab@R2# set interface ge-0/0/3 unit 600 encapsulation vlan-vpls

[edit]
lab@R2# set interface ge-0/0/3 unit 601 vlan-id 601

[edit]
lab@R2# set interface ge-0/0/3 unit 601 encapsulation vlan-vpls

[edit]
lab@R2# commit
commit complete
```

2) Configure BGP family L2VPN signaling on **R2** and **R4**. You would also have to include R3, R5, and Router Reflector normally, but this was done previously for the L2VPNs.

R2

```
[edit]
lab@R2# set protocols bgp group ibgp family l2vpn signaling
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# commit
commit complete
```

R4

```
[edit]
lab@R4# set protocols bgp group ibgp family l2vpn signaling

[edit]
lab@R4# commit
commit complete
```

3) Configure customer C5 VPLS.

a) Configure VPLS routing instance on R2.

```
[edit]
lab@R2# set routing-instances C5-vpls instance-type vpls

[edit]
lab@R2# set routing-instances C5-vpls vlan-id all
```

The `vlan-id all` configuration creates a separate bridge domain for all VLANs. This will automatically enable L2 learning for VLAN600 and VLAN601.

```
[edit]
lab@R2# set routing-instances C5-vpls interface ge-0/0/3.600

[edit]
lab@R2# set routing-instances C5-vpls interface ge-0/0/3.601

[edit]
lab@R2# set routing-instances C5-vpls vrf-target target:54591:501
```

Configure `protocol vpls` under the routing instance.

```
[edit]
lab@R2# set routing-instances C5-vpls protocols vpls site-range 8
```

The `site-range` specifies the upper limit of the maximum site identifiers that can be accepted for VPLS connections.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set routing-instances C5-vpls protocols vpls no-tunnel-services
```

The no-tunnel-services option plays the same role as vrf-table-label with L3VPNs. It enables a virtual LSI interface for additional look-up. In the case with VPLS this look-up is in the MAC table. A Unique LSI interface is created for each remote site. On MX devices this command must be configured for VPLS operation.

```
[edit]
lab@R2# set routing-instances C5-vpls protocols vpls site site-4 site-identifier 4

[edit]
lab@R2# commit
commit complete
```

b) Configure a VPLS routing instance on R3.

The customer site connected to R3 is multi-homed with R4. To prevent L2 loops, both routing instances are configured with different preference values.

```
[edit]
lab@R3# set routing-instances C5-vpls instance-type vpls

[edit]
lab@R3# set routing-instances C5-vpls vlan-id all

[edit]
lab@R3# set routing-instances C5-vpls interface ge-0/0/3.600

[edit]
lab@R3# set routing-instances C5-vpls interface ge-0/0/3.601

[edit]
lab@R3# set routing-instances C5-vpls vrf-target target:54591:501

[edit]
lab@R3# set routing-instances C5-vpls protocols vpls site-range 8

[edit]
lab@R3# set routing-instances C5-vpls protocols vpls no-tunnel-services

[edit]
lab@R3# set routing-instances C5-vpls protocols vpls site site-5 site-identifier 5

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set routing-instances C5-vpls protocols vpls site site-5 multi-homing
```

The multi-homing configuration is enabling to protect against a corner case when all the BGP peers in the core are no available. In such a case this knob disables all VPLS interfaces to prevent the local site to become primary for the CE devices, which results in blackholing.

```
[edit]
lab@R3# set routing-instances C5-vpls protocols vpls site site-5 site-preference
primary
```

R3 is explicitly configured as primary. This automatically increases the preference value of the site to the maximum of 65535. This is changing the local-preference of the Multiprotocol BGP announcements making the primary site routes more preferred by other PEs. The site-identifier should be the same on both multi-homed PEs.

```
[edit]
lab@R3# commit
commit complete
```

c) Configure VPLS routing instance on R4.

R4 is the backup for the multi-homed site.

```
[edit]
lab@R4# set routing-instances C5-vpls instance-type vpls

[edit]
lab@R4# set routing-instances C5-vpls vlan-id all

[edit]
lab@R4# set routing-instances C5-vpls interface ge-0/0/3.601

[edit]
lab@R4# set routing-instances C5-vpls interface ge-0/0/3.600

[edit]
lab@R4# set routing-instances C5-vpls vrf-target target:54591:501

[edit]
lab@R4# set routing-instances C5-vpls protocols vpls site-range 8

[edit]
lab@R4# set routing-instances C5-vpls protocols vpls no-tunnel-services

[edit]
lab@R4# set routing-instances C5-vpls protocols vpls site site-5 site-identifier
5
```


© 2020 Juniper Networks, Inc. All rights reserved.

The same site-identifier as on R3 is configured on R4.

```
[edit]
lab@R4# set routing-instances C5-vpls protocols vpls site site-5 multi-homing

[edit]
lab@R4# set routing-instances C5-vpls protocols vpls site site-5 site-preference backup
```

The backup parameter sets the preference and respectively the local preference of the announcements from R4 to 1, which is lower than the value used by R3.

```
[edit]
lab@R4# commit
commit complete
```

d) Configure VPLS routing instance on R5.

```
[edit]
lab@R5# set routing-instances C5-vpls instance-type vpls

[edit]
lab@R5# set routing-instances C5-vpls vlan-id all

[edit]
lab@R5# set routing-instances C5-vpls interface ge-0/0/3.600

[edit]
lab@R5# set routing-instances C5-vpls interface ge-0/0/3.601

[edit]
lab@R5# set routing-instances C5-vpls vrf-target target:54591:501

[edit]
lab@R5# set routing-instances C5-vpls protocols vpls site-range 8

[edit]
lab@R5# set routing-instances C5-vpls protocols vpls no-tunnel-services

[edit]
lab@R5# set routing-instances C5-vpls protocols vpls site site-6 site-identifier 6

[edit]
lab@R5# commit
commit complete
```

Verify the operation of customer C5 BGP VPLS on R2.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show vpls connections
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           HS -- Hot-standby Connection

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

Instance: C5-vpls
Edge protection: Not-Primary
  Local site: site-4 (4)
    connection-site           Type  St     Time last up          # Up trans
    5                         rmt   Up     Apr  4 17:57:31 2019           1
      Remote PE: 172.30.5.3, Negotiated control-word: No
      Incoming label: 262149, Outgoing label: 262148
      Local interface: lsi.1048577, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 4 remote site 5
      Flow Label Transmit: No, Flow Label Receive: No
    6                         rmt   Up     Apr  4 17:57:31 2019           1
      Remote PE: 172.30.5.5, Negotiated control-word: No
      Incoming label: 262150, Outgoing label: 262148
      Local interface: lsi.1048576, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 4 remote site 6
      Flow Label Transmit: No, Flow Label Receive: No
```


© 2020 Juniper Networks, Inc. All rights reserved.

The output looks more or less the same as with BGP L2VPNs. The difference is that the `connection-site` combines all interfaces on the remote site. A separate LSI interface is created for each remote site and a different Incoming label is calculated.

The connection to the R4 is not UP, because its announcements has a lower local preference value of 1.

R4

```
[edit]
lab@R4# run show route advertising-protocol bgp 172.30.5.41 table C5-vpls.l2vpn.0
extensive

C5-vpls.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
*  172.30.5.4:10:5:1/96 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.30.5.4:10
     Label-base: 262145, range: 8, offset: 1
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 1
     AS path: [54591] I
     Communities: target:54591:501 Layer2-info: encaps: VPLS, control flags:[0x0]
, mtu: 0, site preference: 1
```

R3

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.41 table C5-vpls.l2vpn.0
extensive


C5-vpls.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
*  172.30.5.3:12:5:1/96 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.30.5.3:12
     Label-base: 262145, range: 8, offset: 1
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 65535
     AS path: [54591] I
     Communities: target:54591:501 Layer2-info: encaps: VPLS, control flags:[0x0]
, mtu: 0, site preference: 65535
```

The announcement from R3 carries the maximum local preference value of 65535. The same value is set for the site preference.

© 2020 Juniper Networks, Inc. All rights reserved.

Verify the operation of C5 VPLS on R3.

```
[edit]
lab@R3# run 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           HS -- Hot-standby Connection

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

Instance: C5-vpls
Edge protection: Not-Primary
  Local site: site-5 (5)
    Number of local interfaces: 2
    Number of local interfaces up: 2
    IRB interface present: no
    ge-0/0/3.600
    ge-0/0/3.601
    lsi.1048577         4         Intf - vpls C5-vpls local site 5 remote site 4
    lsi.1048576         6         Intf - vpls C5-vpls local site 5 remote site 6
    Label-base        Offset     Size  Range     Preference
    262145            1          8      8         65535
    connection-site           Type  St     Time last up          # Up trans
    4                         rmt   Up     Apr  4 17:57:31 2019           1
      Remote PE: 172.30.5.2, Negotiated control-word: No
      Incoming label: 262148, Outgoing label: 262149
      Local interface: lsi.1048577, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 5 remote site 4
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Flow Label Transmit: No, Flow Label Receive: No
Connection History:
    Apr  4 17:57:31 2019  status update timer
    Apr  4 17:57:31 2019  loc intf up                    lsi.1048577
    Apr  4 17:57:31 2019  PE route changed
    Apr  4 17:57:31 2019  Out lbl Update                      262149
    Apr  4 17:57:31 2019  In lbl Update                       262148
    Apr  4 17:57:31 2019  loc intf down
5                         rmt   RN
6                         rmt   Up     Apr  4 17:55:28 2019           1
  Remote PE: 172.30.5.5, Negotiated control-word: No
  Incoming label: 262150, Outgoing label: 262149
  Local interface: lsi.1048576, Status: Up, Encapsulation: VPLS
    Description: Intf - vpls C5-vpls local site 5 remote site 6
  Flow Label Transmit: No, Flow Label Receive: No
Connection History:
    Apr  4 17:55:28 2019  status update timer
    Apr  4 17:55:28 2019  loc intf up                    lsi.1048576
    Apr  4 17:55:28 2019  PE route changed
    Apr  4 17:55:28 2019  Out lbl Update                      262149
    Apr  4 17:55:28 2019  In lbl Update                       262150
    Apr  4 17:55:28 2019  loc intf down
```

The connections to the remote sites R2 and R5 are UP. The connection to site 5, with R4 is shown as `RN -- remote site not designated.`

```
[edit]
lab@R3# run show route forwarding-table family vpls
Routing table: C5-vpls.vpls
VPLS:
Enabled protocols: ACKed by all peers,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd     673     1
lsi.1048576        intf     0                    indr  1048640     6
                                                 ulst  1048667     2
                             172.30.0.13         Push 262149, Push 354848, Push
340336(top)      879     2 ge-0/0/4.123
                             172.30.0.22         Push 262149, Push 354848, Push
355184(top)      849     2 ge-0/0/4.134
lsi.1048577          intf     0                    indr  1048582     6
                             172.30.0.13         Push 262149      637     2 ge-
0/0/4.123
0x30003/51           user     0                    comp      860     2
ge-0/0/3.600         intf     0                    ucst      706     4 ge-0/0/3.600
ge-0/0/3.601         intf     0                    ucst      708     4 ge-0/0/3.601
0x30002/51           user     0                    comp      858     2
00:50:56:a9:13:d8/48 user       0                 indr  1048640     6
                                                 ulst  1048667     2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                                  172.30.0.13          Push 262149, Push 354848, Push
340336(top)      879     2 ge-0/0/4.123
                                  172.30.0.22          Push 262149, Push 354848, Push
355184(top)      849     2 ge-0/0/4.134
00:50:56:a9:17:ac/48 user     0                        indr  1048582     6
                                  172.30.0.13          Push 262149      637     2 ge-
0/0/4.123
00:50:56:a9:54:cf/48 user     0                        ucst      706     4 ge-0/0/3.600
00:50:56:a9:13:d8/48 user     0                        indr  1048640     6
                                                       ulst  1048667     2
                                  172.30.0.13          Push 262149, Push 354848, Push
340336(top)      879     2 ge-0/0/4.123
                                  172.30.0.22          Push 262149, Push 354848, Push
355184(top)      849     2 ge-0/0/4.134
00:50:56:a9:17:ac/48 user     0                        indr  1048582     6
                                  172.30.0.13          Push 262149      637     2 ge-
0/0/4.123
00:50:56:a9:54:cf/48 user     0                        ucst      708     4 ge-0/0/3.601
0x30001/51           user     0                        comp      857     2
```

The VPLS forwarding table contains the local interfaces and the dynamically learned MAC addresses. The `default` destination is used for packets which are classified for broadcast, unknown unicast, and multicast (BUM) transmission. These packets are flooded to all interfaces.

Verify the operation of C5 VPLS on R4.

R4 is the backup site of R3 for the multi-homed CE attached to both PEs.

```
[edit]
lab@R4# run 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
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
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                  HS -- Hot-standby Connection

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

Instance: C5-vpls
Edge protection: Not-Primary
  Local site: site-5 (5)
    Number of local interfaces: 2
    Number of local interfaces up: 2
    IRB interface present: no
    ge-0/0/3.600
        Interface flags: VC-Down
    ge-0/0/3.601
        Interface flags: VC-Down
    Label-base        Offset     Size  Range     Preference
    262145            1          8     8         1
    connection-site           Type  St     Time last up          # Up trans
    4                         rmt   LN
    5                         rmt   LN
    6                         rmt   LN
```

All sites are in the state `LN -- local site not designated.`

If you deactivate the VPLS routing instance on R3 (to simulate the failure of R3) the sites should come UP.

R3

```
[edit]
lab@R3# deactivate routing-instances C5-vpls

[edit]
lab@R3# commit
commit complete
```

Verify the table with VPLS connections on R4 again.

```
 [edit]
lab@R4# run show vpls connections extensive
Layer-2 VPN connections:
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
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           HS -- Hot-standby Connection

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

Instance: C5-vpls
Edge protection: Not-Primary
  Local site: site-5 (5)
    Number of local interfaces: 2
    Number of local interfaces up: 2
    IRB interface present: no
    ge-0/0/3.600
    ge-0/0/3.601
    lsi.1048576        4         Intf - vpls C5-vpls local site 5 remote site 4
    lsi.1048577        6         Intf - vpls C5-vpls local site 5 remote site 6
    Label-base       Offset     Size  Range     Preference
    262145           1          8      8         1
    connection-site           Type  St     Time last up          # Up trans
    4                         rmt   Up     Apr  4 18:19:24 2019           1
      Remote PE: 172.30.5.2, Negotiated control-word: No
      Incoming label: 262148, Outgoing label: 262149
      Local interface: lsi.1048576, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 5 remote site 4
      Flow Label Transmit: No, Flow Label Receive: No
    Connection History:
      Apr  4 18:19:24 2019  status update timer
      Apr  4 18:19:24 2019  loc intf up                  lsi.1048576
      Apr  4 18:19:23 2019  PE route changed
      Apr  4 18:19:23 2019  Out lbl Update                   262149
      Apr  4 18:19:23 2019  In lbl Update                    262148
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Apr  4 18:19:23 2019  loc intf down
6                         rmt   Up     Apr  4 18:19:24 2019           1
  Remote PE: 172.30.5.5, Negotiated control-word: No
  Incoming label: 262150, Outgoing label: 262149
  Local interface: lsi.1048577, Status: Up, Encapsulation: VPLS
    Description: Intf - vpls C5-vpls local site 5 remote site 6
  Flow Label Transmit: No, Flow Label Receive: No
Connection History:
    Apr  4 18:19:24 2019  status update timer
    Apr  4 18:19:24 2019  loc intf up                  lsi.1048577
    Apr  4 18:19:23 2019  PE route changed
    Apr  4 18:19:23 2019  Out lbl Update                    262149
    Apr  4 18:19:23 2019  In lbl Update                     262150
    Apr  4 18:19:23 2019  loc intf down
```

Both connections to the remote sites R2 and R5 are now UP.

Verify the operation of C5 VPLS on R5.

```
[edit]
lab@R5# run show vpls connections
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           HS -- Hot-standby Connection

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

Instance: C5-vpls
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Edge protection: Not-Primary
  Local site: site-6 (6)
    connection-site           Type  St     Time last up          # Up trans
    4                         rmt   Up     Apr  4 17:57:31 2019           1
      Remote PE: 172.30.5.2, Negotiated control-word: No
      Incoming label: 262148, Outgoing label: 262150
      Local interface: lsi.1048577, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 6 remote site 4
      Flow Label Transmit: No, Flow Label Receive: No
    5                         rmt   Up     Apr  4 17:55:28 2019           1
      Remote PE: 172.30.5.4, Negotiated control-word: No
      Incoming label: 262149, Outgoing label: 262150
      Local interface: lsi.1048576, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 6 remote site 5
```

Reenable the configuration on R3.

```
[edit]
lab@R3# rollback 1
load complete

[edit]
lab@R3# commit
commit complete
```

4) Now configure the additional PE-CE interfaces on **R1, R6, R7** and **R8** for LDP VPLS.

```
[edit]
lab@R1# set interfaces ge-0/0/3.700 vlan-id 700

[edit]
lab@R1# set interfaces ge-0/0/3.700 encapsulation vlan-vpls

[edit]
lab@R1# set interfaces ge-0/0/3.701 vlan-id 701

[edit]
lab@R1# set interfaces ge-0/0/3.701 encapsulation vlan-vpls

[edit]
lab@R1# commit
commit complete
```

5) Configure the customer C6 LDP VPLS instances.

Unlike the l2circuit configuration, LDP VPLS is configured as a routing-instance type `vpls`.

© 2020 Juniper Networks, Inc. All rights reserved.

a) Configure R1.

```
[edit]
lab@R1# set routing-instances C6-vpls instance-type vpls

[edit]
lab@R1# set routing-instances C6-vpls vlan-id all

[edit]
lab@R1# set routing-instances C6-vpls interface ge-0/0/3.700

[edit]
lab@R1# set routing-instances C6-vpls interface ge-0/0/3.701

[edit]
lab@R1# set routing-instances C6-vpls protocols vpls no-tunnel-services
```

Note: The `no-tunnel-services` parameter is needed also.

```
[edit]
lab@R1# set routing-instances C6-vpls protocols vpls encapsulation-type ethernet-vlan

[edit]
lab@R1# set routing-instances C6-vpls protocols vpls vpls-id 600
```

The `vpls-id` parameter is needed to identify the virtual circuit for the VPLS site. It should match on the remote sites.

```
[edit]
lab@R1# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.6
```

The neighbors should be configured manually.

```
[edit]
lab@R1# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.7 backup-
neighbor 172.30.5.8
```

The site attached to R7 is multihomed with R8. With LDP VPLS the backup is configured on the remote sites. With LDP VPLS, multihomed sites should not be configured with a neighbor relationship between each other.

The remote sites R1 and R8 are configured as backup site of R7.

```
[edit]
lab@R1# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.7 revert-time 60
```


© 2020 Juniper Networks, Inc. All rights reserved.

The `revert-time` parameter will delay the switching from the backup to the primary site, once the primary becomes available again (preempt delay).

```
[edit]
lab@R1# commit
commit complete
```

b) Configure R6.

```
[edit]
lab@R6# set routing-instances C6-vpls instance-type vpls

[edit]
lab@R6# set routing-instances C6-vpls vlan-id all

[edit]
lab@R6# set routing-instances C6-vpls interface ge-0/0/3.700

[edit]
lab@R6# set routing-instances C6-vpls interface ge-0/0/3.701

[edit]
lab@R6# set routing-instances C6-vpls protocols vpls no-tunnel-services

[edit]
lab@R6# set routing-instances C6-vpls protocols vpls vpls-id 600

[edit]
lab@R6# set routing-instances C6-vpls protocols vpls encapsulation-type ethernet-vlan
```

The `encapsulation-type` is set to `vlan`. By default, it is set to `ethernet`. Remember that this encapsulation is signaled by LDP (vc-type 4 or vc-type 5)

```
[edit]
lab@R6# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.1

[edit]
lab@R6# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.7 backup-
neighbor 172.30.5.8

[edit]
lab@R6# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.7 revert-time 60

[edit]
lab@R6# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

c) Configure R7.

CE attached to R7 is multi-homed on R8 for redundancy.

```
[edit]
lab@R7# set routing-instances C6-vpls instance-type vpls

[edit]
lab@R7# set routing-instances C6-vpls vlan-id all

[edit]
lab@R7# set routing-instances C6-vpls interface ge-0/0/3.700

[edit]
lab@R7# set routing-instances C6-vpls interface ge-0/0/3.701

[edit]
lab@R7# set routing-instances C6-vpls protocols vpls no-tunnel-services

[edit]
lab@R7# set routing-instances C6-vpls protocols vpls encapsulation-type ethernet-vlan

[edit]
lab@R7# set routing-instances C6-vpls protocols vpls vpls-id 600

[edit]
lab@R7# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.1

[edit]
lab@R7# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.6

[edit]
lab@R7# commit
commit complete
```

Unlike BGP VPLS with LDP VPLS no multihomed configuration is needed on multihomed sites.

d) Configure R8.

R8 is the other multi-homed site PE.

```
[edit]
lab@R8# set routing-instances C6-vpls instance-type vpls

[edit]
lab@R8# set routing-instances C6-vpls vlan-id all
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set routing-instances C6-vpls interface ge-0/0/3.700

[edit]
lab@R8# set routing-instances C6-vpls interface ge-0/0/3.701

[edit]
lab@R8# set routing-instances C6-vpls protocols vpls no-tunnel-services

[edit]
lab@R8# set routing-instances C6-vpls protocols vpls encapsulation-type ethernet-vlan

[edit]
lab@R8# set routing-instances C6-vpls protocols vpls vpls-id 600

[edit]
lab@R8# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.1

[edit]
lab@R8# set routing-instances C6-vpls protocols vpls neighbor 172.30.5.6

[edit]
lab@R8# commit
commit complete
```

On R8 no configuration for multihoming is applied, similar to R7.

Verify the operation of customer C6 LDP VPLS on R1.

```
[edit]
lab@R1# run show ldp neighbor
Address            Interface          Label space ID         Hold time
172.30.5.6         lo0.0              172.30.5.6:0             43
172.30.5.7         lo0.0              172.30.5.7:0             37
172.30.5.8         lo0.0              172.30.5.8:0             41
172.30.0.2         ae0.0              172.30.5.2:0             14
172.30.0.6         ge-0/0/4.114       172.30.5.4:0             12
```

Extended LDP sessions are established to R6, R7 and R8.

```
[edit]
lab@R1# run show vpls connections
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
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
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            HS -- Hot-standby Connection

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

Instance: C6-vpls
  VPLS-id: 600
    Neighbor                  Type  St     Time last up          # Up trans
    172.30.5.6(vpls-id 600)   rmt   Up     Apr  4 18:29:45 2019           1
      Remote PE: 172.30.5.6, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262145
      Negotiated PW status TLV: No
      Local interface: lsi.1048576, Status: Up, Encapsulation: VLAN
        Description: Intf - vpls C6-vpls neighbor 172.30.5.6 vpls-id 600
      Flow Label Transmit: No, Flow Label Receive: No
    172.30.5.7(vpls-id 600)   rmt   Up     Apr  4 18:31:07 2019           1
      Remote PE: 172.30.5.7, Negotiated control-word: No
      Incoming label: 262146, Outgoing label: 262145
      Negotiated PW status TLV: No
      Local interface: lsi.1048577, Status: Up, Encapsulation: VLAN
        Description: Intf - vpls C6-vpls neighbor 172.30.5.7 vpls-id 600
      Flow Label Transmit: No, Flow Label Receive: No
    172.30.5.8(vpls-id 600)   rmt   BK
```

Neighbor R6 and R7 are UP, but neighbor R8 is in state `BK -- Backup connection`. This is expected because neighbor R8 is configured as backup for neighbor R7.

```
[edit]
lab@R1# run show ldp database l2circuit session 172.30.5.6
Input label database, 172.30.5.1:0--172.30.5.6:0
Labels received: 11
  Label     Prefix
 262145      L2CKT NoCtrlWord VLAN VC 600
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 363872      L2CKT CtrlWord VLAN VC 513

Output label database, 172.30.5.1:0--172.30.5.6:0
Labels advertised: 10
   Label     Prefix
 262145      L2CKT NoCtrlWord VLAN VC 600
 378880      L2CKT CtrlWord VLAN VC 513
```

```
[edit]
lab@R1# run show ldp database l2circuit session 172.30.5.7
Input label database, 172.30.5.1:0--172.30.5.7:0
Labels received: 10
   Label     Prefix
 262145      L2CKT NoCtrlWord VLAN VC 600

Output label database, 172.30.5.1:0--172.30.5.7:0
Labels advertised: 9
   Label     Prefix
 262146      L2CKT NoCtrlWord VLAN VC 600
```

```
[edit]
lab@R1# run show ldp database l2circuit session 172.30.5.8
Input label database, 172.30.5.1:0--172.30.5.8:0
Labels received: 11
   Label     Prefix
 262145      L2CKT NoCtrlWord VLAN VC 600
 371904      L2CKT CtrlWord VLAN VC 512

Output label database, 172.30.5.1:0--172.30.5.8:0
Labels advertised: 9
   Label     Prefix
 378864      L2CKT CtrlWord VLAN VC 512
```

The Input and Output label are exchanged for VC 600 between the neighbors.

Note: The output label to R8 is not sent, because R8 is the backup neighbor of R7. R8 sends a label because R1 and R6 are configured as regular neighbors in the VPLS routing instance.

```
[edit]
lab@R1# run show route forwarding-table family vpls
Routing table: C6-vpls.vpls
VPLS:
Enabled protocols: ACKed by all peers,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
default             perm     0                      dscd      769     1
lsi.1048576         intf     0                      indr  1048579     6
                                                    ulst  1048578     2
                                 172.30.0.6         Push 262145, Push 338064(top)
650     2 ge-0/0/4.114
                                 172.30.0.6         Push 262145, Push 337312(top)
655     2 ge-0/0/4.114
                                 172.30.0.10        Push 262145, Push 349936(top)
656     2 ge-0/0/4.118
                                 172.30.0.10        Push 262145, Push 348752(top)
659     2 ge-0/0/4.118
lsi.1048577         intf     0                      indr  1048576     6
...
```

The forwarding table of the customer VPLS table is shown above.

Verify the operation of customer C6 LDP VPLS on R7.

```
[edit]
lab@R7# run show vpls connections
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           HS -- Hot-standby Connection

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

Instance: C6-vpls
  VPLS-id: 600
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Neighbor                       Type  St     Time last up          # Up trans
172.30.5.1(vpls-id 600)        rmt   Up     Apr  4 18:31:02 2019           1
  Remote PE: 172.30.5.1, Negotiated control-word: No
  Incoming label: 262145, Outgoing label: 262146
  Negotiated PW status TLV: No
  Local interface: lsi.1048576, Status: Up, Encapsulation: VLAN
    Description: Intf - vpls C6-vpls neighbor 172.30.5.1 vpls-id 600
  Flow Label Transmit: No, Flow Label Receive: No
172.30.5.6(vpls-id 600)        rmt   Up     Apr  4 18:31:02 2019           1
  Remote PE: 172.30.5.6, Negotiated control-word: No
  Incoming label: 262146, Outgoing label: 262146
  Negotiated PW status TLV: No
  Local interface: lsi.1048577, Status: Up, Encapsulation: VLAN
    Description: Intf - vpls C6-vpls neighbor 172.30.5.6 vpls-id 600
  Flow Label Transmit: No, Flow Label Receive: No
```

On R7 only neighbor R1 and neighbor R6 are in the UP state. The neighbor with R8 is not present as it is a backup site.

Verify the operation of the customer C6 LDP VPLS on R8.

```
[edit]
lab@R8# run show vpls connections
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           HS -- Hot-standby Connection

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


© 2020 Juniper Networks, Inc. All rights reserved.

```
Instance: C6-vpls
  VPLS-id: 600
    Neighbor                      Type  St     Time last up          # Up trans
    172.30.5.1(vpls-id 600)       rmt   OL
    172.30.5.6(vpls-id 600)       rmt   OL
```

On R8, the neighbors with R1 and R6 are in the state: `OL -- no outgoing label`. R8 is sending labels as you saw in the LDP database on R1, however on R1 neighbor R8 is configured as a backup of R7 and no label is sent until the neighborship with R7 is down.

Verify the redundancy with manual deactivation of R7s VPLS routing instance.

```
[edit]
lab@R7# deactivate routing-instances C6-vpls

[edit]
lab@R7# commit
commit complete
```

Verify the vpls connection table on R8 again.

```
 [edit]
lab@R8# run 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           HS -- Hot-standby Connection

Legend for interface status
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Up -- operational
Dn -- down

Instance: C6-vpls
  VPLS-id: 600
    Number of local interfaces: 2
    Number of local interfaces up: 2
    ge-0/0/3.700
    ge-0/0/3.701
    lsi.1048577                  Intf - vpls C6-vpls neighbor 172.30.5.1 vpls-id
600
    lsi.1048576                  Intf - vpls C6-vpls neighbor 172.30.5.6 vpls-id
600
    Neighbor                     Type  St     Time last up          # Up trans
    172.30.5.1(vpls-id 600)      rmt   Up     Apr  4 18:47:59 2019           1
      Remote PE: 172.30.5.1, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262147
      Negotiated PW status TLV: No
      Local interface: lsi.1048577, Status: Up, Encapsulation: VLAN
        Description: Intf - vpls C6-vpls neighbor 172.30.5.1 vpls-id 600
      Flow Label Transmit: No, Flow Label Receive: No
    Connection History:
        Apr  4 18:47:59 2019  status update timer
        Apr  4 18:47:59 2019  PE route changed
        Apr  4 18:47:59 2019  Out lbl Update                   262147
        Apr  4 18:47:59 2019  In lbl Update                    262145
        Apr  4 18:47:59 2019  loc intf up                 lsi.1048577
    172.30.5.6(vpls-id 600)      rmt   Up     Apr  4 18:47:58 2019           1
      Remote PE: 172.30.5.6, Negotiated control-word: No
      Incoming label: 262146, Outgoing label: 262147
      Negotiated PW status TLV: No
      Local interface: lsi.1048576, Status: Up, Encapsulation: VLAN
        Description: Intf - vpls C6-vpls neighbor 172.30.5.6 vpls-id 600
      Flow Label Transmit: No, Flow Label Receive: No
    Connection History:
        Apr  4 18:47:58 2019  status update timer
        Apr  4 18:47:58 2019  PE route changed
        Apr  4 18:47:58 2019  Out lbl Update                   262147
        Apr  4 18:47:58 2019  In lbl Update                    262146
        Apr  4 18:47:58 2019  loc intf up                 lsi.1048576
```

Now both neighbors to R1 and R6 are in UP state.

Rollback your configuration changes on R7.

```
[edit]
lab@R7# rollback 1
load complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# commit
commit complete
```

6) Configure the MAC table size for customer C5 VPLS on **R2**, **R3**, **R4** and **R5**.

Providing VPLS services for many customers could cause a PE device to run out of its resources, due to a high number of MAC table entries learned. Junos provides a capability for limiting the number of MAC address table entries per routing instance.

By default, each routing instance is allowed to have 512 MAC addresses in the table. By default, If the table is full, then packets with an unknown source MAC address are flooded.

On **R2**, **R3**, **R4**, and **R5:**

```
[edit]
lab@R2# set routing-instances C5-vpls protocols vpls mac-table-size 200

[edit]
lab@R2# set routing-instances C5-vpls protocols vpls interface-mac-limit 200

[edit]
lab@R2# commit
commit complete
```

7) Configure a MAC table size limit for customer C6 on **R1**, **R6**, **R7** and **R8**.

For customer C6 when the mac-limit is reached, new packets with an unknown MAC address must be dropped.

On **R1**, **R6**, **R7**, and **R8**:

```
[edit]
lab@R1# set routing-instances C6-vpls protocols vpls mac-table-size 100

[edit]
lab@R1# set routing-instances C6-vpls protocols vpls interface-mac-limit 100

[edit]
lab@R1# set routing-instances C6-vpls protocols vpls mac-table-size packet-action drop

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

8) Configure customer C5 L2VPN and VPLS interworking on R7.

Customer C5 has the requirement to stitching the L2VPN service with the VPLS service. VLAN600 on CE5-1 attached to R7 should be transported over the BGP L2VPN and connected to the VPLS routing instance on R2.

Configure the additional interface on R7 in VLAN600.

```
[edit]
lab@R7# set interfaces ge-0/0/3.600 vlan-id 600

[edit]
lab@R7# set interfaces ge-0/0/3.600 encapsulation vlan-ccc

[edit]
lab@R7# commit
commit complete
```

Modify the L2VPN routing instance.

```
[edit]
lab@R7# set routing-instances C5-l2vpn interface ge-0/0/3.600

[edit]
lab@R7# set routing-instances C5-l2vpn protocols l2vpn site site-1 interface ge-
0/0/3.600
```

R7 local site is 1, this means that this interface will be connected to the first interface on site 4. Remember that the place of the interface in the hierarchy is important for the auto-discovery of remote sites.

```
[edit]
lab@R7# commit
commit complete
```

9) Configure customer C5 L2VPN and VPLS internetworking on R2.

R2 is where the interconnection between the BGP L2VPN and the BGP VPLS routing instance will be stitched. You need an interface to terminate the BGP L2VPN on VLAN600 and loop it to the VPLS instance. One way is to use a physical loop between two interfaces on the device. Junos also provides an internal mechanism for interconnecting routing instances with a logical tunnel interface.

Logical tunnel interface units accept different encapsulations and families. Two logical tunnel units must be peered to establish the communication between them

© 2020 Juniper Networks, Inc. All rights reserved.

a) Configure an lt- interface.

You need two logical (lt-) tunnel interface units, one with `family vlan-ccc` inside the L2VPN routing instance and the second one with `family vlan-vpls` inside the VPLS routing instance.

```
[edit]
lab@R2# set interfaces lt-0/0/0.0 vlan-id 600
```

Both lt- units are configured to accept 802.1q tagged packets.

```
[edit]
lab@R2# set interfaces lt-0/0/0.0 encapsulation vlan-ccc
```

The `unit 0` is configured for `vlan-ccc` encapsulation. This unit will be attached to the L2VPN routing instance.

```
[edit]
lab@R2# set interfaces lt-0/0/0.0 peer-unit 1
```

The `unit 0` is peered with `unit 1`.

```
[edit]
lab@R2# set interfaces lt-0/0/0.1 vlan-id 600

[edit]
lab@R2# set interfaces lt-0/0/0.1 encapsulation vlan-vpls
```

The `unit 1` is configured for `vlan-vpls` encapsulation. This unit will be attached to the VPLS routing instance.

```
[edit]
lab@R2# set interfaces lt-0/0/0.1 peer-unit 0
```

Both units are peered with each other.

On R7 L2VPN routing instance was configured in Step 1, but on R2 a new routing instance must be configured to terminate the remote site.

```
[edit]
lab@R2# set routing-instances C5-l2vpn instance-type l2vpn

[edit]
lab@R2# set routing-instances C5-l2vpn interface lt-0/0/0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

The `unit 0` of **`lt-`** is the only interface assigned to that routing instance

```
[edit]
lab@R2# set routing-instances C5-l2vpn vrf-target target:54591:500
```

The route target matches the route target configured on R7.

```
[edit]
lab@R2# set routing-instances C5-l2vpn protocols l2vpn encapsulation-type ethernet-vlan

[edit]
lab@R2# set routing-instances C5-l2vpn protocols l2vpn site site-4 site-identifier 4
```

The local site identifier is configured as 4.

```
[edit]
lab@R2# set routing-instances C5-l2vpn protocols l2vpn site site-4 interface lt-0/0/0.0
```

`lt-0/0/0.0` is the only PE-CE interface for this site. It will be connected to the fourth interface on R7 site 1.

```
[edit]
lab@R2# commit
commit complete
```

Verify the l2vpn connection table.

```
[edit]
lab@R2# run show l2vpn 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
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
RS -- remote site standby          SN -- Static Neighbor
LB -- Local site not best-site     RB -- Remote site not best-site
VM -- VLAN ID mismatch             HS -- Hot-standby Connection

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

Instance: C5-l2vpn
Edge protection: Not-Primary
  Local site: site-4 (4)
    Number of local interfaces: 1
    Number of local interfaces up: 1
    lt-0/0/0.0          1
    Label-base        Offset     Size  Range     Preference
    800000            1          2     1         100
      status-vector:  0
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Apr  5 08:21:11 2019           1
      Remote PE: 172.30.5.7, Negotiated control-word: Yes (Null)
      Incoming label: 800000, Outgoing label: 800007
      Local interface: lt-0/0/0.0, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
    Connection History:
        Apr  5 08:21:11 2019  status update timer
        Apr  5 08:21:11 2019  PE route changed
        Apr  5 08:21:11 2019  Out lbl Update                   800007
        Apr  5 08:21:11 2019  In lbl Update                    800000
        Apr  5 08:21:11 2019  loc intf up                  lt-0/0/0.0
    2                         rmt   CN
    3                         rmt   CN
```

The connection with R7 is in the **UP** state. The local interface `lt-0/0/0.0` is configured.

The other two sites are in state `OR -- out of range`, which means that the calculated label is out of the range of the remote site.

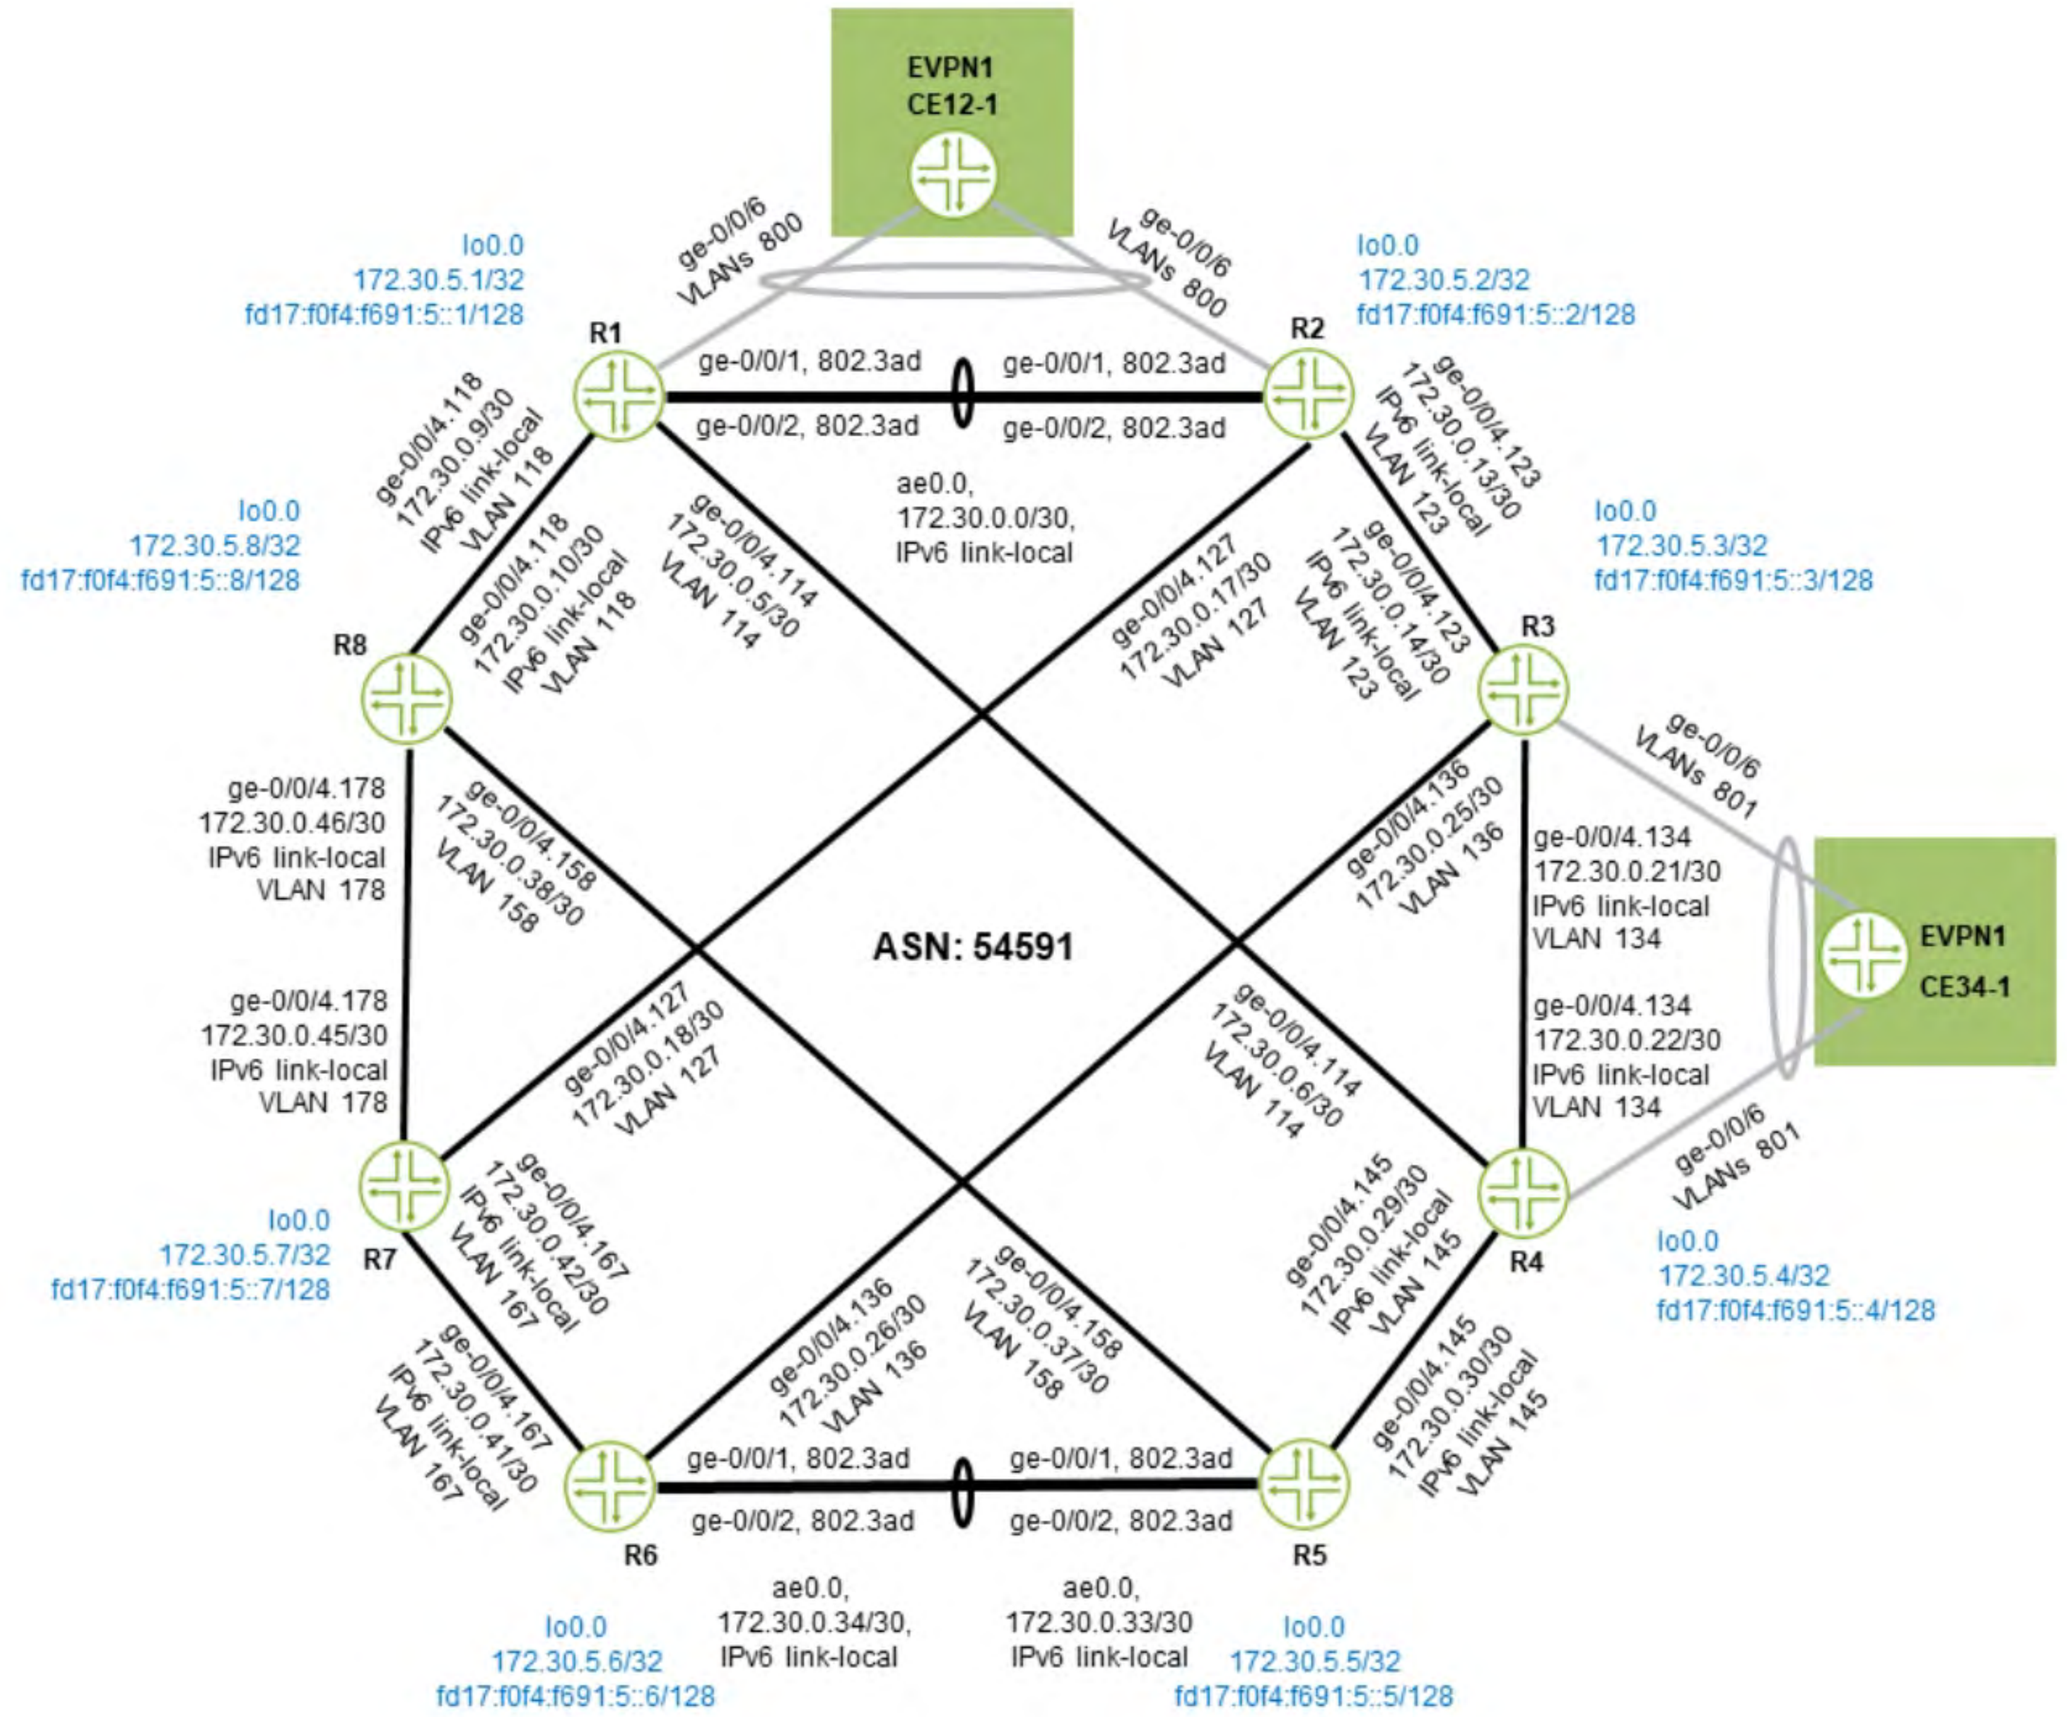
The last step to configure is the `lt-0/0/0.1` for the C5-vpls routing instance. This will stitch the L2VPN routing instance with the VPLS routing instance using a logical tunnel interface.

```
[edit]
lab@R2# set routing-instances C5-vpls interface lt-0/0/0.1

[edit]
lab@R2# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 3. EVPN Configuration



© 2020 Juniper Networks, Inc. All rights reserved.

CE Router	CE L3 address	PE router	CE facing interface	LACP	Vlan ID	Default gateway
CE12-1	192.168.2.12/24	R1	ae1.800	Enabled	800	192.168.2.254/24
		R2				
CE34-1	192.168.2.34/24	R3	ae1.801		801	
		R4				

Table 14

1) Configure an EVPN service between site1 and site 2. Ensure that both sites support active-active multihoming for unicast traffic. Use the information that is provided in Table 14.

Enable EVPN signaling by exchanging EVPN NLRI in BGP for the required routers.

Route Reflector

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family evpn signaling

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family evpn signaling

[edit]
lab@route-reflector# commit
commit complete
```

R1

```
[edit]
lab@R1# set protocols bgp group ibgp family evpn signaling

[edit]
lab@R1# commit
commit complete
```

R2

```
[edit]
lab@R2# set protocols bgp group ibgp family evpn signaling

[edit]
lab@R2# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

R3

```
[edit]
lab@R3# set protocols bgp group ibgp family evpn signaling

[edit]
lab@R3# commit
commit complete
```

R4

```
[edit]
lab@R4# set protocols bgp group ibgp family evpn signaling

[edit]
lab@R4# commit
commit complete
```

Verify the exchange of EVPN NLRI in BGP.

```
lab@R1# run show bgp summary
Groups: 5 Peers: 6 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
…
…
…
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591        598        578       0       1     1:06:13 Establ
  bgp.rtarget.0: 11/17/17/0
  inet.0: 216/216/216/0
  inet6.0: 32/32/32/0
  bgp.l3vpn.0: 16/16/16/0
  bgp.mvpn.0: 4/4/4/0
  C2-spoke.inet.0: 12/16/16/0
  C2-spoke.mvpn.0: 4/4/4/0
  bgp.evpn.0: 0/0/0/0
...
```

At this stage it is expected that no EVPN routes are exchanged as no EVPN service is configured yet.

Configure the aggregate (member) interfaces and the IRB interface on R1, R2, R3 and R4. To support active-active multihoming in EVPN we need to make sure that the “all-active” option is configured. Also make sure the LACP system id match and the ESI values match

© 2020 Juniper Networks, Inc. All rights reserved.

on each ethernet segment. The ESI value is signaled in BGP to ensure BUM traffic is not forwarded to ethernet segments that have the same ESI value configured. For this reason, the ESI values should match on each multihomed ethernet segment.

Note: Because we have a group defining interface families, we will need to alter this group configuration slightly to avoid applying the MPLS and ISO families to ae1.

R1, R2:

```
set chassis aggregated-devices ethernet device-count 2
rename groups if-families interfaces <ae*> to <ae0*>

set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 esi 00:00:00:00:00:00:00:00:00:12
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:12
set interfaces ae1 unit 800 encapsulation vlan-bridge
set interfaces ae1 unit 800 vlan-id 800
set interfaces irb unit 1 family inet address 192.168.2.254/24

set interfaces ge-0/0/6 gigether-options 802.3ad ae1

set routing-instances evpn1 instance-type evpn
set routing-instances evpn1 vlan-id 800
set routing-instances evpn1 interface ae1.800
set routing-instances evpn1 routing-interface irb.1
set routing-instances evpn1 vrf-target target:54591:1000
set routing-instances evpn1 protocols evpn interface ae1.800
```

R1:

```
set routing-instances evpn1 route-distinguisher 172.30.5.1:1000
commit
```

R2:

```
set routing-instances evpn1 route-distinguisher 172.30.5.2:1000
commit
```

R3, R4:

```
set chassis aggregated-devices ethernet device-count 2
rename groups if-families interfaces <ae*> to <ae0*>
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 esi 00:00:00:00:00:00:00:00:00:34
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:34
set interfaces ae1 unit 801 encapsulation vlan-bridge
set interfaces ae1 unit 801 vlan-id 801
set interfaces irb unit 1 family inet address 192.168.2.254/24

set interfaces ge-0/0/6 gigether-options 802.3ad ae1

set routing-instances evpn1 instance-type evpn
set routing-instances evpn1 vlan-id 800
set routing-instances evpn1 interface ae1.801
set routing-instances evpn1 routing-interface irb.1
set routing-instances evpn1 vrf-target target:54591:1000
set routing-instances evpn1 protocols evpn interface ae1.801
```

R3:

```
set routing-instances evpn1 route-distinguisher 172.30.5.3:1000
commit
```

R4:

```
set routing-instances evpn1 route-distinguisher 172.30.5.4:1000
commit
```

Verification on R1:

Interface ae1 should be up and LACP PDUs should be send and received.

```
lab@R1# run show interfaces terse | match ae1
ge-0/0/6.800                up    up   aenet    --> ae1.800
ge-0/0/6.32767              up    up   aenet    --> ae1.32767
ae1                         up    up
ae1.800                     up    up   bridge
ae1.32767                   up    up   multiservice

lab@R1# run show lacp statistics interfaces ae1
Aggregated interface: ae1
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-0/0/6                1292        1275            0            0
```

Verify if EVPN NLRI is exchanged

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# run show bgp summary group ibgp | match evpn
bgp.evpn.0
  bgp.evpn.0: 20/20/20/0
  evpn1.evpn.0: 17/17/17/0
  __default_evpn__.evpn.0: 3/3/3/0
```

R2, R3 and R4 should provide similar outputs.

Verify connectivity between CE12-1 and CE34-4. Also verify reachability towards the default gateway.

VR-DEVICE2:

```
root@vr-device-2# run show route table CE12-1.inet.0

CE12-1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:01:24
                    > to 192.168.2.254 via ae1.0
192.168.2.0/24     *[Direct/0] 00:01:24
                    > via ae1.0
192.168.2.12/32    *[Local/0] 00:01:24
                      Local via ae1.0

root@vr-device-2# run show route table CE34-1.inet.0

CE34-1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:01:35
                    > to 192.168.2.254 via ae2.0
192.168.2.0/24     *[Direct/0] 00:01:35
                    > via ae2.0
192.168.2.34/32    *[Local/0] 00:01:35
                      Local via ae2.0
```

Ping CE34-1 to verify end to end connectivity

```
root@vr-device-2# run ping 192.168.2.34 routing-instance CE12-1
PING 192.168.2.34 (192.168.2.34): 56 data bytes
64 bytes from 192.168.2.34: icmp_seq=0 ttl=64 time=5.830 ms
64 bytes from 192.168.2.34: icmp_seq=1 ttl=64 time=5.424 ms
```

Verify reachability towards the default gateway

© 2020 Juniper Networks, Inc. All rights reserved.

```
root@vr-device-2# run ping 192.168.2.254 routing-instance CE12-1
PING 192.168.2.254 (192.168.2.254): 56 data bytes
64 bytes from 192.168.2.254: icmp_seq=0 ttl=64 time=4.837 ms
64 bytes from 192.168.2.254: icmp_seq=1 ttl=64 time=4.380 ms
```

Verify the EVPN database:

R1:

```
lab@R1# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                  Timestamp        IP
address
800             2c:6b:f5:10:01:f0  172.30.5.3                     Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:1a:b8:f0  172.30.5.2                     Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:38:46:c0  00:00:00:00:00:00:00:00:00:12  Mar 05 14:44:28
192.168.2.12
800             2c:6b:f5:38:46:c1  00:00:00:00:00:00:00:00:00:34  Mar 05 14:44:27
192.168.2.34
800             2c:6b:f5:54:61:f0  172.30.5.4                     Mar 05 14:44:32
192.168.2.254
800             2c:6b:f5:86:c6:f0  irb.1                          Mar 05 14:44:24
192.168.2.254
```

R2:

```
lab@R2# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                  Timestamp        IP
address
800             2c:6b:f5:10:01:f0  172.30.5.3                     Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:1a:b8:f0  irb.1                          Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:38:46:c0  00:00:00:00:00:00:00:00:00:12  Mar 05 14:44:28
192.168.2.12
800             2c:6b:f5:38:46:c1  00:00:00:00:00:00:00:00:00:34  Mar 05 14:44:27
192.168.2.34
800             2c:6b:f5:54:61:f0  172.30.5.4                     Mar 05 14:44:32
192.168.2.254
800             2c:6b:f5:86:c6:f0  172.30.5.1                     Mar 05 14:44:25
192.168.2.254
```

R3:

© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                      Timestamp        IP
address
800             2c:6b:f5:10:01:f0  irb.1                              Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:1a:b8:f0  172.30.5.2                         Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:38:46:c0  00:00:00:00:00:00:00:00:00:12      Mar 05 14:44:28
192.168.2.12
800             2c:6b:f5:38:46:c1  00:00:00:00:00:00:00:00:00:34      Mar 05 14:44:26
192.168.2.34
800             2c:6b:f5:54:61:f0  172.30.5.4                         Mar 05 14:44:31
192.168.2.254
800             2c:6b:f5:86:c6:f0  172.30.5.1                         Mar 05 14:44:24
192.168.2.254
```

R4:

```
lab@R4# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                      Timestamp        IP
address
800             2c:6b:f5:10:01:f0  172.30.5.3                         Mar 05 14:44:33
192.168.2.254
800             2c:6b:f5:1a:b8:f0  172.30.5.2                         Mar 05 14:44:32
192.168.2.254
800             2c:6b:f5:38:46:c1  00:00:00:00:00:00:00:00:00:34      Mar 05 14:44:33
192.168.2.34
800             2c:6b:f5:54:61:f0  irb.1                              Mar 05 14:44:23
192.168.2.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

Appendix - Chapter 7: Inter-Provider VPN Configuration

Inter-Provider VPN Overview

RFC4364 is the updated RFC for BGP/MPLS IP Virtual Private Networks. The Multi-AS Backbone section defines three options for Inter-AS VPN configuration. If a customer requires VPN connectivity across both providers ASs one of these three options can be used.

- *Option A* – VRF-to-VRF at AS border routers. All the VPN VRFs that need inter-as connectivity are configured on the AS border routers. Unlabeled traffic is exchanged between the routers. This option is least the scalable, however easy to configure and secure.
- *Option B* – a multiprotocol EBGP session is established between AS border routers. VPNv4 routes with VPN labels are exchanged over that session. Different LSPs transport the traffic between the ASBR and the PE in each AS. Care should be taken when selecting route targets to be used. Scalability might also be an issue, because all routes are kept in the bgp.l3vpn.0 table.
- *Option C* – This most scalable solution. EBGP session with family labeled-unicast is established between the ASBR routers exchanging loopback IP addresses of the PE routers between both ASs. This "stitches" the LSP between the ASBRs. a Multihop EBGP session is established directly between the remote PE routes to exchange the VPN labels.

Task 1: Inter-Provider VPN Option B

Customer C2 has a remote site attached the neighboring AS 43208.365. The requirement is to connect that site to your AS with Inter-Provider VPN option B.

Neighbor P3-1 is the AS border router for AS 43208.365. It is connected to R3 in AS 54591. An EBGP session with `family inet-vpn unicast` is needed between P3-1 and R3.

1) Configure family MPLS on the P3-1 facing interface on R3.

Traffic between both AS border routers will be label switched. On the AS borders, the transport label will be popped, but the VPN label will flow between both ASs.

```
[edit]
lab@R3# set interfaces ge-0/0/5.302 family mpls
```

2) Configure the P3-1 facing interface under the `protocols mpls` hierarchy.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols mpls interface ge-0/0/5.302
```

3) Configure family inet-vpn unicast on P3-1 EBGP session.

```
[edit]
lab@R3# set protocols bgp group P3-1 family inet-vpn unicast

[edit]
lab@R3# commit
commit complete
```

Verify the BGP session table.

```
[edit]
lab@R3# run show bgp summary
Groups: 5 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     546        518          0          0          0          0
bgp.l3vpn.0
                      22         21          0          0          0          0
inet6.0
                      48         48          0          0          0          0
bgp.l2vpn.0
                       6          6          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591        413         95       0       1          25
Establ
  inet.0: 449/449/449/0
  bgp.l3vpn.0: 19/19/19/0
  inet6.0: 32/32/32/0
  bgp.l3vpn-inet6.0: 10/10/10/0
  bgp.l2vpn.0: 6/6/6/0
  bgp.rtarget.0: 13/17/17/0
  C3.inet6.0: 9/10/10/0
  C1.inet.0: 9/18/18/0
  C3.inet.0: 1/1/1/0
  C5-l2vpn.l2vpn.0: 3/3/3/0
  C5-vpls.l2vpn.0: 3/3/3/0
192.168.0.2      3521382357         129        428       0       0        22:50
Establ
  inet.0: 69/97/69/0
192.168.0.6      2831679853           5         36       0       0           45
Establ
  bgp.l3vpn.0: 2/3/2/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
fc09:c0:ffee::a       64601          52          100        0        0        22:38
Establ
  C3.inet6.0: 8/8/8/0
fe80::223:9c01:2d8b:6c81  3521382357          52          60        0        0
22:27 Establ
  inet6.0: 16/16/16/0
```

The `bgp.l3vpn.0` routing table receives routes from neighbor R3, however `family inet unicast` is missing. Remember that the `family inet unicast` is enabled by default, but adding other families disables the default behavior and `family inet unicast` must be enabled explicitly.

```
[edit]
lab@R3# set protocols bgp group P3-1 family inet unicast

[edit]
lab@R3# commit
commit complete
```

```
[edit]
lab@R3# run show bgp summary
Groups: 5 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     660        598          0          0          0          0
bgp.l3vpn.0
                      22         21          0          0          0          0
inet6.0
                      48         48          0          0          0          0
bgp.l2vpn.0
                       6          6          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591        498        236       0       1        2:52
Establ
  inet.0: 447/447/447/0
  bgp.l3vpn.0: 19/19/19/0
  inet6.0: 32/32/32/0
  bgp.l3vpn-inet6.0: 10/10/10/0
  bgp.l2vpn.0: 6/6/6/0
  bgp.rtarget.0: 13/17/17/0
  C3.inet6.0: 9/10/10/0
  C1.inet.0: 9/18/18/0
  C3.inet.0: 1/1/1/0
  C5-l2vpn.l2vpn.0: 3/3/3/0
  C5-vpls.l2vpn.0: 3/3/3/0
192.168.0.2      3521382357        135        852       0       0       25:17
Establ
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  inet.0: 69/97/69/0
192.168.0.6       2831679853          92         431        0          0           15
Establ
  inet.0: 82/116/82/0
  bgp.l3vpn.0: 2/3/2/0
fc09:c0:ffee::a        64601          57         113        0          0        25:05
Establ
  C3.inet6.0: 8/8/8/0
fe80::223:9c01:2d8b:6c81  3521382357          57         71         0          0
24:54 Establ
  inet6.0: 16/16/16/0
```

4) Modify the BGP import policy for neighbor P3-1.

So far, the policy is only taking the `family inet unicast` prefixes into account. Now you must modify the policy to also allow `family inet-vpn unicast` prefixes.

```
lab@R3# show policy-options policy-statement P3-filter
term 1 {
    from {
        as-path P3-local-routes;
        route-filter 0.0.0.0/0 prefix-length-range /32-/32;
    }
    then {
        local-preference 200;
        community set P3;
        accept;
    }
}
term 2 {
    from {
        route-filter 0.0.0.0/0 prefix-length-range /8-/24;
    }
    then {
        local-preference 200;
        community set P3;
        accept;
    }
}
term 3 {
    then reject;
}
```

See above how the current import policy looks. You must restrict term 2 and term 3 for `family inet` only. You will also need to restrict term 1 to apply only to BGP routes.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set policy-options policy-statement P3-filter term 2 from family inet

[edit]
lab@R3# set policy-options policy-statement P3-filter term 1 from protocol bgp

[edit]
lab@R3# set policy-options policy-statement P3-filter term 3 from family inet

[edit]
lab@R3# commit
commit complete
```

Let's verify the routes received from neighbor 192.168.0.6

```
[edit]
lab@R3# run show route receive-protocol bgp 192.168.0.6 table bgp.l3vpn.0

bgp.l3vpn.0: 34 destinations, 34 routes (34 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.17.47.2:200:172.31.78.0/24
*                         192.168.0.6                             2831679853
64600 I
  172.17.47.2:200:172.31.79.0/24
*                         192.168.0.6                             2831679853
64600 I
  172.17.47.2:200:192.168.0.100/30
*                         192.168.0.6                             2831679853 I
```

Three prefixes are received and accepted from neighbor P3-1 in the `bgp.l3vpn.0` table.

The routes in `bgp.l3vpn.0` match the default BGP export policy and are advertised to P3-1.

```
[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.6 table bgp.l3vpn.0

bgp.l3vpn.0: 41 destinations, 41 routes (41 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.30.5.3:4:0.0.0.0/0
*                         Self                                    I
  172.30.5.3:4:172.30.5.17/32
*                         Self                                    I
  172.30.5.3:4:172.31.48.0/30
*                         Self                 6                  I
  172.30.5.3:4:172.31.48.4/30
*                         Self                 2                  I
  172.30.5.3:4:172.31.48.8/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
*                              Self                2                   I
  172.30.5.3:4:172.31.48.12/30
*                              Self                3                   I
...
```

The VPNv4 routes are advertised to P3-1.

5) Configure the BGP route target family `advertise-default` option on R3.

Verify if the remote VPNv4 routes are advertised to the route reflector.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.41 table bgp.l3vpn.0

bgp.l3vpn.0: 34 destinations, 34 routes (34 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.30.5.3:4:0.0.0.0/0
*                         Self                         100        I
  172.30.5.3:4:172.30.5.17/32
*                         Self                         100        I
  172.30.5.3:4:172.31.48.0/30
*                         Self                 6       100        I
  172.30.5.3:4:172.31.48.4/30
*                         Self                 2       100        I
  172.30.5.3:4:172.31.48.8/30
*                         Self                 2       100        I
  172.30.5.3:4:172.31.48.12/30
*                         Self                 3       100        I
  172.30.5.3:4:172.31.63.2/32
...
```

The routes received from P3-1 are missing from the announcements. The reason for this is that the BGP `family route-target` feature is configured. R3 does not receive route-target announcements for the VPNv4 routes from P3-1. So R3 suppresses the VPNv4 announces and does not send them to route reflector because no other PE requested these route targets.

Activating the `advertise-default` feature under `family route-target` on R3 sends a single default route-target route to the route reflector, which automatically send all routes from `bgp.l3vpn.0` regardless of the route-target value.

```
[edit]
lab@R3# set protocols bgp group ibgp family route-target advertise-default
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# commit
commit complete
```

To check the route target attached to the routes received from P3-1 you can use the "extensive" option for the show route command.

```
[edit]
lab@R3# run show route receive-protocol bgp 192.168.0.6 table bgp.l3vpn.0
extensive

bgp.l3vpn.0: 70 destinations, 70 routes (70 active, 0 holddown, 0 hidden)
* 172.17.47.2:200:172.31.78.0/24 (1 entry, 1 announced)
     Accepted
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 299792
     Nexthop: 192.168.0.6
     AS path: 2831679853 64600 I
     Communities: target:43208:200

* 172.17.47.2:200:172.31.79.0/24 (1 entry, 1 announced)
     Accepted
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 299792
     Nexthop: 192.168.0.6
     AS path: 2831679853 64600 I
     Communities: target:43208:200
...
```

Route target target:43208:200 is attached to the remote routes.

6) Normalize the VPN route target.

VPN C2 uses multiple route targets as the VPN is based on a hub and spoke topology. To normalize the route targets, you must change the extended community on the routes in each direction. The remote site must have the spoke community attached to the routes coming from AS 43208.365 and the hub route target must be set on the routes send to AS 43208.365

Let's create the policies.

```
[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-import term 1 from protocol
bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-import term 1 from community
CE2-remote

[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-import term 1 then community
delete CE2-remote
```

Delete the existing community.

```
[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-import term 1 then community
add CE2-spoke
```

Apply the new, normalized community.

```
[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-import term 1 then accept

[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-export term 1 from protocol
bgp

[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-export term 1 from community
CE2-hub

[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-export term 1 then community
delete CE2-hub

[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-export term 1 then community
add CE2-remote

[edit]
lab@R3# set policy-options policy-statement C2-vpn-target-export term 1 then accept
```

Configure needed extended communities.

```
[edit]
lab@R3# set policy-options community CE2-remote members target:43208:200

[edit]
lab@R3# set policy-options community CE2-spoke members target:54591:201

[edit]
lab@R3# set policy-options community CE2-hub members target:54591:200
```


© 2020 Juniper Networks, Inc. All rights reserved.

Apply the policy on EBGP session to P3-1.

```
[edit]
lab@R3# set protocols bgp group P3-1 import C2-vpn-target-import

[edit]
lab@R3# set protocols bgp group P3-1 export C2-vpn-target-export

[edit]
lab@R3# commit
commit complete
```

Verify the normalization of the route targets.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.41 table bgp.l3vpn.0
extensive

bgp.l3vpn.0: 68 destinations, 68 routes (68 active, 0 holddown, 0 hidden)
...
* 172.17.47.2:200:172.31.78.0/24 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 314128
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [54591] 2831679853 64600 I
     Communities: target:54591:201

* 172.17.47.2:200:172.31.79.0/24 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 314128
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [54591] 2831679853 64600 I
     Communities: target:54591:201
* 172.17.47.2:200:192.168.0.100/30 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 300864
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [54591] 2831679853 I
     Communities: target:54591:201
```


© 2020 Juniper Networks, Inc. All rights reserved.

The routes received from P3-1 and sent to the route reflector are normalized with the spoke route target.

```
[edit]
lab@R3# run show route advertising-protocol bgp 192.168.0.6 table
bgp.l3vpn.0 extensive

bgp.l3vpn.0: 34 destinations, 34 routes (34 active, 0 holddown, 0 hidden)
...
* 172.30.5.3:4:0.0.0.0/0 (1 entry, 1 announced)
 BGP group P3-1 type External
     Route Distinguisher: 172.30.5.3:4
     VPN Label: 16
     Nexthop: Self
     Flags: Nexthop Change
     AS path: [54591] I
     Communities: target:43208:200

* 172.30.5.3:4:172.30.5.17/32 (1 entry, 1 announced)
 BGP group P3-1 type External
     Route Distinguisher: 172.30.5.3:4
     VPN Label: 16
     Nexthop: Self
     Flags: Nexthop Change
     AS path: [54591] I
     Communities: target:43208:200

* 172.30.5.3:4:172.31.48.0/30 (1 entry, 1 announced)
 BGP group P3-1 type External
     Route Distinguisher: 172.30.5.3:4
     VPN Label: 16
     Nexthop: Self
     Flags: Nexthop Change
     MED: 6
     AS path: [54591] I
     Communities: target:43208:200 rte-type:0.0.0.0:1:0

* 172.30.5.3:4:172.31.48.4/30 (1 entry, 1 announced)
 BGP group P3-1 type External
     Route Distinguisher: 172.30.5.3:4
     VPN Label: 16
     Nexthop: Self
     Flags: Nexthop Change
     MED: 2
     AS path: [54591] I
     Communities: target:43208:200 rte-type:0.0.0.0:1:0
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

The remote route target is replaced on the advertised routes to P3-1.

7) Verify the operation of the Inter-provider VPN - Option B.

The routes received from P3-1 are coming with VPN label of 299792.

```
[edit]
lab@R3# run show route receive-protocol bgp 192.168.0.6 table bgp.l3vpn.0
extensive

bgp.l3vpn.0: 68 destinations, 68 routes (68 active, 0 holddown, 0 hidden)
* 172.17.47.2:200:172.31.78.0/24 (1 entry, 1 announced)
     Accepted
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 299792
     Nexthop: 192.168.0.6
     AS path: 2831679853 64600 I
     Communities: target:43208:200
...
```

Because the session is EBGP the next hop is changed. When the next hop is changed the label will also be changed.

Verify the label advertised to the route reflector.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.41 table bgp.l3vpn.0
extensive

bgp.l3vpn.0: 68 destinations, 68 routes (68 active, 0 holddown, 0 hidden)
...
* 172.17.47.2:200:172.31.78.0/24 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 314208
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
     AS path: [54591] 2831679853 64600 I
     Communities: target:54591:201
...
```

Label 314208 is advertised to the PE routers in AS 54591.

```
[edit]
lab@R3# run show route table mpls.0 label 314208
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
mpls.0: 43 destinations, 43 routes (43 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

314208             *[VPN/170] 00:00:39
                    > to 192.168.0.6 via ge-0/0/5.302, Swap 299792
```

When traffic with VPN label 314208 arrives at R3 it is swapped with 299792 and send to P3-1.

Task 2: Inter-Provider VPN Option C

Customer C5 VPLS service is extended to a remote site attached to neighboring AS 43208.365. You must use Inter-provider Option C for this task.

This method is most scalable Inter Provider VPN option, because no VPNv4 routes are stored on the AS border routers. a labeled IPv4 EBGP session is established between the ASs. Loopback IP addresses are exchanged over this session to setup the transport LSPs between both ASs. Having an end-to-end transport LSP between the PEs in both ASs, you can create a BGP session with `family inet-vpn unicast enabled` to exchange the VPNv4 routes and VPN label. As customer C5 is a VPLS you need to enable `family l2vpn signaling`, because the VPLS service should be extended over both ASs.

To eliminate the need for an EBGP session between each PE, an EBGP session will be established with the route reflector.

1) Configure BGP family labeled unicast on all local VPLS PEs and the route reflector.

```
[edit]
lab@R2# set protocols bgp group ibgp family inet labeled-unicast rib inet.3

[edit]
lab@R2# commit
commit complete
```

Adding `rib inet.3` instructs Junos to associate `inet.3` as a primary route table for storing labeled BGP routes and exporting labeled routes from that table.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family inet labeled-unicast rib
inet.3

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family inet labeled-unicast rib
inet.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# commit
commit complete
```

The route reflector is also configured for the BGP `family inet labeled unicast`.

On R3, you must enable `family inet labeled-unicast` BGP on the EBGP session to neighbor P3-1. This is the session used to interconnect both ASs.

```
[edit]
lab@R3# set protocol bgp group ibgp family inet labeled-unicast rib inet.3

[edit]
lab@R3# set protocol bgp group P3-1 family inet labeled-unicast rib inet.3

[edit]
lab@R3# commit
commit complete
```

You must change the P3-filter policy to accept the labeled BGP routes from P3-1.

```
[edit]
lab@R3# set policy-options policy-statement P3-filter term 3 from family inet

[edit]
lab@R3# commit
commit complete
```

Verify the received labeled BGP route, which is the remote PE loopback IP address from P3-1.

```
[edit]
lab@R3# run show route 172.17.47.3

inet.0: 762 destinations, 766 routes (691 active, 0 holddown, 73 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.47.3/32     *[BGP/170] 00:00:02, localpref 100
                      AS path: 2831679853 I
                    > to 192.168.0.6 via ge-0/0/5.302

inet.3: 16 destinations, 22 routes (10 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.47.3/32     *[BGP/170] 00:00:02, localpref 100
                      AS path: 2831679853 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 192.168.0.6 via ge-0/0/5.302
...
```

2) Modify the P3-1 export policy on R3.

You must modify the export policy to advertise the loopback IP addresses from the `inet.3` table to P3-1. This is needed to stitch the labeled path between the PE routers in both ASs.

```
[edit]
lab@R3# run show route table inet.3

inet.3: 16 destinations, 22 routes (10 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.47.3/32     *[BGP/170] 00:18:49, localpref 100
                      AS path: 2831679853 I
                    > to 192.168.0.6 via ge-0/0/5.302
172.30.5.1/32      *[LDP/9] 00:21:50, metric 15
                    > to 172.30.0.13 via ge-0/0/4.123, Push 299776
172.30.5.2/32      *[LDP/9] 00:21:50, metric 10
                    > to 172.30.0.13 via ge-0/0/4.123, Push 0
172.30.5.4/32      *[LDP/9] 00:21:49, metric 10
                    > to 172.30.0.22 via ge-0/0/4.134, Push 0
172.30.5.5/32      *[LDP/9] 00:20:33, metric 15
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-
R6
172.30.5.6/32      *[RSVP/7/1] 00:21:54, metric 10
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13
                    [LDP/9] 00:20:33, metric 10
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-
R6
172.30.5.7/32      *[LDP/9] 00:20:33, metric 20
                    > to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R6
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-
R6
172.30.5.8/32      *[RSVP/7/1] 00:22:11, metric 25
                      to 172.30.0.13 via ge-0/0/4.123, label-switched-path R3-to-
R8-1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.26 via ge-0/0/4.136, label-switched-path R3-to-
R8-2
                      to 172.30.0.26 via ge-0/0/4.136, label-switched-path
Bypass->172.30.0.13->172.30.0.1
                    [LDP/9] 00:19:42, metric 25
                    > to 172.30.0.13 via ge-0/0/4.123, Push 300256
172.30.5.41/32      *[IS-IS/15] 00:00:18, metric 20
                    > to 172.30.0.13 via ge-0/0/4.123
192.168.1.0/24      *[LDP/9] 00:21:50, metric 20
                    > to 172.30.0.13 via ge-0/0/4.123, Push 0
```

The routes in `inet.3` are learned from LDP and the RSVP label distribution protocols.

Note: The local loopback IP address is missing in the table. This route is direct and is present in `inet.0` only.

One of the VPLS PEs is R3, you need to advertise its loopback IP address with a labeled BGP route to the neighboring AS.

You can use the rib-group which copies the loopback IP address of the route reflector to inet.3. You also must modify the policy to include the loopback IP address.

```
[edit]
lab@R3# set policy-options policy-statement rr-loopback-to-inet3 term 1 from route-
filter 172.30.5.3/32 exact

[edit]
lab@R3# commit
commit complete
```

```
[edit]
lab@R3# run show route table inet.3 172.30.5.3/32 exact

inet.3: 17 destinations, 23 routes (11 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.3/32       *[Direct/0] 00:00:53
                     > via lo0.0
```

Now you can modify P3-export policy and advertise the /32 IP addresses from `inet.3` to P3-1 neighbor.

```
[edit]
lab@R3# set policy-options policy-statement P3-export term 2 from rib inet.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set policy-options policy-statement P3-export term 2 from route-filter
172.30.5.0/24 prefix-length-range /32-/32

[edit]
lab@R3# set policy-options policy-statement P3-export term 2 then accept

[edit]
lab@R3# commit
commit complete
```

3) Configure an EBGP session on the route reflector with the remote PE router.

The EBGP session with the remote PE will be established with the route reflector. The l2vpn signaled routes will be reflected to the local PE routers.

RR

```
[edit]
lab@route-reflector# run show route 172.17.47.3

inet.0: 669 destinations, 1070 routes (669 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.47.3/32     *[BGP/170] 00:00:22, localpref 100, from 172.30.5.3
                      AS path: 23456 I
                    > to 172.30.0.69 via ge-0/0/1.207

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.47.3/32     *[BGP/170] 00:00:22, localpref 100, from 172.30.5.3
                      AS path: 23456 I
                      Indirect
```

The remote PE loopback IP address is present in the routing table. You must establish a multihop EBGP session with it.

```
[edit]
lab@route-reflector# set protocols bgp group P3-remote-pe multihop
```

The EBGP session should be configured as `multihop`. This will allow the router to send packets with a TTL greater than 1 to establish the BGP TCP session with the remote PE.

```
[edit]
lab@route-reflector# set protocols bgp group P3-remote-pe multihop no-nexthop-change
```


© 2020 Juniper Networks, Inc. All rights reserved.

By default, BGP changes the next-hop attribute on all EBGP sessions. Changing the next-hop in this specific case will redirect all traffic through the route reflector which is not desired. Usually devices used for Route Reflection are capable of carrying a high number of prefixes in the control plane but are limited in terms of forwarding capacity. With the `no-nexthop-change` knob you will keep the original next-hop. The PE routers in local AS 54591 have a labeled BGP route with the remote next-hop from AS 43208.365.

```
[edit]
lab@route-reflector# set protocols bgp group P3-remote-pe local-address 172.30.5.41
```

Specify the source address the BGP TCP session.

```
[edit]
lab@route-reflector# set protocols bgp group P3-remote-pe family l2vpn signaling
```

Configure family l2vpn signaling for BGP VPLS control traffic exchange with the remote AS.

```
[edit]
lab@route-reflector# set protocols bgp group P3-remote-pe peer-as 23456
```

Because the peer AS number is 4-byte, and the router reflector is configured to only supports 2-byte AS numbers you have to configure the peering with AS_TRANS 23456. This reserved AS number can be used only by peers that support 4-byte AS.

```
[edit]
lab@route-reflector# set protocols bgp group P3-remote-pe neighbor 172.17.47.3

[edit]
lab@route-reflector# commit
commit complete
```

Verify the BGP session table for the remote EBGP peer.

```
[edit]
lab@route-reflector# run show bgp summary
Groups: 3 Peers: 9 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0              1002        601          0          0          0          0
bgp.l3vpn.0          105        105          0          0          0          0
inet6.0               64         48          0          0          0          0
bgp.l2vpn.0            9          9          0          0          0          0
bgp.mvpn.0             5          5          0          0          0          0
inet.3                 1          1          0          0          0          0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Peer                 AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
172.17.47.3       23456          3         15       0       0          48 Establ
  bgp.l2vpn.0: 1/1/1/0
...
```

Verify the single prefix received from the remote neighbor in the `bgp.l2vpn.0` table.

```
[edit]
lab@route-reflector# run show route receive-protocol bgp 172.17.47.3 extensive
...
bgp.l2vpn.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
*  172.17.47.3:500:7:1/96 (1 entry, 1 announced)
     Accepted
     Route Distinguisher: 172.17.47.3:500
     Label-base: 262145, range: 8, offset: 1
     Nexthop: 172.17.47.3
     AS path:
     AS2 PA[1]: AS_TRANS
     AS4 PA[1]: 2831679853
     Merged[1]: 2831679853 I
     Communities: target:43208:500 Layer2-info: encaps: VPLS, control flags:[0x0]
, mtu: 0, site preference: 100
...
```

The route target community target:43208:500 attached to the route is different than the route target used for customer C5 VPLS service. You must normalize the route targets attached to the prefixes in both directions.

```
[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-import term 1
from protocol bgp

[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-import term 1
from community CE5-remote

[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-import term 1
then community delete CE5-remote

[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-import term 1
then community add CE5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-import term 1
then accept

[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-export term 1
from protocol bgp

[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-export term 1
from community CE5

[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-export term 1
then community delete CE5

[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-export term 1
then community add CE5-remote

[edit]
lab@route-reflector# set policy-options policy-statement C5-vpn-target-export term 1
then accept

[edit]
lab@route-reflector# set policy-options community CE5-remote members target:43208:500

[edit]
lab@route-reflector# set policy-options community CE5 members target:54591:501

[edit]
lab@route-reflector# set protocols bgp group P3-remote-pe import C5-vpn-target-import

[edit]
lab@route-reflector# set protocols bgp group P3-remote-pe export C5-vpn-target-export

[edit]
lab@route-reflector# commit
commit complete
```

The policies delete the extended route target community and set the correct one between the remote and local AS.

© 2020 Juniper Networks, Inc. All rights reserved.

Verify the operation of the Inter-provider option C VPLS service.

```
[edit]
lab@R2# run show route table bgp.l2vpn.0 extensive
...
172.17.47.3:500:7:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 172.17.47.3:500
                Next hop type: Indirect, Next hop index: 0
                Address: 0xdd378b0
                Next-hop reference count: 5
                Source: 172.30.5.41
                Protocol next hop: 172.17.47.3
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                State: <Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 2:23       Metric2: 1
                Validation State: unverified
                Task: BGP_54591.172.30.5.41
                AS path:
                AS2 PA[1]: AS_TRANS
                AS4 PA[1]: 2831679853
                Merged[1]: 2831679853 I
                Communities: target:54591:501 Layer2-info: encaps: VPLS, control
flags:[0x0] , mtu: 0, site preference: 100
                Import Accepted
                Label-base: 262145, range: 8, offset: 1
                Localpref: 100
                Router ID: 172.30.5.41
                Secondary Tables: C5-vpls.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 172.17.47.3 Metric: 10
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 172.30.0.14 via ge-0/0/4.123
                                Session Id: 0x0
                                172.17.47.3/32 Originating RIB: inet.3
                                  Metric: 10    Node path count: 1
                                  Indirect nexthops: 1
                                        Protocol Nexthop: 172.30.5.3 Metric: 10
Push 305872
                                        Indirect nexthop: 0xc2bb300 - INH Session
ID: 0x0
                                         Path forwarding nexthops link: 0xc7a9f68
                                         Path inh link: 0x0
                                        Indirect path forwarding nexthops: 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                                        Nexthop: 172.30.0.14 via ge-
0/0/4.123
                                        Session Id: 0
                                172.30.5.3/32 Originating RIB: inet.3
                                  Metric: 10    Node path count: 1
                                  Forwarding nexthops: 1
                                        Nexthop: 172.30.0.14 via ge-
0/0/4.123
                                        Session Id: 14e
```

The L2VPN routes are received from the route reflector with the correct route target. The transport label (318064) is resolved in inet.3.

```
[edit]
lab@R2# run show route table inet.3 172.17.47.3

inet.3: 13 destinations, 18 routes (9 active, 0 holddown, 7 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.47.3/32     *[BGP/170] 01:01:20, localpref 100, from 172.30.5.41
                      AS path: 2831679853 I
                    > to 172.30.0.14 via ge-0/0/4.123, Push 318064
```

The remote PE loopback IP address is learned by labeled BGP.

© 2020 Juniper Networks, Inc. All rights reserved.

Verify the VPLS connection table.

```
[edit]
lab@R2# run show vpls connections
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
VM -- VLAN ID mismatch
Legend for interface status
Up -- operational
Dn -- down
Instance: C5-vpls
  Local site: site-4 (4)
    connection-site      Type  St     Time last up          # Up trans
    5                    rmt   Up     Nov  4 20:24:03 2013           1
      Remote PE: 172.30.5.3, Negotiated control-word: No
      Incoming label: 262149, Outgoing label: 262148
      Local interface: lsi.1048663, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 4 remote site 5
    6                    rmt   Up     Nov  4 19:26:51 2013           1
      Remote PE: 172.30.5.5, Negotiated control-word: No
      Incoming label: 262150, Outgoing label: 262148
      Local interface: lsi.1048650, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 4 remote site 6
    7                    rmt   Up     Nov  4 20:28:18 2013           1
      Remote PE: 172.17.47.3, Negotiated control-word: No
      Incoming label: 262151, Outgoing label: 262148
      Local interface: lsi.1048664, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 4 remote site 7
```

The connection to site 7 is established.

© 2020 Juniper Networks, Inc. All rights reserved.

On R3 the labeled BGP route 172.17.47.3 is received from P3-1 with an Implicit Null label value 3. This is because the remote VPLS site is directly attached on P3-1.

```
[edit]
lab@R3# run show route table inet.3 172.17.47.3/32 extensive

inet.3: 17 destinations, 23 routes (11 active, 0 holddown, 10 hidden)
172.17.47.3/32 (1 entry, 1 announced)
TSI:
Page 0 idx 1 Type 1 val 18e8604
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [54591] 2831679853 I
    Communities:
Path 172.17.47.3 from 192.168.0.6 Vector len 4.  Val: 1
        *BGP    Preference: 170/-101
                Next hop type: Router
                Address: 0x18f20b8
                Next-hop reference count: 5
                Source: 192.168.0.6
                Next hop: 192.168.0.6 via ge-0/0/5.302, selected
                State: <Active Ext>
                Local AS: 54591 Peer AS: 2831679853
                Age: 1:12:29
                Task: BGP_2831679853.192.168.0.6+179
                Announcement bits (3): 1-Resolve tree 1 3-BGP_RT_Background 4-
Resolve tree 4
                AS path: 2831679853 I
                Accepted
                Route Label: 3
                Localpref: 100
                Router ID: 172.17.47.1
```

R3 advertises the route with new label to the RR, because the next-hop of the route is changed.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.41 table inet.3
172.17.47.3/32 extensive

inet.3: 17 destinations, 23 routes (11 active, 0 holddown, 10 hidden)
* 172.17.47.3/32 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Label: 318064
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
     AS path: [54591] 2831679853 I
```

This label is used by other the PE routers to reaching the remote PE router.

```
[edit]
lab@R3# run show route table mpls.0 label 318064

mpls.0: 62 destinations, 62 routes (62 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

318064             *[VPN/170] 01:06:43
                    > to 192.168.0.6 via ge-0/0/5.302, Pop
318064(S=0)        *[VPN/170] 01:06:43
                    > to 192.168.0.6 via ge-0/0/5.302, Pop
```

In the MPLS table this label is popped because P3-1 advertises an Implicit Null label for this route.

P3 is also a PE router with a direct attached site of customer C5 VPLS service.

```
[edit]
lab@R3# run show vpls connections
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
VM -- VLAN ID mismatch

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


© 2020 Juniper Networks, Inc. All rights reserved.

```
Instance: C5-vpls
  Local site: site-5 (5)
    connection-site           Type  St     Time last up          # Up trans
    4                         rmt   Up     Nov  4 20:02:03 2013           1
      Remote PE: 172.30.5.2, Negotiated control-word: No
      Incoming label: 262148, Outgoing label: 262149
      Local interface: lsi.1048686, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 5 remote site 4
    5                         rmt   RN
    6                         rmt   Up     Nov  4 19:58:28 2013           4
      Remote PE: 172.30.5.5, Negotiated control-word: No
      Incoming label: 262150, Outgoing label: 262149
      Local interface: lsi.1048651, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 5 remote site 6
    7                         rmt   Up     Nov  4 20:28:18 2013           1
      Remote PE: 172.17.47.3, Negotiated control-word: No
      Incoming label: 262151, Outgoing label: 262149
      Local interface: lsi.1048687, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls C5-vpls local site 5 remote site 7
```

As on R2, the VPLS connection to the remote site is also established on R3.

This concludes Chapter Seven.

© 2020 Juniper Networks, Inc. All rights reserved.

Appendix - Chapter 8: Class of Service

In this chapter, you are going to configure CoS in the network. You will configure classifiers, policers, forwarding classes, queues and schedulers, rewrite markers, and RED drop profiles.

Class of Service overview

In today's networks many different services are delivered over the same network infrastructure. Some of these services have very strict quality of service requirements.

Both IPv4 and IPv6 have a special field in their header to classify packets based on their service type. Such fields are also available in the MPLS header/label and in 802.1q. When traffic is properly marked, the devices in the network can easily provide the required Class of Service. CoS configuration is unidirectional, the settings must be explicitly configured in both directions. CoS behavior is hop-by-hop, each device on the path must be configured with the correct settings providing end-to-end service.

The following main CoS features can be enabled on a Junos device.

- *Policing*- policing and rate-limiting provides an ingress traffic limit at a configured level. This function protects the device and network resources for unusual peaks.
- *Classification* – this is a critical function, which provides traffic recognition and mapping it to correct classes to ensuring proper processing by other CoS tools.
- *Scheduling* – Egress traffic is mapped into a queue. Each queue can be configured for scheduling and shaping depending on certain CoS service requirements. The Random early detection (RED) algorithm is able to control congestion by selectively dropping certain packets based to avoid TCP global synchronization (tail drop).
- *Rewrite Rules* – rewrites the Class of Service bits. Traffic is re-marked with the correct before send to the next hop.

Keep in mind that many CoS features depend on the availability of specific hardware, which is often different between Junos platforms.

Task 1: Forwarding Classes, Queues and Schedulers

The following configuration should be configured on all devices.

1) Configure the forwarding classes.

Traffic is classified in forwarding classes. They are directly mapped to the output queues.

© 2020 Juniper Networks, Inc. All rights reserved.

Four predefined classes are available on each Junos devices.

```
[edit]
lab@R1# run show class-of-service forwarding-class
Forwarding class                          ID      Queue  Restricted queue  Fabric
priority  Policing priority   SPU priority
  best-effort                             0       0          0              low
normal            low
  expedited-forwarding                    1       1          1              low
normal            low
  assured-forwarding                      2       2          2              low
normal            low
  network-control                         3       3          3              low
normal            low
```

You have to change the name of some the classes based on Table 15.

Forwarding Class	Queue	Scheduler
best-effort	0	be-sc-q0
vpn	1	vpn-sc-q1
vpn-priority	2	vpn-pri-sc-q2
nc	3	nc-sc-q3

Table 15

```
[edit]
lab@R1# set class-of-service forwarding-classes queue 0 best-effort

[edit]
lab@R1# set class-of-service forwarding-classes queue 1 vpn

[edit]
lab@R1# set class-of-service forwarding-classes queue 2 vpn-priority

[edit]
lab@R1# set class-of-service forwarding-classes queue 3 nc

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Again, verify the forwarding classes configuration.

```
[edit]
lab@R1# run show class-of-service forwarding-class
Forwarding class                          ID      Queue  Restricted queue  Fabric
priority  Policing priority   SPU priority
  best-effort                             0       0          0              low
normal            low
  vpn                                     1       1          1              low
normal            low
  vpn-priority                            2       2          2              low
normal            low
  nc                                      3       3          3              low
normal            low
```

2) Configure drop profiles outlined in Table 16 and Table 17.

Name: `low-drop`

Fill Level	Drop Probability
25	5
50	15
75	40

Table 16

Name: `high-drop`

Fill Level	Drop Probability
25	5
50	15
75	40

Table 17

Drop profiles provide random early detection (RED) of the traffic. They selectively drop traffic to avoid congestion in the queues. a weight can be given to each type of traffic within each class by Configure different drop profiles.

© 2020 Juniper Networks, Inc. All rights reserved.

A drop profile defines at what buffer level, the probability for traffic to be dropped. Traffic market with high loss priority has a higher chance to be dropped.

A default Drop Profile is pre-defined in Junos and used in case of congestion. It has two fullness levels, which are 0 and 100 and two drop probabilities which are 0 percent and 100 percent. This means that no traffic is dropped at fullness level of 0, and 100 percent is dropped at when the buffer fullness is 100.

```
drop-profiles {
    terminal {
        fill-level 100 drop-probability 100;
    }
}
```

Configure both drop profiles based on the task requirements.

Two types of profiles can be configured in Junos. Segmented and Interpolated. By default, a drop profile is segmented, which means that based on the configured values a step-by-step graph is created. The drop probability is constant until the next configured level of buffer fullness is reached.

The interpolate profile is smoother, because the graph is created automatically based on 64 fill levels. You can configure some of the fill levels and drop probabilities and the software will generate a graduated map.

```
[edit]
lab@R1# set class-of-service drop-profiles low-drop interpolate fill-level [ 25 50 75 ]

[edit]
lab@R1# set class-of-service drop-profiles low-drop interpolate drop-probability [ 5 15
40 ]

[edit]
lab@R1# set class-of-service drop-profiles high-drop interpolate fill-level [ 25 50 75
]

[edit]
lab@R1# set class-of-service drop-profiles high-drop interpolate drop-probability [ 10
30 65 ]

[edit]
lab@R1# commit
commit complete
```

3) Configure the schedulers based on the information in Table 18.

© 2020 Juniper Networks, Inc. All rights reserved.

Scheduler	Parameter	Value
be-sc-q0	Priority	low
	Transmit rate	remainder
	Buffer size	remainder
	Drop profile LP any	high-drop
vpn-sc-q1	Priority	medium-low
	Transmit rate	20%
	Buffer size	20%
	Drop profile LP low	low-drop
	Drop profile LP high	high-drop
vpn-pri-sc-q2	Priority	medium-high
	Transmit rate	10%
	Buffer size	5 msec
nc-sc-q3	Priority	high
	Transmit rate	5%
	Buffer size	5%

Table 18

Schedulers define the way how queues will treat traffic. They determine the transmission order of the packets, the rate at which the traffic is transmitted, the number of packets that can buffered and in the event of congestion, the treatment for each traffic class.

Schedulers assign relative priorities to the queues. Five priorities are available on Junos devices.

- Low
- Medium-low
- Medium-high
- High
- Strict-high

Strict-high is a special queue which is always in profile. It is used for low latency traffic. Traffic for this queue must be rate-limited, as it is always preferred over traffic in other queues. This might lead to traffic starvation for the other non strict-high queues.

Two default schedulers are configured in Junos.

© 2020 Juniper Networks, Inc. All rights reserved.

```
schedulers {
   network-control {
      transmit-rate percent 5;
      buffer-size percent 5;
      priority low;
      drop-profile-map loss-priority any protocol any drop-profile terminal;
      }
   best-effort {
      transmit-rate percent 95;
      buffer-size percent 95;
      priority low;
      drop-profile-map loss-priority any protocol any drop-profile terminal;
      }
   }
```

Configure four schedulers, which will be assigned a scheduler-map.

```
[edit]
lab@R1# set class-of-service schedulers be-sc-q0 transmit-rate remainder
```

This scheduler will be assigned to the best-effort forwarding-class. It can use all the left bandwidth when it is in profile.

```
[edit]
lab@R1# set class-of-service schedulers be-sc-q0 buffer-size remainder
```

It can use the remaining buffer of the queue.

```
[edit]
lab@R1# set class-of-service schedulers be-sc-q0 priority low
```

Low priority is assigned to that queue. This queue will have the lowest priority in case of congestion.

```
[edit]
lab@R1# set class-of-service schedulers be-sc-q0 drop-profile-map loss-priority any
protocol any drop-profile high-drop
```

An aggressive drop-profile is assigned to all of the traffic.

```
[edit]
lab@R1# set class-of-service schedulers vpn-sc-q1 transmit-rate percent 20
```

The scheduler will be assigned to the VPN forwarding class. It will be able to use 20 percent of the interface bandwidth. Transmit-rate is an interface shaper. Traffic is limited to 20 percent of the bandwidth during congestion. When the interface is not congested, this

© 2020 Juniper Networks, Inc. All rights reserved.

class can reach the maximum bandwidth, so it is not hard limited. Using the `rate-limit` option under `transmit-rate` limits the forwarding-class to the configured value.

```
[edit]
lab@R1# set class-of-service schedulers vpn-sc-q1 buffer-size percent 20
```

20 percent of the buffer is dedicated for that queue.

```
[edit]
lab@R1# set class-of-service schedulers vpn-sc-q1 priority medium-low
```

The priority is set to medium-low. This queue will have a relatively high priority compared to the best-effort forwarding-class when both classes are out of profile.

```
[edit]
lab@R1# set class-of-service schedulers vpn-sc-q1 drop-profile-map loss-priority low
protocol any drop-profile low-drop
```

Two different drop profiles are used based on the packet loss priority (PLP). PLP marks traffic and can be used to control the drop probability of that traffic. It is set by the classifiers. Four PLP values can be configured: low, medium-low, medium-high and high. By default, only the low and high value are used.

```
[edit]
lab@R1# set class-of-service schedulers vpn-sc-q1 drop-profile-map loss-priority high
protocol any drop-profile high-drop
```

```
[edit]
lab@R1# set class-of-service schedulers vpn-prio-sc-q2 transmit-rate percent 10
```

```
[edit]
lab@R1# set class-of-service schedulers vpn-prio-sc-q2 buffer-size temporal 5k
```

The `buffer-size` for the `vpn-priority` class is configured with the `temporal` option. This is used to reduce traffic latency for a given forwarding class. 5K is 5000 microseconds. With the `temporal` parameter, the buffer has a strict upper limit and cannot expand.

```
[edit]
lab@R1# set class-of-service schedulers vpn-prio-sc-q2 priority medium-high
```

```
[edit]
lab@R1# set class-of-service schedulers nc-sc-q3 transmit-rate percent 5
```

```
[edit]
lab@R1# set class-of-service schedulers nc-sc-q3 buffer-size percent 5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set class-of-service schedulers nc-sc-q3 priority high

[edit]
lab@R1# commit
commit complete
```

4) Configure the scheduler maps.

The schedulers configured in the previous step have to be assigned to forwarding classes. This is done with schedulers maps, which are then applied to the interfaces.

```
[edit]
lab@R1# set class-of-service scheduler-maps core-interfaces forwarding-class best-
effort scheduler be-sc-q0

[edit]
lab@R1# set class-of-service scheduler-maps core-interfaces forwarding-class nc
scheduler nc-sc-q3

[edit]
lab@R1# set class-of-service scheduler-maps core-interfaces forwarding-class vpn
scheduler vpn-sc-q1

[edit]
lab@R1# set class-of-service scheduler-maps core-interfaces forwarding-class vpn-
priority scheduler vpn-prio-sc-q2

[edit]
lab@R1# commit
commit complete
```

Verify the configured scheduler-map.

```
[edit]
lab@R1# run show class-of-service scheduler-map core-interfaces | no-more

Scheduler map: core-interfaces, Index: 58651

  Scheduler: be-sc-q0, Forwarding class: best-effort, Index: 9240
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         48162    high-drop
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Medium low      any          48162    high-drop
    Medium high     any          48162    high-drop
    High            any          48162    high-drop

Scheduler: vpn-sc-q1, Forwarding class: vpn, Index: 37515
  Transmit rate: 20 percent, Rate Limit: none, Buffer size: 20 percent,
  Buffer Limit: none, Priority: medium-low
  Excess Priority: unspecified
  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             any         59912    low-drop
    Medium low      any             1    <default-drop-profile>
    Medium high     any             1    <default-drop-profile>
    High            any         48162    high-drop

Scheduler: vpn-prio-sc-q2, Forwarding class: vpn-priority, Index: 57617
  Transmit rate: 10 percent, Rate Limit: none, Buffer size: 5000 us,
  Buffer Limit: none, Priority: medium-high
  Excess Priority: unspecified
  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             any             1    <default-drop-profile>
    Medium low      any             1    <default-drop-profile>
    Medium high     any             1    <default-drop-profile>
    High            any             1    <default-drop-profile>

Scheduler: nc-sc-q3, Forwarding class: nc, Index: 42106
  Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
  Buffer Limit: none, Priority: high
  Excess Priority: unspecified
  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             any             1    <default-drop-profile>
    Medium low      any             1    <default-drop-profile>
    Medium high     any             1    <default-drop-profile>
    High            any             1    <default-drop-profile>
```

5) Apply the scheduler-map to the interfaces.

```
[edit]
lab@R1# set class-of-service interfaces ge-0/0/4 unit * scheduler-map core-interfaces

[edit]
lab@R1# set class-of-service interfaces ae0 unit * scheduler-map core-interfaces
```

Apply the scheduler-map to the core interfaces. You can use regular-expressions to specify all units or a group of interface units.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
[edit class-of-service interfaces]
  'ge-0/0/4'
    'per-unit-scheduler', 'hierarchical-scheduler', or 'shared-scheduler' for this
interface is required for scheduler map on interface unit
[edit class-of-service interfaces]
  'ae0'
    'per-unit-scheduler', 'hierarchical-scheduler', or 'shared-scheduler' for this
interface is required for scheduler map on interface unit
error: commit failed: (statements constraint check failed)
```

The commit fails. By default, physical interfaces support only one scheduler map. If you want to ensure that each logical interface can have a different scheduler applied, you must use the `per-unit-scheduler` feature.

```
[edit]
lab@R1# set interfaces ge-0/0/4 per-unit-scheduler

[edit]
lab@R1# set interfaces ae0 per-unit-scheduler

[edit]
lab@R1# commit
commit complete
```

After enabling `per-unit-scheduler` commit is successful.

Verify the interface queues configured in the scheduler-map.

```
[edit]
lab@R1# run show interfaces queue ae0.0
  Logical interface ae0.0 (Index 102) (SNMP ifIndex 571)
    Description: R2 connection
    Flags: SNMP-Traps 0x0 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          2155         41        149881        17352
        Output:          1629          3        179473         1648
Forwarding classes: 8 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              :              106                    0 pps
    Bytes                :            15160                  592 bps
  Transmitted:
    Packets              :              106                    0 pps
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Bytes                  :            15160                592 bps
    Tail-dropped packets   :                0                  0 pps
    RED-dropped packets    :                0                  0 pps
    RED-dropped bytes      :                0                  0 bps
Queue: 1, Forwarding classes: vpn
  Queued:
    Packets                :                0                  0 pps
    Bytes                  :                0                  0 bps
  Transmitted:
    Packets                :                0                  0 pps
    Bytes                  :                0                  0 bps
    Tail-dropped packets   :                0                  0 pps
    RED-dropped packets    :                0                  0 pps
    RED-dropped bytes      :                0                  0 bps
Queue: 2, Forwarding classes: vpn-priority
  Queued:
    Packets                :                0                  0 pps
    Bytes                  :                0                  0 bps
  Transmitted:
    Packets                :                0                  0 pps
    Bytes                  :                0                  0 bps
    Tail-dropped packets   :                0                  0 pps
    RED-dropped packets    :                0                  0 pps
    RED-dropped bytes      :                0                  0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets                :             1720                 32 pps
    Bytes                  :           174177              19904 bps
  Transmitted:
    Packets                :             1720                 32 pps
    Bytes                  :           174177              19904 bps
    Tail-dropped packets   :                0                  0 pps
    RED-dropped packets    :                0                  0 pps
    RED-dropped bytes      :                0                  0 bps
```

Task 2: Classification, Policing and Marking

1) Configure a firewall classifier on R3 and R8 to classify traffic from customer 3 sites.

Multi-field classifiers provide granular classification of the incoming traffic. Firewall filters are used to match the desired traffic and assigned it to the given forwarding-class.

Note: You will need to change the interface that the filter is applied to for R8 (ge-0/0/5.325).

Configure on both R3 and R8.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set firewall family inet filter C3-classifier term 1 from dscp be

[edit]
lab@R3# set firewall family inet filter C3-classifier term 1 then forwarding-class vpn

[edit]
lab@R3# set firewall family inet filter C3-classifier term 1 then accept
```

Best-effort traffic is assigned to vpn forwarding-class. A dscp value is configured with the name be. This is called a code-point-alias in Junos.

```
[edit]
lab@R3# set firewall family inet filter C3-classifier term 2 then forwarding-class vpn-
priority

[edit]
lab@R3# set firewall family inet filter C3-classifier term 2 then accept
```

The rest of the traffic is assigned to the vpn-priority class.

Assign the filter to the C3 site facing interface.

```
[edit]
lab@R3# set interface ge-0/0/5.317 family inet filter input C3-classifier
```

2) Configure the core point aliases on all routers.

Code-point-aliases assign friendly names to class-of-service code-points.

```
[edit]
lab@R3# set class-of-service code-point-aliases dscp vpn-low 001010
```

Here dscp is the type of the alias, `vpn-low` is the alias name and 001010 is the code-point.

```
[edit]
lab@R3# set class-of-service code-point-aliases dscp vpn-high 001100

[edit]
lab@R3# set class-of-service code-point-aliases dscp vpn-priority 101110

[edit]
lab@R3# set class-of-service code-point-aliases dscp be 000000

[edit]
lab@R3# set class-of-service code-point-aliases dscp nc 110000
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set class-of-service code-point-aliases exp vpn-low 010

[edit]
lab@R3# set class-of-service code-point-aliases exp vpn-high 011

[edit]
lab@R3# set class-of-service code-point-aliases exp vpn-priority 101

[edit]
lab@R3# set class-of-service code-point-aliases exp be 000

[edit]
lab@R3# commit
commit complete
```

3) Configure a next-hop map on R3 and R8.

CoS based forwarding provides the ability to select a forwarding path based on the given forwarding-class.

First you have to configure a CoS forwarding policy, where you specify the next-hop to be mapped to a given forwarding-class.

Note: Ensure you change the LSP name when applying to R8.

```
[edit]
lab@R3# set class-of-service forwarding-policy next-hop-map cbf-map forwarding-class
vpn lsp-next-hop R3-to-R8-2

[edit]
lab@R3# set class-of-service forwarding-policy next-hop-map cbf-map forwarding-class
vpn-priority lsp-next-hop R3-to-R8-1
```

4) Configure the next-hop mapping policy on R3 and R8.

The interesting traffic that should use the CoS based forwarding feature is identified with a routing policy. The terminating action of this policy is `cos-next-hop-map`.

Note: Interesting traffic for R8 should be defined as “fd18:cccc:dddd:5:8::/77”.

```
[edit]
lab@R3# set policy-options policy-statement cbf-map term 1 from route-filter
fd18:cccc:dddd:5:0::/77 longer
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set policy-options policy-statement cbf-map term 1 then cos-next-hop-map cbf-
map
```

5) Apply the routing policy to the forwarding table.

The last step is to apply the routing policy to the forwarding-table, the same way as applying a load balancing policy.

```
[edit]
lab@R3# set routing-options forwarding-table export cbf-map
```

```
[edit]
lab@R3# commit
commit complete
```

Verify the next-hop installed in the forwarding table.

```
[edit]
lab@R3# run show route forwarding-table destination fd18:cccc:dddd:5:0::/77
Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    44     1

Routing table: __master.anon__.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   568     1

Routing table: C1.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   576     1

Routing table: C3.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   584     1
fd18:cccc:dddd:5::/80 user     0                 indr 262154     9
                                                 idxd   798     2
                   idx:2      172.30.0.13       Push 349008, Push 312016(top)
615     2 ge-0/0/4.123
                   idx:2      172.30.0.13       Push 349008, Push 311888(top)
720     2 ge-0/0/4.123
                   idx:2      172.30.0.22       Push 349008, Push 310160, Push
310976(top)   779     2 ge-0/0/4.134
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    idx:2       172.30.0.22         Push 349008, Push 309712, Push
310976(top)   780      2 ge-0/0/4.134
                    idx:xx      172.30.0.26         Push 349008, Push 324720(top)
620      2 ge-0/0/4.136
```

6) Configure the policers on R3 and R8.

The policers were configured in the first Chapter to limit ICMP traffic to the Routing Engine. Here again policers are used to limit or change the loss priority of the traffic.

Two policers have to be created, one soft policer which will increase the packet loss priority of the exceeding traffic. The other policer will be a hard policer, which drops traffic exceeding the limit.

```
[edit]
lab@R3# set firewall policer vpn-policer if-exceeding bandwidth-limit 60m
```

With `bandwidth-limit` you set the committed information rate (CIR). In this particular case it is set to 60 Mbps.

```
[edit]
lab@R3# set firewall policer vpn-policer if-exceeding burst-size-limit 62k
```

The **`burst-size`** is setting the size, in bytes, that can pass through the policer during a burst. It should be higher than the MTU of the interfaces. A minimal recommended size is 10 times the MTU. For an optimal **`burst-size`** you can use the following formula:

burst-size = interface bandwidth x allowable time for bust traffic / 8

These two policers will be applied to the LSPs. The bandwidth of the LSP is 60m. Using the formula above you can calculate the burst for 10ms. The result is 62Kb.

```
[edit]
lab@R3# set firewall policer vpn-policer then loss-priority high
```

The soft policers change the loss priority of the exceeding traffic.

```
[edit]
lab@R3# set firewall policer vpn-priority-policer if-exceeding bandwidth-limit 60m

[edit]
lab@R3# set firewall policer vpn-priority-policer if-exceeding burst-size-limit 62k
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set firewall policer vpn-priority-policer then discard
```

The hard policers drop the exceeding traffic.

```
[edit]
lab@R3# set firewall family any filter vpn-filter term 1 then policer vpn-policer

[edit]
lab@R3# set firewall family any filter vpn-filter term 1 then accept

[edit]
lab@R3# set firewall family any filter vpn-priority-filter term 1 then policer vpn-
priority-policer

[edit]
lab@R3# set firewall family any filter vpn-priority-filter term 1 then accept
```

Configure a firewall filters for protocol-independent traffic with a policer as the terminating action. This is needed as the policers will be applied to LSPs, traffic entering the LSP can be from different protocol families: IPv4, IPv6, MPLS

7) Apply the policing filters to the LSPs.

Note: You must ensure you specify the correct LSPs on R8.

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-1 policing filter vpn-priority-
filter

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8-2 policing filter vpn-filter

[edit]
lab@R3# commit
commit complete
```

8) Configure the behavior aggregate classifiers on all routers using Table 19.

Forwarding Class	Loss Priority	DSCP Value	EXP Value
best-effort	any	0b000000	0b000
vpn	low	0b001010	0b010
	high	0b001100	0b011
vpn-priority	any	0b101110	0b101

© 2020 Juniper Networks, Inc. All rights reserved.

nc	any	0b110000	

*Table 19*38

Unlike Multifield classifiers, where a filter is configured for granular matching, BA classifiers are a simple way of classifying the traffic received on interfaces. They classify the traffic based on the existing (CoS) markings.

You must configure the DSCP and EXP BA classifiers on all routers.

Repeat this on all eight devices.

Define the code point names.

```
[edit]
lab@R1# set class-of-service code-point-aliases dscp vpn-low 001010

[edit]
lab@R1# set class-of-service code-point-aliases dscp vpn-high 001100

[edit]
lab@R1# set class-of-service code-point-aliases dscp vpn-priority 101110

[edit]
lab@R1# set class-of-service code-point-aliases dscp be 000000

[edit]
lab@R1# set class-of-service code-point-aliases dscp nc 110000

[edit]
lab@R1# set class-of-service code-point-aliases exp vpn-low 010

[edit]
lab@R1# set class-of-service code-point-aliases exp vpn-high 011

[edit]
lab@R1# set class-of-service code-point-aliases exp vpn-priority 101

[edit]
lab@R1# set class-of-service code-point-aliases exp be 000
```

Classify `dscp-classifier` from type `dscp` classifies code-points `be` to forwarding-class `best-effort`. Create the other classifiers as defined in Table 19.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set class-of-service classifiers dscp dscp-classifier forwarding-class best-
effort loss-priority low code-points be

[edit]
lab@R1# set class-of-service classifiers dscp dscp-classifier forwarding-class vpn
loss-priority low code-points vpn-low

[edit]
lab@R1# set class-of-service classifiers dscp dscp-classifier forwarding-class vpn
loss-priority high code-points vpn-high

[edit]
lab@R1# set class-of-service classifiers dscp dscp-classifier forwarding-class vpn-
priority loss-priority low code-points vpn-priority

[edit]
lab@R1# set class-of-service classifiers dscp dscp-classifier forwarding-class nc loss-
priority low code-points nc

[edit]
lab@R1# set class-of-service classifiers exp mpls-classifier forwarding-class best-
effort loss-priority low code-points be

[edit]
lab@R1# set class-of-service classifiers exp mpls-classifier forwarding-class vpn loss-
priority low code-points vpn-low

[edit]
lab@R1# set class-of-service classifiers exp mpls-classifier forwarding-class vpn loss-
priority high code-points vpn-high

[edit]
lab@R1# set class-of-service classifiers exp mpls-classifier forwarding-class vpn-
priority loss-priority low code-points vpn-priority
```

9) Apply the classifiers on the interface units.

BA classifiers should be applied to the logical units of the interfaces to classify the traffic received on that interfaces.

```
[edit]
lab@R1# set class-of-service interfaces ge-0/0/4 unit * classifiers dscp dscp-
classifier

[edit]
lab@R1# set class-of-service interfaces ge-0/0/4 unit * classifiers exp mpls-classifier
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set class-of-service interfaces ae0 unit * classifiers dscp dscp-classifier

[edit]
lab@R1# set class-of-service interfaces ae0 unit * classifiers exp mpls-classifier

[edit]
lab@R1# commit
commit complete
```

Both types of classifiers, DSCP and EXP, are applied to all units of the core facing interfaces. EXP classifies packets with an mpls header, DSCP classifies IPv4 traffic.

In RFC5462 the MPLS EXP field has been renamed to TC (Traffic Class). The main reason is that the EXP field is not "experimental" anymore and is commonly used to classify MPLS traffic. For this reason, it has been renamed to Traffic Class. However, in Junos the TC bits must still be classified with EXP classifiers.

Note: As with schedulers-maps, when applying the same classifier to many units you can use a regular expressions value like '*'.

10) Configure rewrite rules on all routers.

Rewrite rules generally perform the opposite task of BA classifiers. They assign CoS code-points based on forwarding classes and the PLP of the packets. They provide consistent CoS throughout the network. Based on the PLP they can change the CoS bits of the traffic that exceeds the configured CIR of a policer or shaper.

Repeat this on all eight devices.

```
[edit]
lab@R1# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class best-
effort loss-priority low code-point be

[edit]
lab@R1# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class vpn
loss-priority low code-point vpn-low

[edit]
lab@R1# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class vpn
loss-priority high code-point vpn-high

[edit]
lab@R1# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class vpn-
priority loss-priority low code-point vpn-priority
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class nc loss-
priority low code-point nc

[edit]
lab@R1# set class-of-service rewrite-rules exp mpls-rewriter forwarding-class best-
effort loss-priority low code-point be

[edit]
lab@R1# set class-of-service rewrite-rules exp mpls-rewriter forwarding-class vpn loss-
priority low code-point vpn-low

[edit]
lab@R1# set class-of-service rewrite-rules exp mpls-rewriter forwarding-class vpn loss-
priority high code-point vpn-high

[edit]
lab@R1# set class-of-service rewrite-rules exp mpls-rewriter forwarding-class vpn-
priority loss-priority low code-point vpn-priority
```

11) Apply the rewrite rules to the core interface on all routers.

```
[edit]
lab@R1# set class-of-service interfaces ge-0/0/4 unit * rewrite-rules dscp dscp-
rewriter

[edit]
lab@R1# set class-of-service interfaces ge-0/0/4 unit * rewrite-rules exp mpls-rewriter
protocol mpls-inet-both

[edit]
lab@R1# set class-of-service interfaces ae0 unit * rewrite-rules dscp dscp-rewriter

[edit]
lab@R1# set class-of-service interfaces ae0 unit * rewrite-rules exp mpls-rewriter
protocol mpls-inet-both

[edit]
lab@R1# commit
commit complete
```

Verify the interface class-of-services configuration.

```
[edit]
lab@R1# run show class-of-service interface ge-0/0/4
Physical interface: ge-0/0/4, Index: 138
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Congestion-notification: Disabled

Logical interface: ge-0/0/4.114, Index: 82
  Object             Name                    Type                    Index
  Rewrite            dscp-rewriter           dscp                    20901
  Rewrite            exp-default             exp (mpls-any)             33
  Rewrite            mpls-rewriter           exp (mpls-inet-both)    10617
  Classifier         dscp-classifier         dscp                    51090
  Classifier         dscp-ipv6-compatibility dscp-ipv6                   9
  Classifier         mpls-classifier         exp                     48975

Logical interface: ge-0/0/4.118, Index: 83
  Object             Name                    Type                    Index
  Rewrite            dscp-rewriter           dscp                    20901
  Rewrite            exp-default             exp (mpls-any)             33
  Rewrite            mpls-rewriter           exp (mpls-inet-both)    10617
  Classifier         dscp-classifier         dscp                    51090
  Classifier         dscp-ipv6-compatibility dscp-ipv6                   9
  Classifier         mpls-classifier         exp                     48975

Logical interface: ge-0/0/4.206, Index: 84
  Object             Name                    Type                    Index
  Rewrite            dscp-rewriter           dscp                    20901
  Rewrite            exp-default             exp (mpls-any)             33
  Rewrite            mpls-rewriter           exp (mpls-inet-both)    10617
  Classifier         dscp-classifier         dscp                    51090
  Classifier         mpls-classifier         exp                     48975

Logical interface: ge-0/0/4.32767, Index: 87
```

This concludes Chapter Eight!

© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Super Lab 1

This lab is designed to combine many of the solutions discussed in the previous chapters into a complete practice lab exam. This lab is similar to a real JNCIE-SP lab exam experience. We recommend attempting this lab if you are able to complete the tasks in the previous chapters without any help.

Task 1: Initial System Configuration

Note: The initial system configuration is common for all routers, except the hostnames and the OoB management addresses. This information is provided for each device in Table 20 and 21.

Note: The devices have been preconfigured with hostnames and management addresses. You will need to confirm the current configuration to ensure there are no intentional mistakes. There is also a lab user that is predefined to allow you access to the environment. You can either reconfigure this user or skip this step.

Router	Router Type	Host Name
vmx1	MX	R1
vmx2	MX	R2
vmx3	MX	R3
vmx4	MX	R4
vmx5	MX	R5
vmx6	MX	R6
vmx7	MX	R7
vmx8	MX	R8

Table 20

Device	OoB Interface Name	OoB Interface IP Address
R1	fxp0	10.10.1.1/24
R2	fxp0	10.10.1.2/24
R3	fxp0	10.10.1.3/24
R4	fxp0	10.10.1.4/24
R5	fxp0	10.10.1.5/24
R6	fxp0	10.10.1.6/24
R7	fxp0	10.10.1.7/24
R8	fxp0	10.10.1.8/24

Table 21

© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure the hostnames on the routers according to Table 20.

```
[edit]
lab@R1# set system host-name R1
```

2) Configure the OoB management interface according to Table 21.

```
[edit]
lab@R1# set interfaces fxp0.0 description "OoB Management"

[edit]
lab@R1# set interfaces fxp0.0 family inet address 10.10.1.1/24
```

3) Ensure that each device accepts management connection for the SSH, Telnet and FTP services.

```
[edit]
lab@R1# set system services ftp

[edit]
lab@R1# set system services ssh

[edit]
lab@R1# set system services telnet
```

4) Configure a static route to the remote management network 10.10.10/24 with a next-hop of 10.10.1.254 on all devices. Make sure that this network is never redistributed to any dynamic routing protocol. Ensure that the devices are reachable, while RPD is not running.

```
[edit]
lab@R1# set routing-options static route 10.10.10.0/24 next-hop 10.10.1.254 no-
readvertise
```

Configure a backup-router ensures connectivity to the management subnet when RDP is not running.

```
[edit]
lab@R1# set system backup-router 10.10.1.254 destination 10.10.10.0/24
```

5) Configure server S1 as the DNS server.

```
[edit]
lab@R1# set system name-server 10.10.1.100
```

6) Set the time zone to America/Los_Angeles on all your devices.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set system time-zone America/Los_Angeles
```

7) Ensure that all your devices synchronize their time with the NTP server S1. Configure the devices to synchronize time with the S1 at boot time.

```
[edit]
lab@R1# set system ntp server 10.10.1.100

[edit]
lab@R1# set system ntp boot-server 10.10.1.100
```

8) Configure the authentication method that first tries to authenticate users on the RADIUS server and then if not successful with the local password database. Use S1 as the RADIUS server. Configure the RADIUS server with a retry attempts 1 and a timeout of 2 seconds. Use workbook as the RADIUS shared secret.

```
[edit]
lab@R1# set system authentication-order [ radius password ]

[edit]
lab@R1# set system radius-server 10.10.1.100 secret workbook

[edit]
lab@R1# set system radius-server 10.10.1.100 timeout 2 retry 1
```

9) Configure a new user lab, with the password lab123, that will have super user privileges.

```
[edit]
lab@R1# set system login user lab class super-user authentication plain-text-password
```

10) Configure additional users on all devices according to Table 22.

Username	Password	Privileges
noc	noc123	Class “operator” permissions. Additionally, this user is allowed to read and modify SNMP configuration, execute system maintenance commands but not allowed to execute “start shell” command
tac	tac123	Class “super-user” permissions. Additionally, this user cannot execute the “clear”, “configure” or “edit” commands

Table 22

© 2020 Juniper Networks, Inc. All rights reserved.

The predefined *operator* class has five permissions flags set. These are *clear, network, reset, trace, and view*. You are not allowed to change these predefined classes. You have to create a new class with those five permission flags plus the two requested in the task. You have to deny the users assigned to that class to not be able to execute *"start shell"*.

```
[edit]
lab@R1# set system login class op-plus permissions [ clear maintenance network reset
snmp-control trace view ]

[edit]
lab@R1# set system login class op-plus deny-commands "start shell"
```

You must configure a new class based on the *super-user* class. However, the users of this class are not allowed to execute *"clear"*, *"configure"*, and *"edit"* command.

```
[edit]
lab@R1# set system login class su-minus permissions all deny-commands
"(clear)|(configure)|(edit)"

[edit]
lab@R1# set system login user noc class op-plus authentication plain-text-password

[edit]
lab@R1# set system login user tac class su-minus authentication plain-text-password
```

11) Configure syslog settings on all your devices as described in Table 23.

Receiver	Message Type
File "jncie-sp-messages"	All info level messages
File "firewall.log"	All firewall filter messages
Syslog server S1	Configuration changes
User "noc"	All warning level messages

Table 23

```
[edit]
lab@R1# set system syslog user noc any warning

[edit]
lab@R1# set system syslog host 10.10.1.100 change-log any

[edit]
lab@R1# set system syslog file jncie-sp-messages any info

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set system syslog file firewall.log firewall any
```

12) Configure SNMPv2 for read only access using the community workbook. Make sure that SNMP server S1 is the only server allowed to access the device with this community.

Unlike SNMPv3, SNMPv2 supports limited security features based on community. You can configure different community for read or write access.

```
[edit]
lab@R1# set snmp community workbook authorization read-only

[edit]
lab@R1# set snmp community workbook clients 10.10.1.100/32
```

13) Configure SNMPv2 to send traps to the SNMP server S1 for routing, link and chassis events.

```
[edit]
lab@R1# set snmp trap-group s1 categories chassis link routing

[edit]
lab@R1# set snmp trap-group s1 targets 10.10.1.100
```

14) Apply a firewall filter to protect the Routing Engine.

```
[edit]
lab@R1# set firewall family inet filter protect-re term 1 from source-address
10.10.1.0/24

[edit]
lab@R1# set firewall family inet filter protect-re term 1 from source-address
10.10.10.0/24

[edit]
lab@R1# set firewall family inet filter protect-re term 1 from source-address
172.30.0.0/16

[edit]
lab@R1# set firewall family inet filter protect-re term 1 from source-address
172.17.0.0/16

[edit]
lab@R1# set firewall family inet filter protect-re term 1 from source-address
172.31.0.0/16
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set firewall family inet filter protect-re term 1 from source-address
192.168.0.0/16

[edit]
lab@R1# set firewall family inet filter protect-re term 1 then accept

[edit]
lab@R1# set firewall family inet filter protect-re term 2 then count dropped-packets

[edit]
lab@R1# set firewall family inet filter protect-re term 2 then syslog

[edit]
lab@R1# set firewall family inet filter protect-re term 2 then discard
```

15) Apply the filter to the loopback interface.

```
[edit]
lab@R1# set interfaces lo0.0 family inet filter input protect-re
```

16) Ensure that all devices archive their configuration periodically every 24 hours to the FTP server S1 using the username lab and password lab123.

```
[edit]
lab@R1# set system archival configuration transfer-interval 1440

[edit]
lab@R1# set system archival configuration archive-sites ftp://lab:lab123@10.10.1.100
```

17) Download the op script called “show-interfaces.slax”, commit script called “interface-mask-check.slax” and event script called “ospf_adjacency_flapping.slax from the FTP server S1 to all your routers.

```
lab@R1> file copy ftp://lab:lab123@10.10.1.100/show-interfaces.slax /var/db/scripts/op
/var/db/scripts/...transferring.file.........A81100% of 2787  B 2981 kBps

lab@R1> file copy ftp://lab:lab123@10.10.1.100/interface-mask-check.slax
/var/db/scripts/commit/
/var/db/scripts/...transferring.file.........bHY100% of  980  B  973 kBps

lab@R1> file copy ftp://lab:lab123@10.10.1.100/ospf_adjacency_flapping.slax
/var/db/scripts/event/
/var/db/scripts/...transferring.file.........3PD100% of  716  B  864 kBps
```

18) Enabling the scripts.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set system scripts op file show-interfaces.slax

[edit]
lab@R1# set system scripts commit file interface-mask-check.slax
```

List the content of the event script to check what triggers the script. This is needed to create an event-options policy.

```
[edit]
lab@R1# run file show /var/db/scripts/event/ospf_adjacency_flapping.slax
/*
 * Very simple script which gets triggerd on receiving
 * RPD_OSPF_NBRDOWN event. RPD_OSPF_NBRDOWN is received when OSPF neighbor
 * adjacency was terminated. This script raises an syslog message
 * with high priority value.
 *
 * Due to a known issue (PR 76849) this script will not get event details.
 * So, we can only log a simple log message saying OSPF adjacency happened
 * but not the details.
 */
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

import "../import/junos.xsl";
--- (more)---
```

Create the event policy and apply the event script.

```
[edit]
lab@R1# set event-options policy ospf_adjacency_flapping events rpd_ospf_nbrdown

[edit]
lab@R1# set event-options policy ospf_adjacency_flapping then event-script
ospf_adjacency_flapping.slax

[edit]
lab@R1# set event-options event-script file ospf_adjacency_flapping.slax

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 2: Building the Network

The second task requires the configuration of the interfaces on all routers according Table 24 and 25. As a result the topology on Figure 16 will be build.

Router	Aggregated Ethernet	Interfaces
R1	ae0	ge-0/0/1
		ge-0/0/2
R2	ae0	ge-0/0/1
		ge-0/0/2
R3	ae0	ge-0/0/1
		ge-0/0/2
R4	ae0	ge-0/0/1
		ge-0/0/2
R5	ae0	ge-0/0/1
		ge-0/0/2
R6	ae0	ge-0/0/1
		ge-0/0/2

Table 24

Note: The interface unit numbers match the VLAN tags.

Router	Interface Name	IP Address	IPv6 Address
R1	ae0.0	172.30.0.1/30	link-local
	ge-0/0/4.117	172.30.0.5/30	link-local
	ge-0/0/4.118	172.30.0.9/30	link-local
	ge-0/0/4.206	172.30.0.65/30	
	ge-0/0/5.318	192.168.0.69/30	
	ge-0/0/5.310	192.168.0.37/30	fc09:c0:ffee::5/126
	lo0.0	172.30.5.1/32	fd17:f0f4:f691:5::1/128
R2	ae0.0	172.30.0.2/30	link-local
	ge-0/0/4.126	172.30.0.17/30	link-local
	ge-0/0/4.123	172.30.0.13/30	link-local
	ge-0/0/4.207	172.30.0.69/30	
	ge-0/0/5.303	192.168.0.9/30	IPv4 compatible/126
	ge-0/0/3.601		
	lo0.0	172.30.5.2/32	fd17:f0f4:f691:5::2/128
R3	ge-0/0/4.123	172.30.0.14/30	link-local
	ge-0/0/4.138	172.30.0.33/30	link-local
	ge-0/0/4.137	172.30.0.29/30	link-local

© 2020 Juniper Networks, Inc. All rights reserved.

	ge-0/0/4.135	172.30.0.85/30	link-local
	ae0.0	172.30.0.81/30	link-local
	ge-0/0/5.306	192.168.0.21/30	
	ge-0/0/5.307	192.168.0.25/30	
	ge-0/0/3.600		
	lo0.0	172.30.5.3/32	fd17:f0f4:f691:5::3/128
R4	ge-0/0/4.146	172.30.0.89/30	link-local
	ae0.0	172.30.0.82/30	link-local
	ge-0/0/3.600		
	ge-0/0/5.323	192.168.0.89/30	
	lo0.0	172.30.5.4/32	fd17:f0f4:f691:5::4/128
	lo0.1	172.30.5.21/32	
R5	ge-0/0/4.135	172.30.0.86/30	link-local
	ae0.0	172.30.0.93/30	link-local
	ge-0/0/5.305	192.168.0.17/30	
	ge-0/0/4.202	172.30.0.49/30	
	lo0.0	172.30.5.5/32	fd17:f0f4:f691:5::5/128
R6	ge-0/0/4.126	172.30.0.18/30	link-local
	ge-0/0/4.146	172.30.0.90/30	link-local
	ae0.0	172.30.0.94/30	link-local
	ge-0/0/4.167	172.30.0.45/30	link-local
	ge-0/0/4.168	172.30.0.21/30	link-local
	ge-0/0/4.204	172.30.0.57/30	
	lo0.0	172.30.5.6/32	fd17:f0f4:f691:5::6/128
R7	ge-0/0/4.117	172.30.0.6/30	link-local
	ge-0/0/4.137	172.30.0.30/30	link-local
	ge-0/0/4.167	172.30.0.46/30	link-local
	ge-0/0/5.311	192.168.0.41/30	
	ge-0/0/5.312	192.168.0.45/30	
	ge-0/0/5.324	192.168.0.93/30	
	lo0.0	172.30.5.7/32	fd17:f0f4:f691:5::7/128
	lo0.1	172.30.5.33/32	
	lo0.2	172.30.5.34/32	
R8	ge-0/0/4.118	172.30.0.10/30	link-local
	ge-0/0/4.138	172.30.0.34/30	link-local
	ge-0/0/4.168	172.30.0.22/30	link-local
	ge-0/0/5.308	192.168.0.29/30	fc09:c0:ffee::1/126
	ge-0/0/5.302	192.168.0.5/30	
	lo0.0	172.30.5.8/32	fd17:f0f4:f691:5::8/128

© 2020 Juniper Networks, Inc. All rights reserved.

Table 2539

The aggregated interfaces should use LACP. All interfaces should have descriptions.

You have to specify how many **ae** interfaces you will use under the **chassis** hierarchy.

```
[edit]
lab@R1# set chassis aggregated-devices ethernet device-count 1
```

You have to add physical interfaces to the Aggregated Ethernet bundle interfaces.

```
[edit]
lab@R1# set interfaces ge-0/0/1 gigether-options 802.3ad ae0

[edit]
lab@R1# set interfaces ge-0/0/2 gigether-options 802.3ad ae0
```

Configure Ethernet interfaces and descriptions.

```
[edit]
lab@R1# set interfaces ge-0/0/4 vlan-tagging

[edit]
lab@R1# set interfaces ge-0/0/4.117 description "R7 connection"

[edit]
lab@R1# set interfaces ge-0/0/4.117 vlan-id 117

[edit]
lab@R1# set interfaces ge-0/0/4.117 family inet address 172.30.0.5/30

[edit]
lab@R1# set interfaces ge-0/0/4.117 family inet6

[edit]
lab@R1# set interfaces ge-0/0/4.118 description "R8 connection"

[edit]
lab@R1# set interfaces ge-0/0/4.118 vlan-id 118

[edit]
lab@R1# set interfaces ge-0/0/4.118 family inet address 172.30.0.9/30

[edit]
lab@R1# set interfaces ge-0/0/4.118 family inet6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set interfaces ge-0/0/4.206 description "RR connection"

[edit]
lab@R1# set interfaces ge-0/0/4.206 vlan-id 206

[edit]
lab@R1# set interfaces ge-0/0/4.206 family inet address 172.30.0.65/30
```

Configure Aggregated Ethernet interface and enable LACP.

```
[edit]
lab@R1# set interfaces ae0 aggregated-ether-options lacp active

[edit]
lab@R1# set interfaces ae0.0 description "R2 connection"

[edit]
lab@R1# set interfaces ae0.0 family inet address 172.30.0.1/30

[edit]
lab@R1# set interfaces ae0.0 family inet6
```

Configure loopback interface and enable protect-re firewall filter.

```
[edit]
lab@R1# set interfaces lo0.0 family inet address 172.30.5.1/32

[edit]
lab@R1# set interfaces lo0.0 family inet6 address fd17:f0f4:f691:5::1/128

[edit]
lab@R1# set interfaces lo0.0 family inet filter input protect-re

[edit]
lab@R1# commit
commit complete
```

R2:

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 123 description "R3 connection"
set interfaces ge-0/0/4 unit 123 vlan-id 123
set interfaces ge-0/0/4 unit 123 family inet address 172.30.0.13/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces ge-0/0/4 unit 123 family inet6
set interfaces ge-0/0/4 unit 126 description "R6 connection"
set interfaces ge-0/0/4 unit 126 vlan-id 126
set interfaces ge-0/0/4 unit 126 family inet address 172.30.0.17/30
set interfaces ge-0/0/4 unit 126 family inet6
set interfaces ge-0/0/4 unit 207 description "RR connection"
set interfaces ge-0/0/4 unit 207 vlan-id 207
set interfaces ge-0/0/4 unit 207 family inet address 172.30.0.69/30
set interfaces ge-0/0/5 vlan-tagging
set interfaces ge-0/0/5 unit 303 description "C3-1 connection"
set interfaces ge-0/0/5 unit 303 vlan-id 303
set interfaces ge-0/0/5 unit 303 family inet address 192.168.0.9/30
set interfaces ge-0/0/5 unit 303 family inet6 address ::192.168.0.9/126
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 description "R1 connection"
set interfaces ae0 unit 0 family inet address 172.30.0.2/30
set interfaces ae0 unit 0 family inet6
set interfaces lo0 unit 0 family inet filter input protect-re
set interfaces lo0 unit 0 family inet address 172.30.5.2/32
set interfaces lo0 unit 0 family inet6 address fd17:f0f4:f691:5::2/128
```

R3:

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 123 description "R2 connection"
set interfaces ge-0/0/4 unit 123 vlan-id 123
set interfaces ge-0/0/4 unit 123 family inet address 172.30.0.14/30
set interfaces ge-0/0/4 unit 123 family inet6
set interfaces ge-0/0/4 unit 135 description "R5 connection"
set interfaces ge-0/0/4 unit 135 vlan-id 135
set interfaces ge-0/0/4 unit 135 family inet address 172.30.0.85/30
set interfaces ge-0/0/4 unit 135 family inet6
set interfaces ge-0/0/4 unit 137 description "R7 connection"
set interfaces ge-0/0/4 unit 137 vlan-id 137
set interfaces ge-0/0/4 unit 137 family inet address 172.30.0.29/30
set interfaces ge-0/0/4 unit 137 family inet6
set interfaces ge-0/0/4 unit 138 description "R8 connection"
set interfaces ge-0/0/4 unit 138 vlan-id 138
set interfaces ge-0/0/4 unit 138 family inet address 172.30.0.33/30
set interfaces ge-0/0/4 unit 138 family inet6
set interfaces ge-0/0/5 vlan-tagging
set interfaces ge-0/0/5 unit 306 description "C2-1 connection 1"
set interfaces ge-0/0/5 unit 306 vlan-id 306
set interfaces ge-0/0/5 unit 306 family inet address 192.168.0.21/30
set interfaces ge-0/0/5 unit 307 description "C2-1 connection 2"
set interfaces ge-0/0/5 unit 307 vlan-id 307
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces ge-0/0/5 unit 307 family inet address 192.168.0.25/30
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 description "R4 connection"
set interfaces ae0 unit 0 family inet address 172.30.0.81/30
set interfaces ae0 unit 0 family inet6
set interfaces lo0 unit 0 family inet filter input protect-re
set interfaces lo0 unit 0 family inet address 172.30.5.3/32
set interfaces lo0 unit 0 family inet6 address fd17:f0f4:f691:5::3/128
```

R4:

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 146 description "R6 connection"
set interfaces ge-0/0/4 unit 146 vlan-id 146
set interfaces ge-0/0/4 unit 146 family inet address 172.30.0.89/30
set interfaces ge-0/0/4 unit 146 family inet6
set interfaces ge-0/0/5 vlan-tagging
set interfaces ge-0/0/5 unit 323 description "CE2-2 connection"
set interfaces ge-0/0/5 unit 323 vlan-id 323
set interfaces ge-0/0/5 unit 323 family inet address 192.168.0.89/30
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 description "R3 connection"
set interfaces ae0 unit 0 family inet address 172.30.0.82/30
set interfaces ae0 unit 0 family inet6
set interfaces lo0 unit 0 family inet filter input protect-re
set interfaces lo0 unit 0 family inet address 172.30.5.4/32
set interfaces lo0 unit 0 family inet6 address fd17:f0f4:f691:5::4/128
```

R5:

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 135 description "R3 connection"
set interfaces ge-0/0/4 unit 135 vlan-id 135
set interfaces ge-0/0/4 unit 135 family inet address 172.30.0.86/30
set interfaces ge-0/0/4 unit 135 family inet6
set interfaces ge-0/0/4 unit 202 description "DC1 connection"
set interfaces ge-0/0/4 unit 202 vlan-id 202
set interfaces ge-0/0/4 unit 202 family inet address 172.30.0.49/30
set interfaces ge-0/0/5 vlan-tagging
set interfaces ge-0/0/5 unit 305 description "C1-1 connection"
set interfaces ge-0/0/5 unit 305 vlan-id 305
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces ge-0/0/5 unit 305 family inet address 192.168.0.17/30
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 description "R6 connection"
set interfaces ae0 unit 0 family inet address 172.30.0.93/30
set interfaces ae0 unit 0 family inet6
set interfaces lo0 unit 0 family inet filter input protect-re
set interfaces lo0 unit 0 family inet address 172.30.5.5/32
set interfaces lo0 unit 0 family inet6 address fd17:f0f4:f691:5::5/128
```

R6:

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 126 description "R2 connection"
set interfaces ge-0/0/4 unit 126 vlan-id 126
set interfaces ge-0/0/4 unit 126 family inet address 172.30.0.18/30
set interfaces ge-0/0/4 unit 126 family inet6
set interfaces ge-0/0/4 unit 146 description "R4 connection"
set interfaces ge-0/0/4 unit 146 vlan-id 146
set interfaces ge-0/0/4 unit 146 family inet address 172.30.0.90/30
set interfaces ge-0/0/4 unit 146 family inet6
set interfaces ge-0/0/4 unit 167 description "R7 connection"
set interfaces ge-0/0/4 unit 167 vlan-id 167
set interfaces ge-0/0/4 unit 167 family inet address 172.30.0.45/30
set interfaces ge-0/0/4 unit 167 family inet6
set interfaces ge-0/0/4 unit 168 description "R8 connection"
set interfaces ge-0/0/4 unit 168 vlan-id 168
set interfaces ge-0/0/4 unit 168 family inet address 172.30.0.21/30
set interfaces ge-0/0/4 unit 168 family inet6
set interfaces ge-0/0/4 unit 204 description "DC1 connection"
set interfaces ge-0/0/4 unit 204 vlan-id 204
set interfaces ge-0/0/4 unit 204 family inet address 172.30.0.57/30
set interfaces ae0 aggregated-ether-options lacp passive
set interfaces ae0 unit 0 description "R5 connection"
set interfaces ae0 unit 0 family inet address 172.30.0.94/30
set interfaces ae0 unit 0 family inet6
set interfaces lo0 unit 0 family inet filter input protect-re
set interfaces lo0 unit 0 family inet address 172.30.5.6/32
set interfaces lo0 unit 0 family inet6 address fd17:f0f4:f691:5::6/128
```

R7:

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 117 description "R1 connection"
set interfaces ge-0/0/4 unit 117 vlan-id 117
set interfaces ge-0/0/4 unit 117 family inet address 172.30.0.6/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set interfaces ge-0/0/4 unit 117 family inet6
set interfaces ge-0/0/4 unit 137 description "R3 connection"
set interfaces ge-0/0/4 unit 137 vlan-id 137
set interfaces ge-0/0/4 unit 137 family inet address 172.30.0.30/30
set interfaces ge-0/0/4 unit 137 family inet6
set interfaces ge-0/0/4 unit 167 description "R6 connection"
set interfaces ge-0/0/4 unit 167 vlan-id 167
set interfaces ge-0/0/4 unit 167 family inet address 172.30.0.46/30
set interfaces ge-0/0/4 unit 167 family inet6
set interfaces ge-0/0/5 vlan-tagging
set interfaces ge-0/0/5 unit 311 description "CE2-1 connection hub"
set interfaces ge-0/0/5 unit 311 vlan-id 311
set interfaces ge-0/0/5 unit 311 family inet address 192.168.0.41/30
set interfaces ge-0/0/5 unit 312 description "CE2-1 connection spoke"
set interfaces ge-0/0/5 unit 312 vlan-id 312
set interfaces ge-0/0/5 unit 312 family inet address 192.168.0.45/30
set interfaces ge-0/0/5 unit 324 description "CE1-1 connection"
set interfaces ge-0/0/5 unit 324 vlan-id 324
set interfaces ge-0/0/5 unit 324 family inet address 192.168.0.93/30
set interfaces lo0 unit 0 family inet filter input protect-re
set interfaces lo0 unit 0 family inet address 172.30.5.7/32
set interfaces lo0 unit 0 family inet6 address fd17:f0f4:f691:5::7/128
```

R8:

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 118 description "R1 connection"
set interfaces ge-0/0/4 unit 118 vlan-id 118
set interfaces ge-0/0/4 unit 118 family inet address 172.30.0.10/30
set interfaces ge-0/0/4 unit 118 family inet6
set interfaces ge-0/0/4 unit 138 description "R3 connection"
set interfaces ge-0/0/4 unit 138 vlan-id 138
set interfaces ge-0/0/4 unit 138 family inet address 172.30.0.34/30
set interfaces ge-0/0/4 unit 138 family inet6
set interfaces ge-0/0/4 unit 168 description "R6 connection"
set interfaces ge-0/0/4 unit 168 vlan-id 168
set interfaces ge-0/0/4 unit 168 family inet address 172.30.0.22/30
set interfaces ge-0/0/4 unit 168 family inet6
set interfaces ge-0/0/5 vlan-tagging
set interfaces ge-0/0/5 unit 302 description "CE2-1 connection spoke"
set interfaces ge-0/0/5 unit 302 vlan-id 302
set interfaces ge-0/0/5 unit 302 family inet address 192.168.0.5/30
set interfaces ge-0/0/5 unit 308 description "CE2-1 connection hub"
set interfaces ge-0/0/5 unit 308 vlan-id 308
set interfaces ge-0/0/5 unit 308 family inet address 192.168.0.29/30
set interfaces ge-0/0/5 unit 308 family inet6 address fc09:c0:ffee::1/126
set interfaces lo0 unit 0 family inet filter input protect-re
set interfaces lo0 unit 0 family inet address 172.30.5.8/32
set interfaces lo0 unit 0 family inet6 address fd17:f0f4:f691:5::8/128
```


© 2020 Juniper Networks, Inc. All rights reserved.

After configuring all routers according to the given tables, you should have the topology shown on Figure 16.

Task 3: IGP Configuration

You must configure both OSPFv2 and OSPFv3 according to Table 26. You also need to configure the OSPFv2 only connections to route reflector outlined in Table 27.

Router	Interface		Area
R1	ae0.0		0.0.0.0
	ge-0/0/4.117		0.0.0.0
	ge-0/0/4.118		0.0.0.0
	lo0.0		0.0.0.0
R2	ae0.0		0.0.0.0
	ge-0/0/4.126		0.0.0.0
	ge-0/0/4.123		0.0.0.0
	lo0.0		0.0.0.0
R3	ge-0/0/4.123		0.0.0.0
	ge-0/0/4.138		0.0.0.0
	ge-0/0/4.137		0.0.0.0
	ge-0/0/4.135		0.0.0.1
	ae0.0		0.0.0.1
	lo0.0		0.0.0.0
R4	ge-0/0/4.146		0.0.0.1
	ae0.0		0.0.0.1
	lo0.0		0.0.0.1
R5	ge-0/0/4.135		0.0.0.1
	ae0.0		0.0.0.1
	lo0.0		0.0.0.1
R6	ge-0/0/4.126		0.0.0.0
	ge-0/0/4.146		0.0.0.1
	ae0.0		0.0.0.1
	ge-0/0/4.167		0.0.0.0
	ge-0/0/4.168		0.0.0.0
	lo0.0		0.0.0.0
R7	ge-0/0/4.117		0.0.0.0
	ge-0/0/4.137		0.0.0.0
	ge-0/0/4.167		0.0.0.0
	lo0.0		0.0.0.0

© 2020 Juniper Networks, Inc. All rights reserved.

R8	ge-0/0/4.118		0.0.0.0
	ge-0/0/4.138		0.0.0.0
	ge-0/0/4.168		0.0.0.0
	lo0.0		0.0.0.0

Table 26

Router	Interface	Area
R1	ge-0/0/4.206	0.0.0.0
R2	ge-0/0/4.207	0.0.0.0

Table 27

1) Configure OSPFv2 and OSPFv3 only on the core interface except on the route reflector facing interfaces on R1 and R2. Make sure that Router ID is explicitly configured on all routers. Make sure you do not have Type 2 LSA in your domain.

    a) Configure OSPF on R1.

Enable OSPFv2 to the route reflector router. Setting the interface to point-to-point will disable designated router election and generation of LSA Type 2.

```
[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface ae0 interface-type p2p

[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.117 interface-type p2p

[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.118 interface-type p2p

[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.206 interface-type p2p
```

Configure a manual Router ID based on the loopback IP address.

```
[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface lo0.0

[edit]
lab@R1# set protocols ospf3 area 0.0.0.0 interface ae0.0 interface-type p2p


[edit]
lab@R1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.117 interface-type p2p
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.118 interface-type p2p

[edit]
lab@R1# set protocols ospf3 area 0.0.0.0 interface lo0.0

[edit]
lab@R1# set routing-options router-id 172.30.5.1
```

Configure the reference bandwidth to 10g.

```
[edit]
lab@R1# set protocols ospf reference-bandwidth 10g

[edit]
lab@R1# set protocols ospf3 reference-bandwidth 10g

[edit]
lab@R1# commit
commit complete
```

b) Configure OSPF on R2.

```
[edit]
lab@R2# set protocols ospf area 0.0.0.0 interface ae0.0 interface-type p2p

[edit]
lab@R2# set protocols ospf area 0.0.0.0 interface ge-0/0/4.123 interface-type p2p

[edit]
lab@R2# set protocols ospf area 0.0.0.0 interface ge-0/0/4.126 interface-type p2p

[edit]
lab@R2# set protocols ospf area 0.0.0.0 interface ge-0/0/4.207 interface-type p2p

[edit]
lab@R2# set protocols ospf area 0.0.0.0 interface lo0.0

[edit]
lab@R2# set protocols ospf3 area 0.0.0.0 interface ae0.0 interface-type p2p

[edit]
lab@R2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.123 interface-type p2p

[edit]
lab@R2# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.126 interface-type p2p
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set protocols ospf3 area 0.0.0.0 interface lo0.0

[edit]
lab@R2# set routing-options router-id 172.30.5.2

[edit]
lab@R2# set protocols ospf reference-bandwidth 10g

[edit]
lab@R2# set protocols ospf3 reference-bandwidth 10g

[edit]
lab@R2# commit
commit complete
```

c) Configure OSPF on R3.

```
[edit]
lab@R3# set protocols ospf area 0.0.0.0 interface ge-0/0/4.123 interface-type p2p

[edit]
lab@R3# set protocols ospf area 0.0.0.0 interface ge-0/0/4.137 interface-type p2p

[edit]
lab@R3# set protocols ospf area 0.0.0.0 interface ge-0/0/4.138 interface-type p2p

[edit]
lab@R3# set protocols ospf area 0.0.0.0 interface lo0.0

[edit]
lab@R3# set protocols ospf area 0.0.0.1 nssa
```

R3 is ASBR between Area 0 and Area 1. Area 1 is configured as a NSSA, because no LSA Type 4 and Type 5 should be present in the database. However external routes will be exported by R5 from RIP. NSSA's use a LSA Type 7 instead of using a LSA type 5 to advertise external routes.

```
[edit]
lab@R3# set protocols ospf area 0.0.0.1 interface ae0.0 interface-type p2p

[edit]
lab@R3# set protocols ospf area 0.0.0.1 interface ge-0/0/4.135 interface-type p2p


[edit]
lab@R3# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.123 interface-type p2p
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.137 interface-type p2p

[edit]
lab@R3# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.138 interface-type p2p

[edit]
lab@R3# set protocols ospf3 area 0.0.0.0 interface lo0.0

[edit]
lab@R3# set protocols ospf3 area 0.0.0.1 nssa

[edit]
lab@R3# set protocols ospf3 area 0.0.0.1 interface ae0.0 interface-type p2p

[edit]
lab@R3# set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.135 interface-type p2p

[edit]
lab@R3# set routing-options router-id 172.30.5.3

[edit]
lab@R3# set protocols ospf reference-bandwidth 10g

[edit]
lab@R3# set protocols ospf3 reference-bandwidth 10g

[edit]
lab@R3# commit
commit complete
```

d) Configure OSPF on R4.

R4 is an internal router for Area 0.0.0.1.

```
[edit]
lab@R4# set protocols ospf area 0.0.0.1 nssa

[edit]
lab@R4# set protocols ospf area 0.0.0.1 interface ae0.0 interface-type p2p

[edit]
lab@R4# set protocols ospf area 0.0.0.1 interface ge-0/0/4.146 interface-type p2p


[edit]
lab@R4# set protocols ospf area 0.0.0.1 interface lo0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.146 interface-type p2p

[edit]
lab@R4# set protocols ospf3 area 0.0.0.1 interface ae0.0 interface-type p2p

[edit]
lab@R4# set protocols ospf3 area 0.0.0.1 interface lo0.0

[edit]
lab@R4# set protocols ospf3 area 0.0.0.1 nssa

[edit]
lab@R4# set routing-options router-id 172.30.5.4

[edit]
lab@R4# set protocols ospf reference-bandwidth 10g

[edit]
lab@R4# set protocols ospf3 reference-bandwidth 10g

[edit]
lab@R4# commit
commit complete
```

e) Configure OSPF on R5.

R5 is Internal router for Area 0.0.0.1.

```
[edit]
lab@R5# set protocols ospf area 0.0.0.1 nssa

[edit]
lab@R5# set protocols ospf area 0.0.0.1 interface ae0.0 interface-type p2p

[edit]
lab@R5# set protocols ospf area 0.0.0.1 interface ge-0/0/4.135 interface-type p2p

[edit]
lab@R5# set protocols ospf area 0.0.0.1 interface lo0.0

[edit]
lab@R5# set protocols ospf3 area 0.0.0.1 nssa


[edit]
lab@R5# set protocols ospf3 area 0.0.0.1 interface ae0.0 interface-type p2p
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.135 interface-type p2p

[edit]
lab@R5# set protocols ospf3 area 0.0.0.1 interface lo0.0

[edit]
lab@R5# set routing-options router-id 172.30.5.5

[edit]
lab@R5# set protocols ospf reference-bandwidth 10g

[edit]
lab@R5# set protocols ospf3 reference-bandwidth 10g

[edit]
lab@R5# commit
commit complete
```

f) Configure OSPF on R6.

R6 is ASBR between Area 0 and Area 1.

```
[edit]
lab@R6# set protocols ospf area 0.0.0.0 interface ge-0/0/4.126 interface-type p2p

[edit]
lab@R6# set protocols ospf area 0.0.0.0 interface ge-0/0/4.167 interface-type p2p

[edit]
lab@R6# set protocols ospf area 0.0.0.0 interface ge-0/0/4.168 interface-type p2p

[edit]
lab@R6# set protocols ospf area 0.0.0.0 interface lo0.0

[edit]
lab@R6# set protocols ospf area 0.0.0.1 nssa

[edit]
lab@R6# set protocols ospf area 0.0.0.1 interface ae0.0 interface-type p2p

[edit]
lab@R6# set protocols ospf area 0.0.0.1 interface ge-0/0/4.146 interface-type p2p

[edit]
lab@R6# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.126 interface-type p2p

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R6# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.167 interface-type p2p

[edit]
lab@R6# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.168 interface-type p2p

[edit]
lab@R6# set protocols ospf3 area 0.0.0.0 interface lo0.0

[edit]
lab@R6# set protocols ospf3 area 0.0.0.1 nssa

[edit]
lab@R6# set protocols ospf3 area 0.0.0.1 interface ae0.0 interface-type p2p

[edit]
lab@R6# set protocols ospf3 area 0.0.0.1 interface ge-0/0/4.146 interface-type p2p

[edit]
lab@R6# set routing-options router-id 172.30.5.6

[edit]
lab@R6# set protocols ospf reference-bandwidth 10g

[edit]
lab@R6# set protocols ospf3 reference-bandwidth 10g

[edit]
lab@R6# commit
commit complete
```

g) Configure OSPF on R7.

```
[edit]
lab@R7# set protocols ospf area 0.0.0.0 interface ge-0/0/4.117 interface-type p2p

[edit]
lab@R7# set protocols ospf area 0.0.0.0 interface ge-0/0/4.137 interface-type p2p

[edit]
lab@R7# set protocols ospf area 0.0.0.0 interface ge-0/0/4.167 interface-type p2p

[edit]
lab@R7# set protocols ospf area 0.0.0.0 interface lo0.0


[edit]
lab@R7# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.117 interface-type p2p

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.137 interface-type p2p

[edit]
lab@R7# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.167 interface-type p2p

[edit]
lab@R7# set protocols ospf3 area 0.0.0.0 interface lo0.0

[edit]
lab@R7# set routing-options router-id 172.30.5.7

[edit]
lab@R7# set protocols ospf reference-bandwidth 10g

[edit]
lab@R7# set protocols ospf3 reference-bandwidth 10g

[edit]
lab@R7# commit
commit complete
```

h) Configure OSPF on R8.

```
[edit]
lab@R8# set protocols ospf area 0.0.0.0 interface ge-0/0/4.118 interface-type p2p

[edit]
lab@R8# set protocols ospf area 0.0.0.0 interface ge-0/0/4.138 interface-type p2p

[edit]
lab@R8# set protocols ospf area 0.0.0.0 interface ge-0/0/4.168 interface-type p2p

[edit]
lab@R8# set protocols ospf area 0.0.0.0 interface lo0.0

[edit]
lab@R8# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.118 interface-type p2p

[edit]
lab@R8# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.138 interface-type p2p

[edit]
lab@R8# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.168 interface-type p2p

[edit]
lab@R8# set protocols ospf3 area 0.0.0.0 interface lo0.0

[edit]
lab@R8# set routing-options router-id 172.30.5.8
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set protocols ospf reference-bandwidth 10g

[edit]
lab@R8# set protocols ospf3 reference-bandwidth 10g

[edit]
lab@R8# commit
commit complete
```

i) Configure OSPF on RR.

```
[edit]
lab@route-reflector# set protocols ospf area 0.0.0.0 interface ge-0/0/1.206 interface-
type p2p

[edit]
lab@route-reflector# set protocols ospf area 0.0.0.0 interface ge-0/0/1.207 interface-
type p2p

[edit]
lab@route-reflector# set protocols ospf area 0.0.0.0 interface lo0.0

[edit]
lab@route-reflector# commit
commit complete
```

Verify the operation of OSPFv2 and OSPFv3.

a) R1

Verify OSPFv2 and OSPFv3 interfaces and neighbor databases.

```
[edit]
lab@R1# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.2       ae0.0                  Full      172.30.5.2       128    34
172.30.0.6       ge-0/0/4.117           Full      172.30.5.7       128    38
172.30.0.10      ge-0/0/4.118           Full      172.30.5.8       128    35
172.30.0.66      ge-0/0/4.206           Full      172.30.5.41      128    39
```

All neighbors are in UP state.

```
[edit]
lab@R1# run show ospf interface
Interface          State   Area       DR ID          BDR ID          Nbrs
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ae0.0              PtToPt  0.0.0.0    0.0.0.0        0.0.0.0          1
ge-0/0/4.117       PtToPt  0.0.0.0    0.0.0.0        0.0.0.0          1
ge-0/0/4.118       PtToPt  0.0.0.0    0.0.0.0        0.0.0.0          1
ge-0/0/4.206       PtToPt  0.0.0.0    0.0.0.0        0.0.0.0          1
lo0.0              DR      0.0.0.0    172.30.5.1     0.0.0.0          0
```

No DR and BDR are elected because the interfaces are configured for point-to-point connections.

```
[edit]
lab@R1# run show ospf3 neighbor
ID                  Interface              State     Pri   Dead
172.30.5.2          ae0.0                  Full      128     38
  Neighbor-address fe80::fac0:1ff:fedd:280
172.30.5.7          ge-0/0/4.117           Full      128     32
  Neighbor-address fe80::fac0:100:75dc:3204
172.30.5.8          ge-0/0/4.118           Full      128     32
  Neighbor-address fe80::fac0:100:76dc:2e84
```

All neighbors for OSPFv3 are UP.

```
[edit]
lab@R1# run show ospf3 interface
Interface          State   Area       DR ID          BDR ID          Nbrs
ae0.0              PtToPt  0.0.0.0    0.0.0.0        0.0.0.0            1
ge-0/0/4.117       PtToPt  0.0.0.0    0.0.0.0        0.0.0.0            1
ge-0/0/4.118       PtToPt  0.0.0.0    0.0.0.0        0.0.0.0            1
lo0.0              PtToPt  0.0.0.0    0.0.0.0        0.0.0.0            0
```

Verify the OSPF database.

```
lab@R1# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.1       172.30.5.1      0x80000008   131  0x22 0xfc9  132
Router   172.30.5.2       172.30.5.2      0x80000006   137  0x22 0xea7  132
Router   172.30.5.3       172.30.5.3      0x80000006   141  0x22 0x7312 108
Router   172.30.5.6       172.30.5.6      0x80000006   148  0x22 0xd992 108
Router   172.30.5.7       172.30.5.7      0x80000005   154  0x22 0xe792 108
Router   172.30.5.8       172.30.5.8      0x80000006   151  0x22 0x9df8 108
Router   172.30.5.41      172.30.5.41     0x80000004   395  0x22 0x80b3  84
Summary  172.30.0.80      172.30.5.3      0x80000002  1464  0x22 0xf55a  28
Summary  172.30.0.80      172.30.5.6      0x80000002   635  0x22 0xed5e  28
Summary  172.30.0.84      172.30.5.3      0x80000002  1464  0x22 0xcd7e  28
Summary  172.30.0.84      172.30.5.6      0x80000001   635  0x22 0xc781  28
Summary  172.30.0.88      172.30.5.3      0x80000001  1207  0x22 0xb196  28
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Summary   172.30.0.88      172.30.5.6      0x80000002   640  0x22 0x93b1   28
Summary   172.30.0.92      172.30.5.3      0x80000001  1156  0x22 0x89ba   28
Summary   172.30.0.92      172.30.5.6      0x80000002   640  0x22 0x6bd5   28
Summary   172.30.5.4       172.30.5.3      0x80000001  1207  0x22 0xcdc7   28
Summary   172.30.5.4       172.30.5.6      0x80000002   635  0x22 0xb9d7   28
Summary   172.30.5.5       172.30.5.3      0x80000001  1156  0x22 0xc3d0   28
Summary   172.30.5.5       172.30.5.6      0x80000001   635  0x22 0xb1df   28
```

Routes from Area 0.0.0.1 are present as Summary LSA Type 3.

```
[edit]
lab@R1# run show ospf3 database

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.0          172.30.5.1      0x80000004   147  0x4a2e  72
Router      0.0.0.0          172.30.5.2      0x80000005   153  0x7bff  72
Router      0.0.0.0          172.30.5.3      0x80000006   157  0x5019  72
Router      0.0.0.0          172.30.5.6      0x80000006   164  0xa4be  72
Router      0.0.0.0          172.30.5.7      0x80000006   170  0x165a  72
Router      0.0.0.0          172.30.5.8      0x80000005   167  0x78f4  72
InterArPfx  0.0.0.1          172.30.5.3      0x80000001  1223  0x449   44
InterArPfx  0.0.0.2          172.30.5.3      0x80000001  1173  0x1437  44
InterArPfx  0.0.0.1          172.30.5.6      0x80000002   646  0xef59  44
InterArPfx  0.0.0.2          172.30.5.6      0x80000001   656  0x246   44
IntraArPfx *0.0.0.1          172.30.5.1      0x80000003   147  0xb29b  52
IntraArPfx  0.0.0.1          172.30.5.2      0x80000004   153  0xda6f  52
IntraArPfx  0.0.0.1          172.30.5.3      0x80000003   157  0x741   52
IntraArPfx  0.0.0.1          172.30.5.6      0x80000003   164  0x85b9  52
IntraArPfx  0.0.0.1          172.30.5.7      0x80000003   170  0xaf8c  52
IntraArPfx  0.0.0.1          172.30.5.8      0x80000003   167  0xd95f  52

    OSPF3 Link-Local database, interface ae0.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.2          172.30.5.1      0x80000002  1847  0xc765  44
Link        0.0.0.2          172.30.5.2      0x80000002  1743  0x1610  44

    OSPF3 Link-Local database, interface ge-0/0/4.117 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.3          172.30.5.1      0x80000002  1363  0xd65b  44
Link        0.0.0.2          172.30.5.7      0x80000001   482  0xf107  44

    OSPF3 Link-Local database, interface ge-0/0/4.118 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.4          172.30.5.1      0x80000002   879  0xdf50  44
Link        0.0.0.2          172.30.5.8      0x80000001   279  0xb5c4  44

    OSPF3 Link-Local database, interface lo0.0 Area 0.0.0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          172.30.5.1        0x80000002  2331  0xb0f0  44
```

b) R2

```
[edit]
lab@R2# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.1       ae0.0                  Full      172.30.5.1       128    32
172.30.0.14      ge-0/0/4.123           Full      172.30.5.3       128    36
172.30.0.18      ge-0/0/4.126           Full      172.30.5.6       128    34
172.30.0.70      ge-0/0/4.207           Full      172.30.5.41      128    34
```

```
[edit]
lab@R2# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ae0.0              PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
ge-0/0/4.123       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
ge-0/0/4.126       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
ge-0/0/4.207       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
```

```
[edit]
lab@R2# run show ospf3 neighbor
ID                 Interface              State     Pri   Dead
172.30.5.1         ae0.0                  Full      128     38
  Neighbor-address fe80::226:88ff:feef:7600
172.30.5.3         ge-0/0/4.123           Full      128     32
  Neighbor-address fe80::2e21:7200:7bcd:2684
172.30.5.6         ge-0/0/4.126           Full      128     32
  Neighbor-address fe80::fac0:100:7edc:2c04
```

```
[edit]
lab@R2# run show ospf3 interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ae0.0              PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
ge-0/0/4.123       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
ge-0/0/4.126       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
lo0.0              PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            0
```

Verify the OSPFv2 and OSPFv3 databases.

```
[edit]
lab@R2# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Opt  Cksum  Len
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Router   172.30.5.1      172.30.5.1      0x80000008   206  0x22 0xfc9  132
Router  *172.30.5.2      172.30.5.2      0x80000006   210  0x22 0xea7  132
Router   172.30.5.3      172.30.5.3      0x80000006   214  0x22 0x7312 108
Router   172.30.5.6      172.30.5.6      0x80000006   221  0x22 0xd992 108
Router   172.30.5.7      172.30.5.7      0x80000005   229  0x22 0xe792 108
Router   172.30.5.8      172.30.5.8      0x80000006   226  0x22 0x9df8 108
Router   172.30.5.41     172.30.5.41     0x80000004   469  0x22 0x80b3  84
Summary  172.30.0.80     172.30.5.3      0x80000002  1537  0x22 0xf55a  28
Summary  172.30.0.80     172.30.5.6      0x80000002   708  0x22 0xed5e  28
Summary  172.30.0.84     172.30.5.3      0x80000002  1537  0x22 0xcd7e  28
Summary  172.30.0.84     172.30.5.6      0x80000001   708  0x22 0xc781  28
Summary  172.30.0.88     172.30.5.3      0x80000001  1280  0x22 0xb196  28
Summary  172.30.0.88     172.30.5.6      0x80000002   713  0x22 0x93b1  28
Summary  172.30.0.92     172.30.5.3      0x80000001  1229  0x22 0x89ba  28
Summary  172.30.0.92     172.30.5.6      0x80000002   713  0x22 0x6bd5  28
Summary  172.30.5.4      172.30.5.3      0x80000001  1280  0x22 0xcdc7  28
Summary  172.30.5.4      172.30.5.6      0x80000002   708  0x22 0xb9d7  28
Summary  172.30.5.5      172.30.5.3      0x80000001  1229  0x22 0xc3d0  28
Summary  172.30.5.5      172.30.5.6      0x80000001   708  0x22 0xb1df  28
```

```
[edit]
lab@R2# run show ospf3 database

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.0          172.30.5.1      0x80000004   226  0x4a2e  72
Router     *0.0.0.0          172.30.5.2      0x80000005   230  0x7bff  72
Router      0.0.0.0          172.30.5.3      0x80000006   234  0x5019  72
Router      0.0.0.0          172.30.5.6      0x80000006   241  0xa4be  72
Router      0.0.0.0          172.30.5.7      0x80000006   249  0x165a  72
Router      0.0.0.0          172.30.5.8      0x80000005   246  0x78f4  72
InterArPfx  0.0.0.1          172.30.5.3      0x80000001  1300  0x449   44
InterArPfx  0.0.0.2          172.30.5.3      0x80000001  1250  0x1437  44
InterArPfx  0.0.0.1          172.30.5.6      0x80000002   723  0xef59  44
InterArPfx  0.0.0.2          172.30.5.6      0x80000001   733  0x246   44
IntraArPfx  0.0.0.1          172.30.5.1      0x80000003   226  0xb29b  52
IntraArPfx *0.0.0.1          172.30.5.2      0x80000004   230  0xda6f  52
IntraArPfx  0.0.0.1          172.30.5.3      0x80000003   234  0x741   52
IntraArPfx  0.0.0.1          172.30.5.6      0x80000003   241  0x85b9  52
IntraArPfx  0.0.0.1          172.30.5.7      0x80000003   249  0xaf8c  52
IntraArPfx  0.0.0.1          172.30.5.8      0x80000003   246  0xd95f  52

    OSPF3 Link-Local database, interface ae0.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.2          172.30.5.1      0x80000002  1926  0xc765  44
Link       *0.0.0.2          172.30.5.2      0x80000002  1820  0x1610  44

    OSPF3 Link-Local database, interface ge-0/0/4.123 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Link        *0.0.0.3            172.30.5.2        0x80000002  1370  0x968e  44
Link         0.0.0.3            172.30.5.3        0x80000002   810  0xdfaa  44

    OSPF3 Link-Local database, interface ge-0/0/4.126 Area 0.0.0.0
 Type        ID                 Adv Rtr            Seq          Age  Cksum  Len
Link        *0.0.0.4            172.30.5.2        0x80000002   919  0xc55b  44
Link         0.0.0.3            172.30.5.6        0x80000002    81  0x19db  44

    OSPF3 Link-Local database, interface lo0.0 Area 0.0.0.0
 Type        ID                 Adv Rtr            Seq          Age  Cksum  Len
Link        *0.0.0.1            172.30.5.2        0x80000002  2271  0xfd9c  44
```

c) R3

```
[edit]
lab@R3# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.13      ge-0/0/4.123           Full      172.30.5.2       128    31
172.30.0.30      ge-0/0/4.137           Full      172.30.5.7       128    37
172.30.0.34      ge-0/0/4.138           Full      172.30.5.8       128    35
172.30.0.82      ae0.0                  Full      172.30.5.4       128    38
172.30.0.86      ge-0/0/4.135           Full      172.30.5.5       128    34
```

```
[edit]
lab@R3# run show ospf interface
Interface          State   Area       DR ID           BDR ID          Nbrs
ge-0/0/4.123       PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
ge-0/0/4.137       PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
ge-0/0/4.138       PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
ae0.0              PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
ge-0/0/4.135       PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
```

R3 is an ASBR and has neighbors in both Areas. All neighbors States are Established.

```
[edit]
lab@R3# run show ospf3 neighbor
ID                 Interface              State     Pri   Dead
172.30.5.2         ge-0/0/4.123           Full      128     37
  Neighbor-address fe80::fac0:100:7bdd:204
172.30.5.7         ge-0/0/4.137           Full      128     31
  Neighbor-address fe80::fac0:100:89dc:3204
172.30.5.8         ge-0/0/4.138           Full      128     37
  Neighbor-address fe80::fac0:100:8adc:2e84
172.30.5.4         ae0.0                  Full      128     32
  Neighbor-address fe80::fac0:1ff:fedc:3200
172.30.5.5         ge-0/0/4.135           Full      128     32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Neighbor-address fe80::fac0:100:87dd:384
```

```
[edit]
lab@R3# run show ospf3 interface
Interface          State   Area       DR ID          BDR ID          Nbrs
ge-0/0/4.123       PtToPt  0.0.0.0    0.0.0.0        0.0.0.0            1
ge-0/0/4.137       PtToPt  0.0.0.0    0.0.0.0        0.0.0.0            1
ge-0/0/4.138       PtToPt  0.0.0.0    0.0.0.0        0.0.0.0            1
lo0.0              PtToPt  0.0.0.0    0.0.0.0        0.0.0.0            0
ae0.0              PtToPt  0.0.0.1    0.0.0.0        0.0.0.0            1
ge-0/0/4.135       PtToPt  0.0.0.1    0.0.0.0        0.0.0.0            1
```

Verify the OSPFv2 and OSPFv3 databases.

```
[edit]
lab@R3# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.1       172.30.5.1       0x80000008   270  0x22 0xfc9  132
Router   172.30.5.2       172.30.5.2       0x80000006   274  0x22 0xea7  132
Router  *172.30.5.3       172.30.5.3       0x80000006   276  0x22 0x7312 108
Router   172.30.5.6       172.30.5.6       0x80000006   285  0x22 0xd992 108
Router   172.30.5.7       172.30.5.7       0x80000005   291  0x22 0xe792 108
Router   172.30.5.8       172.30.5.8       0x80000006   288  0x22 0x9df8 108
Router   172.30.5.41      172.30.5.41      0x80000004   533  0x22 0x80b3  84
Summary *172.30.0.80      172.30.5.3       0x80000002  1599  0x22 0xf55a  28
Summary  172.30.0.80      172.30.5.6       0x80000002   772  0x22 0xed5e  28
Summary *172.30.0.84      172.30.5.3       0x80000002  1599  0x22 0xcd7e  28
Summary  172.30.0.84      172.30.5.6       0x80000001   772  0x22 0xc781  28
Summary *172.30.0.88      172.30.5.3       0x80000001  1342  0x22 0xb196  28
Summary  172.30.0.88      172.30.5.6       0x80000002   778  0x22 0x93b1  28
Summary *172.30.0.92      172.30.5.3       0x80000001  1291  0x22 0x89ba  28
Summary  172.30.0.92      172.30.5.6       0x80000002   778  0x22 0x6bd5  28
Summary *172.30.5.4       172.30.5.3       0x80000001  1342  0x22 0xcdc7  28
Summary  172.30.5.4       172.30.5.6       0x80000002   772  0x22 0xb9d7  28
Summary *172.30.5.5       172.30.5.3       0x80000001  1291  0x22 0xc3d0  28
Summary  172.30.5.5       172.30.5.6       0x80000001   772  0x22 0xb1df  28

    OSPF database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.3       172.30.5.3       0x80000006   276  0x20 0x71cc  72
Router   172.30.5.4       172.30.5.4       0x80000005   280  0x20 0x64ee  84
Router   172.30.5.5       172.30.5.5       0x80000005   283  0x20 0xbe81  84
Router   172.30.5.6       172.30.5.6       0x80000005   285  0x20 0x62b4  72
Summary *172.30.0.0       172.30.5.3       0x80000002  1144  0x20 0x4160  28
Summary  172.30.0.0       172.30.5.6       0x80000002   371  0x20 0x2f6f  28
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Summary *172.30.0.4      172.30.5.3      0x80000003   596  0x20 0x1785  28
Summary  172.30.0.4      172.30.5.6      0x80000002   585  0x20 0x793   28
Summary *172.30.0.8      172.30.5.3      0x80000004   389  0x20 0xecaa  28
Summary  172.30.0.8      172.30.5.6      0x80000003   384  0x20 0xdcb8  28
Summary *172.30.0.12     172.30.5.3      0x80000003    50  0x20 0xbcd8  28
Summary  172.30.0.12     172.30.5.6      0x80000002   254  0x20 0xb6db  28
Summary *172.30.0.16     172.30.5.3      0x80000002   681  0x20 0xa0f0  28
Summary  172.30.0.16     172.30.5.6      0x80000002   778  0x20 0x840b  28
Summary *172.30.0.20     172.30.5.3      0x80000003   389  0x20 0x7616  28
Summary  172.30.0.20     172.30.5.6      0x80000002   778  0x20 0x5c2f  28
Summary *172.30.0.28     172.30.5.3      0x80000002  1599  0x20 0x1e68  28
Summary  172.30.0.28     172.30.5.6      0x80000002   585  0x20 0x166c  28
Summary *172.30.0.32     172.30.5.3      0x80000002  1599  0x20 0xf58c  28
Summary  172.30.0.32     172.30.5.6      0x80000003   384  0x20 0xeb91  28
Summary *172.30.0.44     172.30.5.3      0x80000002   596  0x20 0x87ed  28
Summary  172.30.0.44     172.30.5.6      0x80000002   778  0x20 0x6b08  28
Summary *172.30.0.64     172.30.5.3      0x80000003   389  0x20 0xc698  28
Summary  172.30.0.64     172.30.5.6      0x80000003   384  0x20 0xb4a7  28
Summary *172.30.0.68     172.30.5.3      0x80000002   417  0x20 0x96c6  28
Summary  172.30.0.68     172.30.5.6      0x80000002   137  0x20 0x84d5  28
Summary *172.30.5.1      172.30.5.3      0x80000003   389  0x20 0x1087  28
Summary  172.30.5.1      172.30.5.6      0x80000003   384  0x20 0xfd96  28
Summary *172.30.5.2      172.30.5.3      0x80000002   289  0x20 0xfd9a  28
Summary  172.30.5.2      172.30.5.6      0x80000002    21  0x20 0xeba9  28
Summary *172.30.5.3      172.30.5.3      0x80000002   169  0x20 0xe9ae  28
Summary  172.30.5.3      172.30.5.6      0x80000003   384  0x20 0xe9a8  28
Summary *172.30.5.6      172.30.5.3      0x80000003   382  0x20 0xddb4  28
Summary  172.30.5.6      172.30.5.6      0x80000002   778  0x20 0xb9d8  28
Summary *172.30.5.7      172.30.5.3      0x80000001   596  0x20 0xcdc6  28
Summary  172.30.5.7      172.30.5.6      0x80000002   585  0x20 0xb9d6  28
Summary *172.30.5.8      172.30.5.3      0x80000002   389  0x20 0xc1d0  28
Summary  172.30.5.8      172.30.5.6      0x80000002   384  0x20 0xafdf  28
Summary *172.30.5.41     172.30.5.3      0x80000002   559  0x20 0x80ef  28
Summary  172.30.5.41     172.30.5.6      0x80000002   561  0x20 0x6efe  28
```

Routes are exchanged between both Area 0.0.0.0 and Area 0.0.0.1.

```
[edit]
lab@R3# run show ospf3 database

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.0          172.30.5.1      0x80000004   323  0x4a2e  72
Router      0.0.0.0          172.30.5.2      0x80000005   327  0x7bff  72
Router     *0.0.0.0          172.30.5.3      0x80000006   329  0x5019  72
Router      0.0.0.0          172.30.5.6      0x80000006   338  0xa4be  72
Router      0.0.0.0          172.30.5.7      0x80000006   344  0x165a  72
Router      0.0.0.0          172.30.5.8      0x80000005   341  0x78f4  72
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
InterArPfx *0.0.0.1          172.30.5.3       0x80000001  1395  0x449   44
InterArPfx *0.0.0.2          172.30.5.3       0x80000001  1345  0x1437  44
InterArPfx  0.0.0.1          172.30.5.6       0x80000002   820  0xef59  44
InterArPfx  0.0.0.2          172.30.5.6       0x80000001   830  0x246   44
IntraArPfx  0.0.0.1          172.30.5.1       0x80000003   323  0xb29b  52
IntraArPfx  0.0.0.1          172.30.5.2       0x80000004   327  0xda6f  52
IntraArPfx *0.0.0.1          172.30.5.3       0x80000003   329  0x741   52
IntraArPfx  0.0.0.1          172.30.5.6       0x80000003   338  0x85b9  52
IntraArPfx  0.0.0.1          172.30.5.7       0x80000003   344  0xaf8c  52
IntraArPfx  0.0.0.1          172.30.5.8       0x80000003   341  0xd95f  52

    OSPF3 database, Area 0.0.0.1
 Type       ID               Adv Rtr          Seq         Age  Cksum  Len
Router     *0.0.0.0          172.30.5.3       0x80000004   329  0x7bd6  56
Router      0.0.0.0          172.30.5.4       0x80000005   333  0x1a3d  56
Router      0.0.0.0          172.30.5.5       0x80000005   336  0x1245  56
Router      0.0.0.0          172.30.5.6       0x80000004   339  0x9abd  56
InterArPfx *0.0.0.1          172.30.5.3       0x80000003   445  0xb795  44
InterArPfx *0.0.0.2          172.30.5.3       0x80000001  1407  0xc588  44
InterArPfx *0.0.0.3          172.30.5.3       0x80000001  1407  0xcf7d  44
InterArPfx *0.0.0.4          172.30.5.3       0x80000003   437  0x1c29  44
InterArPfx *0.0.0.5          172.30.5.3       0x80000001   647  0x2a1c  44
InterArPfx *0.0.0.6          172.30.5.3       0x80000001   445  0x3a0a  44
InterArPfx  0.0.0.1          172.30.5.6       0x80000001   836  0x2029  44
InterArPfx  0.0.0.2          172.30.5.6       0x80000003   439  0x9bad  44
InterArPfx  0.0.0.3          172.30.5.6       0x80000001   831  0xa9a0  44
InterArPfx  0.0.0.4          172.30.5.6       0x80000003   439  0xbb89  44
InterArPfx  0.0.0.5          172.30.5.6       0x80000002   644  0x162c  44
InterArPfx  0.0.0.6          172.30.5.6       0x80000002   439  0x261a  44
IntraArPfx  0.0.0.1          172.30.5.4       0x80000003   333  0x3114  52
IntraArPfx  0.0.0.1          172.30.5.5       0x80000003   336  0x5be6  52

    OSPF3 Link-Local database, interface ae0.0 Area 0.0.0.1
 Type       ID               Adv Rtr          Seq         Age  Cksum  Len
Link       *0.0.0.7          172.30.5.3       0x80000002    90  0x2068  44
Link        0.0.0.2          172.30.5.4       0x80000002   656  0xd2a4  44

    OSPF3 Link-Local database, interface ge-0/0/4.123 Area 0.0.0.0
 Type       ID               Adv Rtr          Seq         Age  Cksum  Len
Link        0.0.0.3          172.30.5.2       0x80000002  1467  0x968e  44
Link       *0.0.0.3          172.30.5.3       0x80000002   905  0xdfaa  44

    OSPF3 Link-Local database, interface ge-0/0/4.135 Area 0.0.0.1
 Type       ID               Adv Rtr          Seq         Age  Cksum  Len
Link       *0.0.0.8          172.30.5.3       0x80000001  1407  0x88f3  44
Link        0.0.0.3          172.30.5.5       0x80000001  1353  0x7f18  44

    OSPF3 Link-Local database, interface ge-0/0/4.137 Area 0.0.0.0
 Type       ID               Adv Rtr          Seq         Age  Cksum  Len
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Link        *0.0.0.5          172.30.5.3        0x80000002   692  0xd6a3  44
Link         0.0.0.3          172.30.5.7        0x80000001   656  0x657e  44

    OSPF3 Link-Local database, interface ge-0/0/4.138 Area 0.0.0.0
 Type        ID               Adv Rtr           Seq         Age  Cksum  Len
Link        *0.0.0.6          172.30.5.3        0x80000002   488  0xdf98  44
Link         0.0.0.3          172.30.5.8        0x80000001   453  0x293c  44

    OSPF3 Link-Local database, interface lo0.0 Area 0.0.0.0
 Type        ID               Adv Rtr           Seq         Age  Cksum  Len
Link        *0.0.0.1          172.30.5.3        0x80000002  1134  0x47b8  44
```

d) R4

```
[edit]
lab@R4# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.81      ae0.0                  Full      172.30.5.3       128    35
172.30.0.90      ge-0/0/4.146           Full      172.30.5.6       128    31
```

```
[edit]
lab@R4# run show ospf interface
Interface          State   Area       DR ID           BDR ID          Nbrs
ae0.0              PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
ge-0/0/4.146       PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
lo0.0              DR      0.0.0.1    172.30.5.4      0.0.0.0            0
```

R4 is Internal to Area 0.0.0.1, only neighbors in that Area are established.

```
[edit]
lab@R4# run show ospf3 neighbor
ID               Interface              State     Pri   Dead
172.30.5.3       ae0.0                  Full      128     33
  Neighbor-address fe80::2e21:72ff:fecd:2700
172.30.5.6       ge-0/0/4.146           Full      128     34
  Neighbor-address fe80::fac0:100:92dc:2c04
```

```
[edit]
lab@R4# run show ospf3 interface
Interface          State   Area       DR ID           BDR ID          Nbrs
ae0.0              PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
ge-0/0/4.146       PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
lo0.0              DR      0.0.0.1    172.30.5.4      0.0.0.0            0
```

Verify the OSPFv2 and OSPFv3 databases.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show ospf database

    OSPF database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Router   172.30.5.3       172.30.5.3       0x80000006   373  0x20 0x71cc  72
Router  *172.30.5.4       172.30.5.4       0x80000005   375  0x20 0x64ee  84
Router   172.30.5.5       172.30.5.5       0x80000005   380  0x20 0xbe81  84
Router   172.30.5.6       172.30.5.6       0x80000005   381  0x20 0x62b4  72
Summary  172.30.0.0       172.30.5.3       0x80000002  1241  0x20 0x4160  28
Summary  172.30.0.0       172.30.5.6       0x80000002   466  0x20 0x2f6f  28
Summary  172.30.0.4       172.30.5.3       0x80000003   694  0x20 0x1785  28
Summary  172.30.0.4       172.30.5.6       0x80000002   680  0x20 0x793   28
Summary  172.30.0.8       172.30.5.3       0x80000004   486  0x20 0xecaa  28
Summary  172.30.0.8       172.30.5.6       0x80000003   479  0x20 0xdcb8  28
Summary  172.30.0.12      172.30.5.3       0x80000003   147  0x20 0xbcd8  28
Summary  172.30.0.12      172.30.5.6       0x80000002   349  0x20 0xb6db  28
Summary  172.30.0.16      172.30.5.3       0x80000002   778  0x20 0xa0f0  28
Summary  172.30.0.16      172.30.5.6       0x80000002   875  0x20 0x840b  28
Summary  172.30.0.20      172.30.5.3       0x80000003   486  0x20 0x7616  28
Summary  172.30.0.20      172.30.5.6       0x80000002   875  0x20 0x5c2f  28
Summary  172.30.0.28      172.30.5.3       0x80000003    32  0x20 0x1c69  28
Summary  172.30.0.28      172.30.5.6       0x80000002   680  0x20 0x166c  28
Summary  172.30.0.32      172.30.5.3       0x80000002  1697  0x20 0xf58c  28
Summary  172.30.0.32      172.30.5.6       0x80000003   479  0x20 0xeb91  28
Summary  172.30.0.44      172.30.5.3       0x80000002   694  0x20 0x87ed  28
Summary  172.30.0.44      172.30.5.6       0x80000002   875  0x20 0x6b08  28
Summary  172.30.0.64      172.30.5.3       0x80000003   486  0x20 0xc698  28
Summary  172.30.0.64      172.30.5.6       0x80000003   479  0x20 0xb4a7  28
Summary  172.30.0.68      172.30.5.3       0x80000002   514  0x20 0x96c6  28
Summary  172.30.0.68      172.30.5.6       0x80000002   232  0x20 0x84d5  28
Summary  172.30.5.1       172.30.5.3       0x80000003   486  0x20 0x1087  28
Summary  172.30.5.1       172.30.5.6       0x80000003   479  0x20 0xfd96  28
Summary  172.30.5.2       172.30.5.3       0x80000002   387  0x20 0xfd9a  28
Summary  172.30.5.2       172.30.5.6       0x80000002   117  0x20 0xeba9  28
Summary  172.30.5.3       172.30.5.3       0x80000002   267  0x20 0xe9ae  28
Summary  172.30.5.3       172.30.5.6       0x80000003   479  0x20 0xe9a8  28
Summary  172.30.5.6       172.30.5.3       0x80000003   479  0x20 0xddb4  28
Summary  172.30.5.6       172.30.5.6       0x80000002   875  0x20 0xb9d8  28
Summary  172.30.5.7       172.30.5.3       0x80000001   694  0x20 0xcdc6  28
Summary  172.30.5.7       172.30.5.6       0x80000002   680  0x20 0xb9d6  28
Summary  172.30.5.8       172.30.5.3       0x80000002   486  0x20 0xc1d0  28
Summary  172.30.5.8       172.30.5.6       0x80000002   479  0x20 0xafdf  28
Summary  172.30.5.41      172.30.5.3       0x80000002   657  0x20 0x80ef  28
Summary  172.30.5.41      172.30.5.6       0x80000002   657  0x20 0x6efe  28
```

R4 receives Summary LSAs for all Area 0.0.0.0 routes. All the loopback IP addresses, including the one from RR are present in the OSPF database.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show ospf3 database

    OSPF3 database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.0          172.30.5.3      0x80000004   390  0x7bd6  56
Router     *0.0.0.0          172.30.5.4      0x80000005   392  0x1a3d  56
Router      0.0.0.0          172.30.5.5      0x80000005   397  0x1245  56
Router      0.0.0.0          172.30.5.6      0x80000004   398  0x9abd  56
InterArPfx  0.0.0.1          172.30.5.3      0x80000003   506  0xb795  44
InterArPfx  0.0.0.2          172.30.5.3      0x80000001  1469  0xc588  44
InterArPfx  0.0.0.3          172.30.5.3      0x80000001  1469  0xcf7d  44
InterArPfx  0.0.0.4          172.30.5.3      0x80000003   498  0x1c29  44
InterArPfx  0.0.0.5          172.30.5.3      0x80000001   709  0x2a1c  44
InterArPfx  0.0.0.6          172.30.5.3      0x80000001   506  0x3a0a  44
InterArPfx  0.0.0.1          172.30.5.6      0x80000001   897  0x2029  44
InterArPfx  0.0.0.2          172.30.5.6      0x80000003   498  0x9bad  44
InterArPfx  0.0.0.3          172.30.5.6      0x80000001   892  0xa9a0  44
InterArPfx  0.0.0.4          172.30.5.6      0x80000003   498  0xbb89  44
InterArPfx  0.0.0.5          172.30.5.6      0x80000002   704  0x162c  44
InterArPfx  0.0.0.6          172.30.5.6      0x80000002   498  0x261a  44
IntraArPfx *0.0.0.1          172.30.5.4      0x80000003   392  0x3114  52
IntraArPfx  0.0.0.1          172.30.5.5      0x80000003   397  0x5be6  52

    OSPF3 Link-Local database, interface ae0.0 Area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.7          172.30.5.3      0x80000002   152  0x2068  44
Link       *0.0.0.2          172.30.5.4      0x80000002   715  0xd2a4  44

    OSPF3 Link-Local database, interface ge-0/0/4.146 Area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.3          172.30.5.4      0x80000002   142  0xb54   44
Link        0.0.0.4          172.30.5.6      0x80000002    60  0x8061  44

    OSPF3 Link-Local database, interface lo0.0 Area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          172.30.5.4      0x80000002  1423  0xbb30  44
```

e) R5.

```
[edit]
lab@R5# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.94      ae0.0                  Full      172.30.5.6       128    36
172.30.0.85      ge-0/0/4.135           Full      172.30.5.3       128    32
```

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# run show ospf interface
Interface         State   Area       DR ID           BDR ID          Nbrs
ae0.0             PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
ge-0/0/4.135      PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
lo0.0             DR      0.0.0.1    172.30.5.5      0.0.0.0            0
```

```
[edit]
lab@R5# run show ospf3 neighbor
ID                  Interface              State     Pri   Dead
172.30.5.6          ae0.0                  Full      128     34
  Neighbor-address fe80::fac0:1ff:fedc:2c80
172.30.5.3          ge-0/0/4.135           Full      128     34
  Neighbor-address fe80::2e21:7200:87cd:2684
```

```
[edit]
lab@R5# run show ospf3 interface
Interface         State   Area       DR ID           BDR ID          Nbrs
ae0.0             PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
ge-0/0/4.135      PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
lo0.0             DR      0.0.0.1    172.30.5.5      0.0.0.0            0
```

Verify the OSPFv2 and OSPFv3 databases.

```
[edit]
lab@R5# run show ospf database

    OSPF database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Router   172.30.5.3       172.30.5.3       0x80000006   411  0x20 0x71cc  72
Router   172.30.5.4       172.30.5.4       0x80000005   415  0x20 0x64ee  84
Router  *172.30.5.5       172.30.5.5       0x80000005   416  0x20 0xbe81  84
Router   172.30.5.6       172.30.5.6       0x80000005   419  0x20 0x62b4  72
Summary  172.30.0.0       172.30.5.3       0x80000002  1279  0x20 0x4160  28
Summary  172.30.0.0       172.30.5.6       0x80000002   504  0x20 0x2f6f  28
Summary  172.30.0.4       172.30.5.3       0x80000003   732  0x20 0x1785  28
Summary  172.30.0.4       172.30.5.6       0x80000002   718  0x20 0x793   28
Summary  172.30.0.8       172.30.5.3       0x80000004   524  0x20 0xecaa  28
Summary  172.30.0.8       172.30.5.6       0x80000003   517  0x20 0xdcb8  28
Summary  172.30.0.12      172.30.5.3       0x80000003   185  0x20 0xbcd8  28
Summary  172.30.0.12      172.30.5.6       0x80000002   387  0x20 0xb6db  28
Summary  172.30.0.16      172.30.5.3       0x80000002   816  0x20 0xa0f0  28
Summary  172.30.0.16      172.30.5.6       0x80000003    39  0x20 0x820c  28
Summary  172.30.0.20      172.30.5.3       0x80000003   524  0x20 0x7616  28
Summary  172.30.0.20      172.30.5.6       0x80000002   911  0x20 0x5c2f  28
Summary  172.30.0.28      172.30.5.3       0x80000003    70  0x20 0x1c69  28
Summary  172.30.0.28      172.30.5.6       0x80000002   718  0x20 0x166c  28
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Summary  172.30.0.32     172.30.5.3      0x80000002  1734  0x20 0xf58c  28
Summary  172.30.0.32     172.30.5.6      0x80000003   517  0x20 0xeb91  28
Summary  172.30.0.44     172.30.5.3      0x80000002   732  0x20 0x87ed  28
Summary  172.30.0.44     172.30.5.6      0x80000002   911  0x20 0x6b08  28
Summary  172.30.0.64     172.30.5.3      0x80000003   524  0x20 0xc698  28
Summary  172.30.0.64     172.30.5.6      0x80000003   517  0x20 0xb4a7  28
Summary  172.30.0.68     172.30.5.3      0x80000002   552  0x20 0x96c6  28
Summary  172.30.0.68     172.30.5.6      0x80000002   270  0x20 0x84d5  28
Summary  172.30.5.1      172.30.5.3      0x80000003   524  0x20 0x1087  28
Summary  172.30.5.1      172.30.5.6      0x80000003   517  0x20 0xfd96  28
Summary  172.30.5.2      172.30.5.3      0x80000002   425  0x20 0xfd9a  28
Summary  172.30.5.2      172.30.5.6      0x80000002   155  0x20 0xeba9  28
Summary  172.30.5.3      172.30.5.3      0x80000002   306  0x20 0xe9ae  28
Summary  172.30.5.3      172.30.5.6      0x80000003   518  0x20 0xe9a8  28
Summary  172.30.5.6      172.30.5.3      0x80000003   518  0x20 0xddb4  28
Summary  172.30.5.6      172.30.5.6      0x80000002   912  0x20 0xb9d8  28
Summary  172.30.5.7      172.30.5.3      0x80000001   733  0x20 0xcdc6  28
Summary  172.30.5.7      172.30.5.6      0x80000002   719  0x20 0xb9d6  28
Summary  172.30.5.8      172.30.5.3      0x80000002   525  0x20 0xc1d0  28
Summary  172.30.5.8      172.30.5.6      0x80000002   518  0x20 0xafdf  28
Summary  172.30.5.41     172.30.5.3      0x80000002   695  0x20 0x80ef  28
Summary  172.30.5.41     172.30.5.6      0x80000002   695  0x20 0x6efe  28
```

The database is identical with the database on R4. R5 like R4 is an internal router in Area 0.0.0.1.

```
[edit]
lab@R5# run show ospf3 database | no-more

    OSPF3 database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.0          172.30.5.3      0x80000004   452  0x7bd6  56
Router      0.0.0.0          172.30.5.4      0x80000005   456  0x1a3d  56
Router     *0.0.0.0          172.30.5.5      0x80000005   457  0x1245  56
Router      0.0.0.0          172.30.5.6      0x80000004   460  0x9abd  56
InterArPfx  0.0.0.1          172.30.5.3      0x80000003   568  0xb795  44
InterArPfx  0.0.0.2          172.30.5.3      0x80000001  1530  0xc588  44
InterArPfx  0.0.0.3          172.30.5.3      0x80000001  1530  0xcf7d  44
InterArPfx  0.0.0.4          172.30.5.3      0x80000003   560  0x1c29  44
InterArPfx  0.0.0.5          172.30.5.3      0x80000001   771  0x2a1c  44
InterArPfx  0.0.0.6          172.30.5.3      0x80000001   568  0x3a0a  44
InterArPfx  0.0.0.1          172.30.5.6      0x80000001   957  0x2029  44
InterArPfx  0.0.0.2          172.30.5.6      0x80000003   560  0x9bad  44
InterArPfx  0.0.0.3          172.30.5.6      0x80000001   952  0xa9a0  44
InterArPfx  0.0.0.4          172.30.5.6      0x80000003   560  0xbb89  44
InterArPfx  0.0.0.5          172.30.5.6      0x80000002   765  0x162c  44
InterArPfx  0.0.0.6          172.30.5.6      0x80000002   560  0x261a  44
IntraArPfx  0.0.0.1          172.30.5.4      0x80000003   456  0x3114  52
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
IntraArPfx *0.0.0.1          172.30.5.5       0x80000003   457  0x5be6  52

    OSPF3 Link-Local database, interface ae0.0 Area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.2          172.30.5.5       0x80000002   626  0x178c  44
Link        0.0.0.2          172.30.5.6       0x80000002   480  0x53a7  44

    OSPF3 Link-Local database, interface ge-0/0/4.135 Area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.8          172.30.5.3       0x80000002    42  0x86f4  44
Link       *0.0.0.3          172.30.5.5       0x80000002    78  0x7d19  44

    OSPF3 Link-Local database, interface lo0.0 Area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          172.30.5.5       0x80000002  1434  0xff18  44
```

f) R6

```
[edit]
lab@R6# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.17      ge-0/0/4.126           Full      172.30.5.2       128    32
172.30.0.46      ge-0/0/4.167           Full      172.30.5.7       128    39
172.30.0.22      ge-0/0/4.168           Full      172.30.5.8       128    37
172.30.0.93      ae0.0                  Full      172.30.5.5       128    38
172.30.0.89      ge-0/0/4.146           Full      172.30.5.4       128    38
```

```
[edit]
lab@R6# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.126       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
ge-0/0/4.167       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
ge-0/0/4.168       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
lo0.0              DR      0.0.0.0         172.30.5.6      0.0.0.0            0
ae0.0              PtToPt  0.0.0.1         0.0.0.0         0.0.0.0            1
ge-0/0/4.146       PtToPt  0.0.0.1         0.0.0.0         0.0.0.0            1
```

```
[edit]
lab@R6# run show ospf3 neighbor
ID                 Interface              State     Pri   Dead
172.30.5.2         ge-0/0/4.126           Full      128     38
  Neighbor-address fe80::fac0:100:7edd:204
172.30.5.7         ge-0/0/4.167           Full      128     35
  Neighbor-address fe80::fac0:100:a7dc:3204
172.30.5.8         ge-0/0/4.168           Full      128     32
  Neighbor-address fe80::fac0:100:a8dc:2e84
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.5        ae0.0                    Full       128     36
  Neighbor-address fe80::fac0:1ff:fedd:400
172.30.5.4        ge-0/0/4.146             Full       128     39
  Neighbor-address fe80::fac0:100:92dc:3184
```

```
[edit]
lab@R6# run show ospf3 interface
Interface          State   Area       DR ID           BDR ID          Nbrs
ge-0/0/4.126       PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
ge-0/0/4.167       PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
ge-0/0/4.168       PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
lo0.0              DR      0.0.0.0    172.30.5.6      0.0.0.0            0
ae0.0              PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
ge-0/0/4.146       PtToPt  0.0.0.1    0.0.0.0         0.0.0.0            1
```

Verify the OSPFv2 and OSPFv3 databases.

```
[edit]
lab@R6# run show ospf database | no-more

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.1       172.30.5.1       0x80000008   458  0x22 0xfc9  132
Router   172.30.5.2       172.30.5.2       0x80000006   461  0x22 0xea7  132
Router   172.30.5.3       172.30.5.3       0x80000006   465  0x22 0x7312 108
Router  *172.30.5.6       172.30.5.6       0x80000006   471  0x22 0xd992 108
Router   172.30.5.7       172.30.5.7       0x80000005   479  0x22 0xe792 108
Router   172.30.5.8       172.30.5.8       0x80000006   476  0x22 0x9df8 108
Router   172.30.5.41      172.30.5.41      0x80000004   721  0x22 0x80b3  84
Summary  172.30.0.80      172.30.5.3       0x80000002  1789  0x22 0xf55a  28
Summary *172.30.0.80      172.30.5.6       0x80000002   958  0x22 0xed5e  28
Summary  172.30.0.84      172.30.5.3       0x80000002  1789  0x22 0xcd7e  28
Summary *172.30.0.84      172.30.5.6       0x80000001   958  0x22 0xc781  28
Summary  172.30.0.88      172.30.5.3       0x80000001  1532  0x22 0xb196  28
Summary *172.30.0.88      172.30.5.6       0x80000002   963  0x22 0x93b1  28
Summary  172.30.0.92      172.30.5.3       0x80000001  1481  0x22 0x89ba  28
Summary *172.30.0.92      172.30.5.6       0x80000002   963  0x22 0x6bd5  28
Summary  172.30.5.4       172.30.5.3       0x80000001  1532  0x22 0xcdc7  28
Summary *172.30.5.4       172.30.5.6       0x80000002   958  0x22 0xb9d7  28
Summary  172.30.5.5       172.30.5.3       0x80000001  1481  0x22 0xc3d0  28
Summary *172.30.5.5       172.30.5.6       0x80000001   958  0x22 0xb1df  28

    OSPF database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.3       172.30.5.3       0x80000006   466  0x20 0x71cc  72
Router   172.30.5.4       172.30.5.4       0x80000005   468  0x20 0x64ee  84
Router   172.30.5.5       172.30.5.5       0x80000005   471  0x20 0xbe81  84
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Router  *172.30.5.6       172.30.5.6       0x80000005   471  0x20 0x62b4  72
Summary  172.30.0.0       172.30.5.3       0x80000002  1333  0x20 0x4160  28
Summary *172.30.0.0       172.30.5.6       0x80000002   556  0x20 0x2f6f  28
Summary  172.30.0.4       172.30.5.3       0x80000003   786  0x20 0x1785  28
Summary *172.30.0.4       172.30.5.6       0x80000002   770  0x20 0x793   28
Summary  172.30.0.8       172.30.5.3       0x80000004   578  0x20 0xecaa  28
Summary *172.30.0.8       172.30.5.6       0x80000003   570  0x20 0xdcb8  28
Summary  172.30.0.12      172.30.5.3       0x80000003   239  0x20 0xbcd8  28
Summary *172.30.0.12      172.30.5.6       0x80000002   439  0x20 0xb6db  28
Summary  172.30.0.16      172.30.5.3       0x80000002   870  0x20 0xa0f0  28
Summary *172.30.0.16      172.30.5.6       0x80000003    92  0x20 0x820c  28
Summary  172.30.0.20      172.30.5.3       0x80000003   578  0x20 0x7616  28
Summary *172.30.0.20      172.30.5.6       0x80000002   963  0x20 0x5c2f  28
Summary  172.30.0.28      172.30.5.3       0x80000003   124  0x20 0x1c69  28
Summary *172.30.0.28      172.30.5.6       0x80000002   770  0x20 0x166c  28
Summary  172.30.0.32      172.30.5.3       0x80000003     9  0x20 0xf38d  28
Summary *172.30.0.32      172.30.5.6       0x80000003   570  0x20 0xeb91  28
Summary  172.30.0.44      172.30.5.3       0x80000002   786  0x20 0x87ed  28
Summary *172.30.0.44      172.30.5.6       0x80000002   963  0x20 0x6b08  28
Summary  172.30.0.64      172.30.5.3       0x80000003   578  0x20 0xc698  28
Summary *172.30.0.64      172.30.5.6       0x80000003   570  0x20 0xb4a7  28
Summary  172.30.0.68      172.30.5.3       0x80000002   607  0x20 0x96c6  28
Summary *172.30.0.68      172.30.5.6       0x80000002   322  0x20 0x84d5  28
Summary  172.30.5.1       172.30.5.3       0x80000003   578  0x20 0x1087  28
Summary *172.30.5.1       172.30.5.6       0x80000003   570  0x20 0xfd96  28
Summary  172.30.5.2       172.30.5.3       0x80000002   479  0x20 0xfd9a  28
Summary *172.30.5.2       172.30.5.6       0x80000002   207  0x20 0xeba9  28
Summary  172.30.5.3       172.30.5.3       0x80000002   359  0x20 0xe9ae  28
Summary *172.30.5.3       172.30.5.6       0x80000003   570  0x20 0xe9a8  28
Summary  172.30.5.6       172.30.5.3       0x80000003   572  0x20 0xddb4  28
Summary *172.30.5.6       172.30.5.6       0x80000002   963  0x20 0xb9d8  28
Summary  172.30.5.7       172.30.5.3       0x80000001   786  0x20 0xcdc6  28
Summary *172.30.5.7       172.30.5.6       0x80000002   770  0x20 0xb9d6  28
Summary  172.30.5.8       172.30.5.3       0x80000002   578  0x20 0xc1d0  28
Summary *172.30.5.8       172.30.5.6       0x80000002   570  0x20 0xafdf  28
Summary  172.30.5.41      172.30.5.3       0x80000002   749  0x20 0x80ef  28
Summary *172.30.5.41      172.30.5.6       0x80000002   747  0x20 0x6efe  28
```

```
[edit]
lab@R6# run show ospf3 database | no-more

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.0          172.30.5.1       0x80000004  490  0x4a2e  72
Router      0.0.0.0          172.30.5.2       0x80000005  493  0x7bff  72
Router      0.0.0.0          172.30.5.3       0x80000006  497  0x5019  72
Router     *0.0.0.0          172.30.5.6       0x80000006  503  0xa4be  72
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Router      0.0.0.0          172.30.5.7       0x80000006   511  0x165a  72
Router      0.0.0.0          172.30.5.8       0x80000005   508  0x78f4  72
InterArPfx  0.0.0.1          172.30.5.3       0x80000001  1564  0x449   44
InterArPfx  0.0.0.2          172.30.5.3       0x80000001  1514  0x1437  44
InterArPfx *0.0.0.1          172.30.5.6       0x80000002   985  0xef59  44
InterArPfx *0.0.0.2          172.30.5.6       0x80000001   995  0x246   44
IntraArPfx  0.0.0.1          172.30.5.1       0x80000003   490  0xb29b  52
IntraArPfx  0.0.0.1          172.30.5.2       0x80000004   493  0xda6f  52
IntraArPfx  0.0.0.1          172.30.5.3       0x80000003   497  0x741   52
IntraArPfx *0.0.0.1          172.30.5.6       0x80000003   503  0x85b9  52
IntraArPfx  0.0.0.1          172.30.5.7       0x80000003   511  0xaf8c  52
IntraArPfx  0.0.0.1          172.30.5.8       0x80000003   508  0xd95f  52

    OSPF3 database, Area 0.0.0.1
 Type       ID               Adv Rtr          Seq          Age  Cksum  Len
Router      0.0.0.0          172.30.5.3       0x80000004   498  0x7bd6  56
Router      0.0.0.0          172.30.5.4       0x80000005   500  0x1a3d  56
Router      0.0.0.0          172.30.5.5       0x80000005   503  0x1245  56
Router     *0.0.0.0          172.30.5.6       0x80000004   503  0x9abd  56
InterArPfx  0.0.0.1          172.30.5.3       0x80000003   613  0xb795  44
InterArPfx  0.0.0.2          172.30.5.3       0x80000001  1576  0xc588  44
InterArPfx  0.0.0.3          172.30.5.3       0x80000001  1576  0xcf7d  44
InterArPfx  0.0.0.4          172.30.5.3       0x80000003   605  0x1c29  44
InterArPfx  0.0.0.5          172.30.5.3       0x80000001   816  0x2a1c  44
InterArPfx  0.0.0.6          172.30.5.3       0x80000001   613  0x3a0a  44
InterArPfx *0.0.0.1          172.30.5.6       0x80000001  1000  0x2029  44
InterArPfx *0.0.0.2          172.30.5.6       0x80000003   603  0x9bad  44
InterArPfx *0.0.0.3          172.30.5.6       0x80000001   995  0xa9a0  44
InterArPfx *0.0.0.4          172.30.5.6       0x80000003   603  0xbb89  44
InterArPfx *0.0.0.5          172.30.5.6       0x80000002   809  0x162c  44
InterArPfx *0.0.0.6          172.30.5.6       0x80000002   603  0x261a  44
IntraArPfx  0.0.0.1          172.30.5.4       0x80000003   500  0x3114  52
IntraArPfx  0.0.0.1          172.30.5.5       0x80000003   503  0x5be6  52

    OSPF3 Link-Local database, interface ae0.0 Area 0.0.0.1
 Type       ID               Adv Rtr          Seq          Age  Cksum  Len
Link        0.0.0.2          172.30.5.5       0x80000002   671  0x178c  44
Link       *0.0.0.2          172.30.5.6       0x80000002   524  0x53a7  44

    OSPF3 Link-Local database, interface ge-0/0/4.126 Area 0.0.0.0
 Type       ID               Adv Rtr          Seq          Age  Cksum  Len
Link        0.0.0.4          172.30.5.2       0x80000002  1183  0xc55b  44
Link       *0.0.0.3          172.30.5.6       0x80000002   343  0x19db  44

    OSPF3 Link-Local database, interface ge-0/0/4.146 Area 0.0.0.1
 Type       ID               Adv Rtr          Seq          Age  Cksum  Len
Link        0.0.0.3          172.30.5.4       0x80000002   249  0xb54   44
Link       *0.0.0.4          172.30.5.6       0x80000002   165  0x8061  44
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    OSPF3 Link-Local database, interface ge-0/0/4.167 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.5          172.30.5.6        0x80000001  1005  0x15b5  44
Link        0.0.0.4          172.30.5.7        0x80000001   822  0x972d  44

    OSPF3 Link-Local database, interface ge-0/0/4.168 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.6          172.30.5.6        0x80000001  1005  0x1eaa  44
Link        0.0.0.4          172.30.5.8        0x80000001   619  0x5bea  44

    OSPF3 Link-Local database, interface lo0.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          172.30.5.6        0x80000002   965  0x4726  44
```

g) R7

```
[edit]
lab@R7# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.5       ge-0/0/4.117           Full      172.30.5.1       128    38
172.30.0.29      ge-0/0/4.137           Full      172.30.5.3       128    36
172.30.0.45      ge-0/0/4.167           Full      172.30.5.6       128    36
```

```
[edit]
lab@R7# run show ospf interface
Interface          State   Area       DR ID           BDR ID          Nbrs
ge-0/0/4.117       PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
ge-0/0/4.137       PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
ge-0/0/4.167       PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
lo0.0              DR      0.0.0.0    172.30.5.7      0.0.0.0            0
```

```
[edit]
lab@R7# run show ospf3 neighbor
ID                 Interface              State     Pri   Dead
172.30.5.1         ge-0/0/4.117           Full      128     35
  Neighbor-address fe80::226:8800:75ef:7584
172.30.5.3         ge-0/0/4.137           Full      128     32
  Neighbor-address fe80::2e21:7200:89cd:2684
172.30.5.6         ge-0/0/4.167           Full      128     37
  Neighbor-address fe80::fac0:100:a7dc:2c04
```

```
[edit]
lab@R7# run show ospf3 interface
Interface          State   Area       DR ID           BDR ID          Nbrs
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ge-0/0/4.117    PtToPt  0.0.0.0    0.0.0.0          0.0.0.0             1
ge-0/0/4.137    PtToPt  0.0.0.0    0.0.0.0          0.0.0.0             1
ge-0/0/4.167    PtToPt  0.0.0.0    0.0.0.0          0.0.0.0             1
lo0.0           DR      0.0.0.0    172.30.5.7       0.0.0.0             0
```

Verify the OSPFv2 and OSPFv3 databases.

```
[edit]
lab@R7# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.1       172.30.5.1      0x80000008   534  0x22 0xfc9  132
Router   172.30.5.2       172.30.5.2      0x80000006   539  0x22 0xea7  132
Router   172.30.5.3       172.30.5.3      0x80000006   541  0x22 0x7312 108
Router   172.30.5.6       172.30.5.6      0x80000006   549  0x22 0xd992 108
Router  *172.30.5.7       172.30.5.7      0x80000005   555  0x22 0xe792 108
Router   172.30.5.8       172.30.5.8      0x80000006   554  0x22 0x9df8 108
Router   172.30.5.41      172.30.5.41     0x80000004   798  0x22 0x80b3  84
Summary  172.30.0.80      172.30.5.3      0x80000002  1864  0x22 0xf55a  28
Summary  172.30.0.80      172.30.5.6      0x80000002  1037  0x22 0xed5e  28
Summary  172.30.0.84      172.30.5.3      0x80000002  1864  0x22 0xcd7e  28
Summary  172.30.0.84      172.30.5.6      0x80000001  1037  0x22 0xc781  28
Summary  172.30.0.88      172.30.5.3      0x80000001  1607  0x22 0xb196  28
Summary  172.30.0.88      172.30.5.6      0x80000002  1043  0x22 0x93b1  28
Summary  172.30.0.92      172.30.5.3      0x80000001  1556  0x22 0x89ba  28
Summary  172.30.0.92      172.30.5.6      0x80000002  1043  0x22 0x6bd5  28
Summary  172.30.5.4       172.30.5.3      0x80000001  1607  0x22 0xcdc7  28
Summary  172.30.5.4       172.30.5.6      0x80000002  1037  0x22 0xb9d7  28
Summary  172.30.5.5       172.30.5.3      0x80000001  1556  0x22 0xc3d0  28
Summary  172.30.5.5       172.30.5.6      0x80000001  1037  0x22 0xb1df  28
```

```
[edit]
lab@R7# run show ospf3 database | no-more

    OSPF3 database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq        Age  Cksum  Len
Router      0.0.0.0          172.30.5.1      0x80000004   568  0x4a2e  72
Router      0.0.0.0          172.30.5.2      0x80000005   573  0x7bff  72
Router      0.0.0.0          172.30.5.3      0x80000006   575  0x5019  72
Router      0.0.0.0          172.30.5.6      0x80000006   583  0xa4be  72
Router     *0.0.0.0          172.30.5.7      0x80000006   589  0x165a  72
Router      0.0.0.0          172.30.5.8      0x80000005   588  0x78f4  72
InterArPfx  0.0.0.1          172.30.5.3      0x80000001  1641  0x449   44
InterArPfx  0.0.0.2          172.30.5.3      0x80000001  1591  0x1437  44
InterArPfx  0.0.0.1          172.30.5.6      0x80000002  1066  0xef59  44
InterArPfx  0.0.0.2          172.30.5.6      0x80000001  1076  0x246   44
IntraArPfx  0.0.0.1          172.30.5.1      0x80000003   568  0xb29b  52
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
IntraArPfx  0.0.0.1          172.30.5.2       0x80000004   573  0xda6f  52
IntraArPfx  0.0.0.1          172.30.5.3       0x80000003   575  0x741   52
IntraArPfx  0.0.0.1          172.30.5.6       0x80000003   583  0x85b9  52
IntraArPfx *0.0.0.1          172.30.5.7       0x80000003   589  0xaf8c  52
IntraArPfx  0.0.0.1          172.30.5.8       0x80000003   588  0xd95f  52

    OSPF3 Link-Local database, interface ge-0/0/4.117 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.3          172.30.5.1       0x80000002  1783  0xd65b  44
Link       *0.0.0.2          172.30.5.7       0x80000002   144  0xef08  44

    OSPF3 Link-Local database, interface ge-0/0/4.137 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.5          172.30.5.3       0x80000002   938  0xd6a3  44
Link       *0.0.0.3          172.30.5.7       0x80000001   900  0x657e  44

    OSPF3 Link-Local database, interface ge-0/0/4.167 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.5          172.30.5.6       0x80000002    68  0x13b6  44
Link       *0.0.0.4          172.30.5.7       0x80000001   900  0x972d  44

    OSPF3 Link-Local database, interface lo0.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          172.30.5.7       0x80000002   860  0xbfa6  44
```

h) R8.

```
[edit]
lab@R8# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.30.0.9       ge-0/0/4.118           Full      172.30.5.1       128    37
172.30.0.33      ge-0/0/4.138           Full      172.30.5.3       128    39
172.30.0.21      ge-0/0/4.168           Full      172.30.5.6       128    31
```

```
[edit]
lab@R8# run show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/4.118       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
ge-0/0/4.138       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
ge-0/0/4.168       PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1
lo0.0              DR      0.0.0.0         172.30.5.8      0.0.0.0            0
```

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# run show ospf3 neighbor
ID              Interface              State     Pri   Dead
172.30.5.1      ge-0/0/4.118           Full      128     38
  Neighbor-address fe80::226:8800:76ef:7584
172.30.5.3      ge-0/0/4.138           Full      128     36
  Neighbor-address fe80::2e21:7200:8acd:2684
172.30.5.6      ge-0/0/4.168           Full      128     38
  Neighbor-address fe80::fac0:100:a8dc:2c04
```

```
[edit]
lab@R8# run show ospf3 interface
Interface       State   Area       DR ID           BDR ID          Nbrs
ge-0/0/4.118    PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
ge-0/0/4.138    PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
ge-0/0/4.168    PtToPt  0.0.0.0    0.0.0.0         0.0.0.0            1
lo0.0           DR      0.0.0.0    172.30.5.8      0.0.0.0            0
```

Verify the OSPFv2 and OSPFv3 databases.

```
[edit]
lab@R8# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.30.5.1       172.30.5.1       0x80000008   573  0x22 0xfc9  132
Router   172.30.5.2       172.30.5.2       0x80000006   578  0x22 0xea7  132
Router   172.30.5.3       172.30.5.3       0x80000006   580  0x22 0x7312 108
Router   172.30.5.6       172.30.5.6       0x80000006   588  0x22 0xd992 108
Router   172.30.5.7       172.30.5.7       0x80000005   595  0x22 0xe792 108
Router  *172.30.5.8       172.30.5.8       0x80000006   591  0x22 0x9df8 108
Router   172.30.5.41      172.30.5.41      0x80000004   836  0x22 0x80b3  84
Summary  172.30.0.80      172.30.5.3       0x80000003     8  0x22 0xf35b  28
Summary  172.30.0.80      172.30.5.6       0x80000002  1076  0x22 0xed5e  28
Summary  172.30.0.84      172.30.5.3       0x80000002  1905  0x22 0xcd7e  28
Summary  172.30.0.84      172.30.5.6       0x80000001  1076  0x22 0xc781  28
Summary  172.30.0.88      172.30.5.3       0x80000001  1648  0x22 0xb196  28
Summary  172.30.0.88      172.30.5.6       0x80000002  1081  0x22 0x93b1  28
Summary  172.30.0.92      172.30.5.3       0x80000001  1597  0x22 0x89ba  28
Summary  172.30.0.92      172.30.5.6       0x80000002  1081  0x22 0x6bd5  28
Summary  172.30.5.4       172.30.5.3       0x80000001  1648  0x22 0xcdc7  28
Summary  172.30.5.4       172.30.5.6       0x80000002  1076  0x22 0xb9d7  28
Summary  172.30.5.5       172.30.5.3       0x80000001  1597  0x22 0xc3d0  28
Summary  172.30.5.5       172.30.5.6       0x80000001  1076  0x22 0xb1df  28
```

```
[edit]
lab@R8# run show ospf3 database | no-more
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    OSPF3 database, Area 0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.0          172.30.5.1      0x80000004   590  0x4a2e  72
Router      0.0.0.0          172.30.5.2      0x80000005   595  0x7bff  72
Router      0.0.0.0          172.30.5.3      0x80000006   597  0x5019  72
Router      0.0.0.0          172.30.5.6      0x80000006   605  0xa4be  72
Router      0.0.0.0          172.30.5.7      0x80000006   612  0x165a  72
Router     *0.0.0.0          172.30.5.8      0x80000005   608  0x78f4  72
InterArPfx  0.0.0.1          172.30.5.3      0x80000001  1665  0x449   44
InterArPfx  0.0.0.2          172.30.5.3      0x80000001  1615  0x1437  44
InterArPfx  0.0.0.1          172.30.5.6      0x80000002  1088  0xef59  44
InterArPfx  0.0.0.2          172.30.5.6      0x80000001  1098  0x246   44
IntraArPfx  0.0.0.1          172.30.5.1      0x80000003   590  0xb29b  52
IntraArPfx  0.0.0.1          172.30.5.2      0x80000004   595  0xda6f  52
IntraArPfx  0.0.0.1          172.30.5.3      0x80000003   597  0x741   52
IntraArPfx  0.0.0.1          172.30.5.6      0x80000003   605  0x85b9  52
IntraArPfx  0.0.0.1          172.30.5.7      0x80000003   612  0xaf8c  52
IntraArPfx *0.0.0.1          172.30.5.8      0x80000003   608  0xd95f  52

    OSPF3 Link-Local database, interface ge-0/0/4.118 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.4          172.30.5.1      0x80000002  1321  0xdf50  44
Link       *0.0.0.2          172.30.5.8      0x80000001   719  0xb5c4  44

    OSPF3 Link-Local database, interface ge-0/0/4.138 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.6          172.30.5.3      0x80000002   756  0xdf98  44
Link       *0.0.0.3          172.30.5.8      0x80000001   719  0x293c  44

    OSPF3 Link-Local database, interface ge-0/0/4.168 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.6          172.30.5.6      0x80000001  1107  0x1eaa  44
Link       *0.0.0.4          172.30.5.8      0x80000001   719  0x5bea  44

    OSPF3 Link-Local database, interface lo0.0 Area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.1          172.30.5.8      0x80000002   679  0x7078  44
```

Verify connectivity on R1.

```
[edit]
lab@R1# run ping 172.30.5.2 rapid
PING 172.30.5.2 (172.30.5.2): 56 data bytes
!!!!!
--- 172.30.5.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.182/1.915/2.649/0.734 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run ping 172.30.5.3 rapid
PING 172.30.5.3 (172.30.5.3): 56 data bytes
!!!!!
--- 172.30.5.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.542/3.944/7.063/1.774 ms
```

```
[edit]
lab@R1# run ping 172.30.5.4 rapid
PING 172.30.5.4 (172.30.5.4): 56 data bytes
!!!!!
--- 172.30.5.4 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.512/2.909/7.984/2.540 ms
```

```
[edit]
lab@R1# run ping 172.30.5.5 rapid
PING 172.30.5.5 (172.30.5.5): 56 data bytes
!!!!!
--- 172.30.5.5 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.561/2.910/7.882/2.487 ms
```

```
[edit]
lab@R1# run ping 172.30.5.6 rapid
PING 172.30.5.6 (172.30.5.6): 56 data bytes
!!!!!
--- 172.30.5.6 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.304/3.807/7.716/3.016 ms
```

```
[edit]
lab@R1# run ping 172.30.5.7 rapid
PING 172.30.5.7 (172.30.5.7): 56 data bytes
!!!!!
--- 172.30.5.7 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.929/1.241/1.858/0.324 ms
```

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# run ping 172.30.5.8 rapid
PING 172.30.5.8 (172.30.5.8): 56 data bytes
!!!!!
--- 172.30.5.8 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.065/2.463/7.727/2.633 ms
```

```
[edit]
lab@R1# run ping 172.30.5.41 rapid
PING 172.30.5.41 (172.30.5.41): 56 data bytes
!!!!!
--- 172.30.5.41 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.249/2.466/7.073/2.305 ms
```

All IPv4 loopback IP addresses are reachable.

```
[edit]
lab@R1# run ping fd17:f0f4:f691:5::2 rapid
PING6(56=40+8+8 bytes) fd17:f0f4:f691:5::1 --> fd17:f0f4:f691:5::2
!!!!!
--- fd17:f0f4:f691:5::2 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 1.928/4.974/7.757/2.142 ms
```

```
[edit]
lab@R1# run ping fd17:f0f4:f691:5::3 rapid
PING6(56=40+8+8 bytes) fd17:f0f4:f691:5::1 --> fd17:f0f4:f691:5::3
!!!!!
--- fd17:f0f4:f691:5::3 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 1.973/2.258/2.819/0.309 ms
```

```
[edit]
lab@R1# run ping fd17:f0f4:f691:5::4 rapid
PING6(56=40+8+8 bytes) fd17:f0f4:f691:5::1 --> fd17:f0f4:f691:5::4
!!!!!
--- fd17:f0f4:f691:5::4 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 2.214/4.531/7.686/2.535 ms
```

```
[edit]
lab@R1# run ping fd17:f0f4:f691:5::5 rapid
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
PING6(56=40+8+8 bytes) fd17:f0f4:f691:5::1 --> fd17:f0f4:f691:5::5
!!!!!
--- fd17:f0f4:f691:5::5 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 2.244/3.660/8.207/2.313 ms
```

```
[edit]
lab@R1# run ping fd17:f0f4:f691:5::6 rapid
PING6(56=40+8+8 bytes) fd17:f0f4:f691:5::1 --> fd17:f0f4:f691:5::6
!!!!!
--- fd17:f0f4:f691:5::6 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 1.870/5.622/14.014/4.704 ms
```

```
[edit]
lab@R1# run ping fd17:f0f4:f691:5::7 rapid
PING6(56=40+8+8 bytes) fd17:f0f4:f691:5::1 --> fd17:f0f4:f691:5::7
!!!!!
--- fd17:f0f4:f691:5::7 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 1.878/4.378/13.151/4.390 ms
```

```
[edit]
lab@R1# run ping fd17:f0f4:f691:5::8 rapid
PING6(56=40+8+8 bytes) fd17:f0f4:f691:5::1 --> fd17:f0f4:f691:5::8
!!!!!
--- fd17:f0f4:f691:5::8 ping6 statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 1.819/12.025/45.422/16.824 ms
```

All IPv6 loopback IP addresses are also reachable.

Verify routing table on R5 in role of ABR for loopback IP addresses.

```
[edit]
lab@R5# run show route protocol ospf

inet.0: 34 destinations, 34 routes (34 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.0.0/30      *[OSPF/10] 00:00:12, metric 20
                    > to 172.30.0.94 via ae0.0
172.30.0.4/30      *[OSPF/10] 00:00:11, metric 25
                    > to 172.30.0.94 via ae0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.0.8/30      *[OSPF/10] 00:00:06, metric 25
                    > to 172.30.0.94 via ae0.0
172.30.0.12/30     *[OSPF/10] 00:00:12, metric 20
                    > to 172.30.0.85 via ge-0/0/4.135
172.30.0.16/30     *[OSPF/10] 00:00:12, metric 15
                    > to 172.30.0.94 via ae0.0
172.30.0.20/30     *[OSPF/10] 00:00:12, metric 15
                    > to 172.30.0.94 via ae0.0
172.30.0.28/30     *[OSPF/10] 00:00:11, metric 20
                    > to 172.30.0.85 via ge-0/0/4.135
172.30.0.32/30     *[OSPF/10] 00:00:06, metric 20
                    > to 172.30.0.85 via ge-0/0/4.135
172.30.0.44/30     *[OSPF/10] 00:00:12, metric 15
                    > to 172.30.0.94 via ae0.0
172.30.0.64/30     *[OSPF/10] 00:00:06, metric 30
                    > to 172.30.0.94 via ae0.0
172.30.0.68/30     *[OSPF/10] 00:00:12, metric 25
                    > to 172.30.0.94 via ae0.0
172.30.0.80/30     *[OSPF/10] 00:00:12, metric 15
                    > to 172.30.0.85 via ge-0/0/4.135
172.30.0.88/30     *[OSPF/10] 00:00:12, metric 15
                    > to 172.30.0.94 via ae0.0
172.30.5.1/32      *[OSPF/10] 00:00:06, metric 20
                    > to 172.30.0.94 via ae0.0
172.30.5.2/32      *[OSPF/10] 00:00:12, metric 15
                    > to 172.30.0.94 via ae0.0
172.30.5.3/32      *[OSPF/10] 00:00:12, metric 10
                    > to 172.30.0.85 via ge-0/0/4.135
172.30.5.4/32      *[OSPF/10] 00:00:12, metric 15
                      to 172.30.0.94 via ae0.0
                    > to 172.30.0.85 via ge-0/0/4.135
172.30.5.6/32      *[OSPF/10] 00:00:12, metric 5
                    > to 172.30.0.94 via ae0.0
172.30.5.7/32      *[OSPF/10] 00:00:12, metric 15
                    > to 172.30.0.94 via ae0.0
172.30.5.8/32      *[OSPF/10] 00:00:12, metric 15
                    > to 172.30.0.94 via ae0.0
172.30.5.41/32     *[OSPF/10] 00:00:12, metric 25
                    > to 172.30.0.94 via ae0.0
224.0.0.5/32       *[OSPF/10] 00:34:40, metric 1



inet6.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691:5::1/128
                   *[OSPF3/10] 00:00:10, metric 20
                    > to fe80::fac0:1ff:fedc:2c80 via ae0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
fd17:f0f4:f691:5::2/128
                   *[OSPF3/10] 00:00:10, metric 15
                    > to fe80::fac0:1ff:fedc:2c80 via ae0.0
fd17:f0f4:f691:5::3/128
                   *[OSPF3/10] 00:00:10, metric 10
                    > to fe80::2e21:7200:87cd:2684 via ge-0/0/4.135
fd17:f0f4:f691:5::4/128
                   *[OSPF3/10] 00:00:10, metric 15
                      to fe80::fac0:1ff:fedc:2c80 via ae0.0
                    > to fe80::2e21:7200:87cd:2684 via ge-0/0/4.135
fd17:f0f4:f691:5::6/128
                   *[OSPF3/10] 00:00:15, metric 5
                    > to fe80::fac0:1ff:fedc:2c80 via ae0.0
fd17:f0f4:f691:5::7/128
                   *[OSPF3/10] 00:00:10, metric 15
                    > to fe80::fac0:1ff:fedc:2c80 via ae0.0
fd17:f0f4:f691:5::8/128
                   *[OSPF3/10] 00:00:10, metric 15
                    > to fe80::fac0:1ff:fedc:2c80 via ae0.0
ff02::5/128        *[OSPF3/10] 00:34:40, metric 1
                      MultiRecv
```

All IPv4 and IPv6 routes from Area 0 are received including the IP address of the route reflector router. This ensures optimal routing for the routers in Area 1.

2) Summarize the Area 1 OSPF internal IPv4 routes to the backbone area.

You have to configure route summarization of the Area 1 internal routes towards Area 0. The Inter Area summarization is done on ASBR routers.

Configure R3 for Area 1 summarization.

```
[edit]
lab@R3# set protocols ospf area 0.0.0.1 area-range 172.30.0.80/28

[edit]
lab@R3# commit
commit complete
```

The **area-range** command is configured under **area** hierarchy, not under **nssa**.

Summarization should be configured on R6 also to provide redundancy if R3 fails.

```
[edit]
lab@R6# set protocols ospf area 0.0.0.1 area-range 172.30.0.80/28

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R6# commit
commit complete
```

Verify the summarization of the Area 1 routes on R1.

```
[edit]
lab@R1# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.1       172.30.5.1       0x80000009  2448  0x22 0xd8c  132
Router   172.30.5.2       172.30.5.2       0x80000007  2405  0x22 0xc6a  132
Router   172.30.5.3       172.30.5.3       0x8000000e    37  0x22 0x8db9 108
Router   172.30.5.6       172.30.5.6       0x8000000a    40  0x22 0xfb36 108
Router   172.30.5.7       172.30.5.7       0x80000006  2397  0x22 0x1033 108
Router   172.30.5.8       172.30.5.8       0x80000007  2396  0x22 0xc599 108
Router   172.30.5.41      172.30.5.41      0x80000005   699  0x22 0x7eb4  84
Summary  172.30.0.80      172.30.5.3       0x80000008    37  0x22 0x2e1a  28
Summary  172.30.0.80      172.30.5.6       0x80000005    40  0x22 0x2226  28
Summary  172.30.5.4       172.30.5.3       0x80000006   543  0x22 0xeba0  28
Summary  172.30.5.4       172.30.5.6       0x80000004  1095  0x22 0x1076  28
Summary  172.30.5.5       172.30.5.3       0x80000005   543  0x22 0x1671  28
Summary  172.30.5.5       172.30.5.6       0x80000003   986  0x22 0xd5b5  28
```

Both ASBRs are advertising a summary route instead of separate routes of the internal routes of Area 1. The loopback IP addresses of R4 and R5 are still advertised because they are not part of summary route.

3) Enable RIP on R5 and R6 towards DC1.

R5 and R6 have RIP connectivity to DC1 router.

Configure the RIP neighbor on R5.

```
[edit]
lab@R5# set protocols rip group dc1 neighbor ge-0/0/4.202

[edit]
lab@R5# commit
commit complete
```

Configure RIP neighbor on R6.

```
[edit]
lab@R6# set protocols rip group dc1 neighbor ge-0/0/4.204

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R6# commit
commit complete
```

Verify the RIP neighbors and if routes are received.

```
[edit]
lab@R5# run show rip neighbor
                Local  Source           Destination      Send   Receive   In
Neighbor        State  Address          Address          Mode   Mode      Met
--------        -----  -------          -----------      ----   -------   ---
ge-0/0/4.202    Up     172.30.0.49      224.0.0.9        mcast  both       1
```

```
[edit]
lab@R5# run show route receive-protocol rip 172.30.0.50

inet.0: 51 destinations, 51 routes (51 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.33.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.34.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.35.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.36.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.37.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.38.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.39.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.40.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.41.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.42.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.43.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.44.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.45.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                    > to 172.30.0.50 via ge-0/0/4.202
172.30.46.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                     > to 172.30.0.50 via ge-0/0/4.202
172.30.47.0/24     *[RIP/100] 00:01:05, metric 2, tag 0
                     > to 172.30.0.50 via ge-0/0/4.202
```

```
[edit]
lab@R6# run show rip neighbor
                Local  Source          Destination      Send   Receive   In
Neighbor        State  Address         Address          Mode   Mode      Met
--------        -----  -------         -----------      ----   -------   ---
ge-0/0/4.204      Up 172.30.0.57      224.0.0.9        mcast  both        1
```

```
[edit]
lab@R6# run show route receive-protocol rip 172.30.0.58

inet.0: 54 destinations, 54 routes (54 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.33.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.34.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.35.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.36.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.37.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.38.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.39.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.40.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.41.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.42.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.43.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.44.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.45.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.46.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                     > to 172.30.0.58 via ge-0/0/4.204
172.30.47.0/24     *[RIP/100] 00:00:33, metric 2, tag 0
```

Junos by default accepts all routes from RIP neighbors but does not advertise any routes. R5 and R6 are receiving the same routes from DC1 via RIP.

4) Redistributing default route to DC1 router from R5 and R6.

R5 has two interfaces to the core network, one to R3 and one to R6. If you create an aggregate default route and advertise it to RIP, it will be advertised also in case that both interfaces to R5 are down. This will lead to the black hole the traffic at R5. It is better if you generate the default route from both ABRs. This way R5 will advertise the default only if it has connectivity to the core.

Generate a Type 7 default route on R3 and R6.

R3:

```
[edit]
lab@R3# set protocols ospf area 0.0.0.1 nssa default-lsa default-metric 10

[edit]
lab@R3# commit
commit complete
```

R6:

```
[edit]
lab@R6# set protocols ospf area 0.0.0.1 nssa default-lsa default-metric 10

[edit]
lab@R6# commit
commit complete
```

This will generate an LSA Type 7 default on both ABRs.

Verify if the default route is present on R5.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# run show route protocol ospf terse

inet.0: 52 destinations, 52 routes (52 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop         AS path
* 0.0.0.0/0          O 150         15             >172.30.0.94
* 172.30.0.0/30      O  10         20             >172.30.0.94
* 172.30.0.4/30      O  10         25             >172.30.0.94
* 172.30.0.8/30      O  10         25             >172.30.0.94
* 172.30.0.12/30     O  10         20             >172.30.0.85
* 172.30.0.16/30     O  10         15             >172.30.0.94
* 172.30.0.20/30     O  10         15             >172.30.0.94
* 172.30.0.28/30     O  10         20             >172.30.0.85
...
```

The default route is present as an external OSPF route received from R6.The Interface towards R6 has a lower cost.

```
[edit]
lab@R6# run show route protocol ospf

inet.0: 55 destinations, 55 routes (55 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.0.0/30      *[OSPF/10] 00:41:13, metric 15
                    > to 172.30.0.17 via ge-0/0/4.126
172.30.0.4/30      *[OSPF/10] 00:41:13, metric 20
                    > to 172.30.0.46 via ge-0/0/4.167
172.30.0.8/30      *[OSPF/10] 00:41:07, metric 20
                    > to 172.30.0.22 via ge-0/0/4.168
172.30.0.12/30     *[OSPF/10] 00:41:13, metric 20
                    > to 172.30.0.17 via ge-0/0/4.126
172.30.0.28/30     *[OSPF/10] 00:41:13, metric 20
                    > to 172.30.0.46 via ge-0/0/4.167
172.30.0.32/30     *[OSPF/10] 00:41:07, metric 20
                    > to 172.30.0.22 via ge-0/0/4.168
...
```

On R6 the default route is not present in the routing table. The reason that the default route is not present on R6 is because of its role in the OSPF network. As an ABR it does not accept LSAs from other ABRs until they are reachable.

You have to create an aggregate default route on R6 to be able to advertise a default route to DC1.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set routing-options aggregate route 0/0

[edit]
lab@R6# commit
commit complete
```

Now you can create policies on R3 and R6 to advertise the default route to DC1. Because DC1 will advertise the default route back to OSPF domain with lower preference of 100, which is the route preference of RIP, you have to make sure that this default is not accepted by R5 and R6. You can use a route tag to mark the routes advertised and filter them based on that tag on input direction.

You have to also make sure that R6 is the preferred gateway to the OSPF network, using higher metric to advertised routes from R5.

```
[edit]
lab@R5# set policy-options policy-statement ospf-to-rip term 1 from route-filter 0/0
exact

[edit]
lab@R5# set policy-options policy-statement ospf-to-rip term 1 then metric 10

[edit]
lab@R5# set policy-options policy-statement ospf-to-rip term 1 then tag 1234

[edit]
lab@R5# set policy-options policy-statement ospf-to-rip term 1 then accept
```

Apply the export policy for the RIP neighbor.

```
[edit]
lab@R5# set protocols rip group dc1 export ospf-to-rip

[edit]
lab@R5# commit
commit complete
```

Configure an export policy on R6 matching the default route and set a lower metric.

```
[edit]
lab@R6# set policy-options policy-statement ospf-to-rip term 1 from route-filter 0/0
exact

[edit]
lab@R6# set policy-options policy-statement ospf-to-rip term 1 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set policy-options policy-statement ospf-to-rip term 1 then tag 1234

[edit]
lab@R6# set policy-options policy-statement ospf-to-rip term 1 then metric 5
```

Apply the policy as export on the RIP neighbor.

```
[edit]
lab@R6# set protocols rip group dc1 export ospf-to-rip

[edit]
lab@R6# commit
commit complete
```

Verify the advertisement of default route to DC1.

```
[edit]
lab@R5# run show route advertising-protocol rip 172.30.0.49

inet.0: 53 destinations, 53 routes (53 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 00:32:29, metric 15, tag 0
                    > to 172.30.0.94 via ae0.0
```

```
[edit]
lab@R6# run show route advertising-protocol rip 172.30.0.57

inet.0: 57 destinations, 74 routes (57 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 00:26:08
                       Reject
```

Depending on how quickly you issue the command R6 may be advertising the statically configured aggregate route. Chances are that the RIP route is being received and takes precedence over the aggregate route, so you will not be advertising it. You need to block the default route being sent by R5.

Configure an import policy for traffic received from the DC1 using the tag applied to the OSPF routes advertised from R5 and R6.

```
[edit]
lab@R5# set policy-options policy-statement rip-filter term 1 from protocol rip
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement rip-filter term 1 from tag 1234

[edit]
lab@R5# set policy-options policy-statement rip-filter term 1 then reject

[edit]
lab@R5# set protocols rip group dc1 import rip-filter

[edit]
lab@R5# commit
commit complete
```

Apply the same policy on R6 also.

```
[edit]
lab@R6# set policy-options policy-statement rip-filter term 1 from protocol rip

[edit]
lab@R6# set policy-options policy-statement rip-filter term 1 from tag 1234

[edit]
lab@R6# set policy-options policy-statement rip-filter term 1 then reject

[edit]
lab@R6# set protocols rip group dc1 import rip-filter

[edit]
lab@R6# commit
commit complete
```

5) Redistribute the RIP routes to OSPF.

You have to redistribute the RIP routes to OSPF domain.

```
[edit]
lab@R5# set policy-option policy-statement rip-to-ospf term 1 from protocol rip

[edit]
lab@R5# set policy-option policy-statement rip-to-ospf term 1 then accept
```

Apply the policy in the export direction on OSPF process.

```
[edit]
lab@R5# set protocols ospf export rip-to-ospf
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# commit
commit complete
```

Verify the result on R1, which is an internal router in Area 0.

```
[edit]
lab@R1# run show ospf database

    OSPF database, Area 0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.1       172.30.5.1       0x8000000a    58  0x22 0xb8d  132
Router   172.30.5.2       172.30.5.2       0x80000008    15  0x22 0xa6b  132
Router   172.30.5.3       172.30.5.3       0x8000000e   647  0x22 0x8db9 108
Router   172.30.5.6       172.30.5.6       0x8000000b   303  0x22 0xf937 108
Router   172.30.5.7       172.30.5.7       0x80000007     7  0x22 0xe34  108
Router   172.30.5.8       172.30.5.8       0x80000008     6  0x22 0xc39a 108
Router   172.30.5.41      172.30.5.41      0x80000005  1309  0x22 0x7eb4  84
Summary  172.30.0.80      172.30.5.3       0x80000008   647  0x22 0x2e1a  28
Summary  172.30.0.80      172.30.5.6       0x80000005   650  0x22 0x2226  28
Summary  172.30.5.4       172.30.5.3       0x80000006  1153  0x22 0xeba0  28
Summary  172.30.5.4       172.30.5.6       0x80000004  1705  0x22 0x1076  28
Summary  172.30.5.5       172.30.5.3       0x80000005  1153  0x22 0x1671  28
Summary  172.30.5.5       172.30.5.6       0x80000003  1596  0x22 0xd5b5  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   172.30.32.0      172.30.5.6       0x80000001    48  0x22 0x9783  36
Extern   172.30.33.0      172.30.5.6       0x80000001    48  0x22 0x8c8d  36
Extern   172.30.34.0      172.30.5.6       0x80000001    48  0x22 0x8197  36
Extern   172.30.35.0      172.30.5.6       0x80000001    48  0x22 0x76a1  36
Extern   172.30.36.0      172.30.5.6       0x80000001    48  0x22 0x6bab  36
Extern   172.30.37.0      172.30.5.6       0x80000001    48  0x22 0x60b5  36
Extern   172.30.38.0      172.30.5.6       0x80000001    48  0x22 0x55bf  36
Extern   172.30.39.0      172.30.5.6       0x80000001    48  0x22 0x4ac9  36
Extern   172.30.40.0      172.30.5.6       0x80000001    48  0x22 0x3fd3  36
Extern   172.30.41.0      172.30.5.6       0x80000001    48  0x22 0x34dd  36
Extern   172.30.42.0      172.30.5.6       0x80000001    48  0x22 0x29e7  36
Extern   172.30.43.0      172.30.5.6       0x80000001    48  0x22 0x1ef1  36
Extern   172.30.44.0      172.30.5.6       0x80000001    48  0x22 0x13fb  36
Extern   172.30.45.0      172.30.5.6       0x80000001    48  0x22 0x806   36
Extern   172.30.46.0      172.30.5.6       0x80000001    48  0x22 0xfc10  36
Extern   172.30.47.0      172.30.5.6       0x80000001    48  0x22 0xf11a  36
```

R1 receives the external RIP routes from R6. It is the selected ABR for translating LSA Type 7 to Type 5 and sending them to Backbone Area 0. This ABR selection is based on a higher Router ID, R3 is suppressing the routes because it receives LSA Type 1 from R6 in Area 0.

© 2020 Juniper Networks, Inc. All rights reserved.

RIP route in R6 are known by RIP and by OSPF.

```
[edit]
lab@R6# run show route 172.30.32.0/20

inet.0: 57 destinations, 74 routes (57 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24     *[RIP/100] 01:02:07, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
                    [OSPF/150] 00:18:50, metric 2, tag 0
                    > to 172.30.0.93 via ae0.0
172.30.33.0/24     *[RIP/100] 01:02:07, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
                    [OSPF/150] 00:18:50, metric 2, tag 0
                    > to 172.30.0.93 via ae0.0
172.30.34.0/24     *[RIP/100] 01:02:07, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
                    [OSPF/150] 00:18:50, metric 2, tag 0
                    > to 172.30.0.93 via ae0.0
172.30.35.0/24     *[RIP/100] 01:02:07, metric 2, tag 0
                    > to 172.30.0.58 via ge-0/0/4.204
                    [OSPF/150] 00:18:50, metric 2, tag 0
                    > to 172.30.0.93 via ae0.0
--- (more) ---
```

If you apply an export policy for advertising all RIP routes you will not be able to summarize the routes to Area 0. This is covered in the next Step.

6) Configure the tightest possible summary route for Area 1 OSPF external IPv4 routes towards the backbone area.

You must summarize all RIP routes and send just one summary route to Area 0. As with summarization of the internal IPv4 routes you can use the area-range command. This time it should be configured under area nssa for summarizing the external RIP routes from the LSA Type 7 to one single LSA Type 5, which suppresses the more specific routes.

```
[edit]
lab@R3# set protocols ospf area 0.0.0.1 nssa area-range 172.30.32.0/20

[edit]
lab@R3# commit
commit complete
```

On R6, you must implement another solution for advertising the summary route because the active routes in the routing table are known by RIP. The RIP routes will not be

© 2020 Juniper Networks, Inc. All rights reserved.

summarized by adding the area-range under area nssa. However, you still must add it to summarizing the RIP routes exported to OSPF case the link to DC1 on R6 fails.

```
[edit]
lab@R6# set protocols ospf area 0.0.0.1 nssa area-range 172.30.32.0/20

[edit]
lab@R6# commit
commit complete
```

Verify the presence of the summary route on R1, and also verify if the more specific routes are suppressed.

```
[edit]
lab@R1# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.1       172.30.5.1      0x8000000a  1381  0x22 0xb8d  132
Router   172.30.5.2       172.30.5.2      0x80000008  1338  0x22 0xa6b  132
Router   172.30.5.3       172.30.5.3      0x8000000f  1329  0x22 0x8bba 108
Router   172.30.5.6       172.30.5.6      0x8000000d   545  0x22 0xf539 108
Router   172.30.5.7       172.30.5.7      0x80000007  1330  0x22 0xe34  108
Router   172.30.5.8       172.30.5.8      0x80000008  1329  0x22 0xc39a 108
Router   172.30.5.41      172.30.5.41     0x80000005  2632  0x22 0x7eb4  84
Summary  172.30.0.80      172.30.5.3      0x80000008  1970  0x22 0x2e1a  28
Summary  172.30.0.80      172.30.5.6      0x80000005  1973  0x22 0x2226  28
Summary  172.30.5.4       172.30.5.3      0x80000007    57  0x22 0xe9a1  28
Summary  172.30.5.4       172.30.5.6      0x80000005   634  0x22 0xe77   28
Summary  172.30.5.5       172.30.5.3      0x80000005  2476  0x22 0x1671  28
Summary  172.30.5.5       172.30.5.6      0x80000004   394  0x22 0xd3b6  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   172.30.32.0       172.30.5.6      0x80000003   545  0x22 0xf506  36
```

Only one external LSA is present on R1, again it is advertised from R6 because of the higher Router id.

To advertise the RIP summary route on R6 to accomplish redundancy requirement while R5 is not available you can create an aggregate route summarizing all RIP routes and advertise it export policy. Aggregate route is active in the routing table as long as the contributing routes are present. This will ensure that the aggregate will not be advertised if the link to DC1 is down.

Configure an aggregate route and routing policy to advertise it to the OSPF domain.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set routing-options aggregate route 172.30.32.0/20

[edit]
lab@R6# set policy-options policy-statement rip-to-ospf term 1 from route-filter
172.30.32.0/20 exact

[edit]
lab@R6# set policy-options policy-statement rip-to-ospf term 1 then accept

[edit]
lab@R6# set protocols ospf export rip-to-ospf

[edit]
lab@R6# commit
commit complete
```

Note: Here the name of the policy is kept the same as configured on R3 for consistency, but basically the policy on R6 should match the aggregate route.

7) No ABR failure should cause the RIP summary route to disappear from OSPF backbone Area.

To simulate ABR failure you can deactivate the **protocol ospf** hierarchy on R6.

```
[edit]
lab@R6# deactivate protocols ospf

[edit]
lab@R6# commit
commit complete
```

Verify the OSPF database on R1 if the summary RIP route is received from R3 in absence of router R6.

```
[edit]
lab@R1# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.30.5.1       172.30.5.1      0x8000000f     1  0x22 0x192  132
Router   172.30.5.2       172.30.5.2      0x8000000b     2  0x22 0x1527 120
Router   172.30.5.3       172.30.5.3      0x8000000f  1385  0x22 0x8bba 108
Router   172.30.5.6       172.30.5.6      0x8000000d   599  0x22 0xf539 108
Router   172.30.5.7       172.30.5.7      0x8000000a     7  0x22 0x9c89  96
Router   172.30.5.8       172.30.5.8      0x8000000b    11  0x22 0x3dec  96
Router   172.30.5.41      172.30.5.41     0x80000007    19  0x22 0x7ab6  84
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Summary  172.30.0.80      172.30.5.3       0x80000008  2026  0x22 0x2e1a  28
Summary  172.30.0.80      172.30.5.6       0x80000005  2029  0x22 0x2226  28
Summary  172.30.5.4       172.30.5.3       0x80000007   112  0x22 0xe9a1  28
Summary  172.30.5.4       172.30.5.6       0x80000005   689  0x22 0xe77   28
Summary  172.30.5.5       172.30.5.3       0x80000005  2532  0x22 0x1671  28
Summary  172.30.5.5       172.30.5.6       0x80000004   449  0x22 0xd3b6  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   172.30.32.0      172.30.5.3       0x80000001    45  0x22 0xcf4   36
Extern   172.30.32.0      172.30.5.6       0x80000003   599  0x22 0xf506  36
```

Note: The LSA from R6 is still shown. You can clear the OSPF database on R1 to wipe out all LSAs. This will force R1 to request database synchronization and the external LSA will not be received.

Enable OSPF on R6

```
[edit]
lab@R6# activate protocols ospf

[edit]
lab@R6# commit
commit complete
```

8) Verify loop free connectivity in the network.

Performing a traceroute from R1 to R4 and R5 to check for loop free connectivity.

```
[edit]
lab@R1# run traceroute 172.30.5.4 no-resolve
traceroute to 172.30.5.4 (172.30.5.4), 30 hops max, 40 byte packets
 1  172.30.0.2  8.693 ms  7.516 ms  6.941 ms
 2  172.30.0.14  13.535 ms  8.926 ms  7.992 ms
 3  172.30.5.4  2.526 ms  7.892 ms  7.918 ms
```

```
[edit]
lab@R1# run traceroute 172.30.5.5 no-resolve
traceroute to 172.30.5.5 (172.30.5.5), 30 hops max, 40 byte packets
 1  172.30.0.2  7.529 ms  8.065 ms  7.902 ms
 2  172.30.0.14  9.163 ms  9.926 ms  7.487 ms
 3  172.30.5.5  8.493 ms  9.225 ms  8.060 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 4: BGP Configuration

In this task, you must configure External BGP sessions with different peers and Customers. First, you must build the Internal BGP network. According to the requirements routing policy should be applied to control the exchange of the routes.

Then, you must configure IBGP and ensure that your configuration follows the requirements. You must use the route reflector with two clusters and only one session from each router to one of the clusters is allowed.

You can split the eight routers in two parts as follows: R1, R6, R7, R8 with sessions to cluster 0.0.0.1 and R2, R3, R4, R5 with sessions to cluster 0.0.0.2.

1) Configure all routers with Internal BGP sessions to the route reflector.

a) R1

The Autonomous system of the network is 54591.

```
[edit]
lab@R1# set routing-options autonomous-system 54591
```

Configure internal BGP groups with name ibgp with the route reflector. Internal BGP groups accept neighbors from the same AS only.

```
[edit]
lab@R1# set protocols bgp group ibgp type internal
```

Configure neighbor changes to be logged by syslog configured under the **system syslog** hierarchy.

```
[edit]
lab@R1# set protocols bgp group ibgp log-updown
```

IBGP sessions are sourced from the loopback IP address.

```
[edit]
lab@R1# set protocols bgp group ibgp local-address 172.30.5.1
```

The RR loopback address is used in the neighbor statement.

```
[edit]
lab@R1# set protocols bgp group ibgp neighbor 172.30.5.41
```

Configure BGP authentication for the session with the route reflector.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols bgp group ibgp authentication-key workbook

[edit]
lab@R1# commit
commit complete
```

b) R2

```
[edit]
lab@R2# set routing-options autonomous-system 54591

[edit]
lab@R2# set protocols bgp group ibgp log-updown

[edit]
lab@R2# set protocols bgp group ibgp type internal

[edit]
lab@R2# set protocols bgp group ibgp local-address 172.30.5.2

[edit]
lab@R2# set protocols bgp group ibgp neighbor 172.30.5.41

[edit]
lab@R2# set protocols bgp group ibgp authentication-key workbook

[edit]
lab@R2# commit
commit complete
```

c) R3

```
[edit]
lab@R3# set routing-options autonomous-system 54591

[edit]
lab@R3# set protocols bgp group ibgp log-updown

[edit]
lab@R3# set protocols bgp group ibgp type internal

[edit]
lab@R3# set protocols bgp group ibgp local-address 172.30.5.3

[edit]
lab@R3# set protocols bgp group ibgp neighbor 172.30.5.41
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols bgp group ibgp authentication-key workbook

[edit]
lab@R3# commit
commit complete
```

R4

```
[edit]
lab@R4# set routing-options autonomous-system 54591

[edit]
lab@R4# set protocols bgp group ibgp log-updown

[edit]
lab@R4# set protocols bgp group ibgp type internal

[edit]
lab@R4# set protocols bgp group ibgp local-address 172.30.5.4

[edit]
lab@R4# set protocols bgp group ibgp neighbor 172.30.5.41

[edit]
lab@R4# set protocols bgp group ibgp authentication-key workbook

[edit]
lab@R4# commit
commit complete
```

d) R5

```
[edit]
lab@R5# set routing-options autonomous-system 54591

[edit]
lab@R5# set protocols bgp group ibgp log-updown

[edit]
lab@R5# set protocols bgp group ibgp type internal

[edit]
lab@R5# set protocols bgp group ibgp local-address 172.30.5.5

[edit]
lab@R5# set protocols bgp group ibgp neighbor 172.30.5.41
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols bgp group ibgp authentication-key workbook

[edit]
lab@R5# commit
commit complete
```

e) R6

```
[edit]
lab@R6# set routing-options autonomous-system 54591

[edit]
lab@R6# set protocols bgp group ibgp log-updown

[edit]
lab@R6# set protocols bgp group ibgp type internal

[edit]
lab@R6# set protocols bgp group ibgp local-address 172.30.5.6

[edit]
lab@R6# set protocols bgp group ibgp neighbor 172.30.5.41

[edit]
lab@R6# set protocols bgp group ibgp authentication-key workbook

[edit]
lab@R6# commit
commit complete
```

f) R7

```
[edit]
lab@R7# set routing-options autonomous-system 54591

[edit]
lab@R7# set protocols bgp group ibgp log-updown

[edit]
lab@R7# set protocols bgp group ibgp type internal

[edit]
lab@R7# set protocols bgp group ibgp local-address 172.30.5.7

[edit]
lab@R7# set protocols bgp group ibgp neighbor 172.30.5.41
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols bgp group ibgp authentication-key workbook

[edit]
lab@R7# commit
commit complete
```

g) R8

```
[edit]
lab@R8# set routing-options autonomous-system 54591

[edit]
lab@R8# set protocols bgp group ibgp log-updown

[edit]
lab@R8# set protocols bgp group ibgp type internal

[edit]
lab@R8# set protocols bgp group ibgp local-address 172.30.5.8

[edit]
lab@R8# set protocols bgp group ibgp neighbor 172.30.5.41

[edit]
lab@R8# set protocols bgp group ibgp authentication-key workbook

[edit]
lab@R8# commit
commit complete
```

h) Route reflector router.

On the route reflector two different groups should be configured, each with four neighbors creating two clusters.

Configure the BGP groups with cluster-ids to reflect the routes between the internal BGP neighbors.

```
[edit]
lab@route-reflector# set routing-options autonomous-system 54591

[edit]
lab@route-reflector# set protocols bgp group cluster-1 type internal

[edit]
lab@route-reflector# set protocols bgp group cluster-1 log-updown
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 local-address 172.30.5.41

[edit]
lab@route-reflector# set protocols bgp group cluster-1 cluster 0.0.0.1
```

Configure routers R1, R6, R7 and R8 as neighbors to group cluster-1.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 172.30.5.1
 [edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 172.30.5.6

[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 172.30.5.7

[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 172.30.5.8
```

Configure BGP authentication for the session with the neighbors in the cluster.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 authentication-key workbook
```

Configure routers R2, R3, R4, R5 as neighbors in group cluster-2.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 type internal

[edit]
lab@route-reflector# set protocols bgp group cluster-2 log-updown

[edit]
lab@route-reflector# set protocols bgp group cluster-2 local-address 172.30.5.41

[edit]
lab@route-reflector# set protocols bgp group cluster-2 cluster 0.0.0.2

[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 172.30.5.2

[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 172.30.5.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 172.30.5.4

[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 172.30.5.5

[edit]
lab@route-reflector# set protocols bgp group cluster-2 authentication-key workbook

[edit]
lab@route-reflector# commit
commit complete
```

Verify the IBGP operation between the routers and the route reflector.

```
[edit]
lab@route-reflector# run show bgp summary
Groups: 2 Peers: 8 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 0          0          0          0          0          0
Peer                 AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
172.30.5.1      54591          8          9       0       0        2:56 Establ
  inet.0: 0/0/0
172.30.5.2      54591          8          9       0       0        2:53 Establ
  inet.0: 0/0/0
172.30.5.3      54591          8          8       0       0        2:45 Establ
  inet.0: 0/0/0
172.30.5.4      54591          8          7       0       0        2:44 Establ
  inet.0: 0/0/0
172.30.5.5      54591          8          7       0       0        2:43 Establ
  inet.0: 0/0/0
172.30.5.6      54591          8          8       0       0        2:42 Establ
  inet.0: 0/0/0
172.30.5.7      54591          7          7       0       0        2:41 Establ
  inet.0: 0/0/0
172.30.5.8      54591          8          7       0       0        2:41 Establ
  inet.0: 0/0/0
```

On the route reflector you can verify the State of all IBGP sessions with a single command. All sessions are established. Still no routes are received by the route reflector. On the other side, for example R1, the session is also Established.

```
[edit]
lab@R1# run show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                        0          0          0          0          0          0
Peer                      AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41            54591          4          6       0       0        1:32
0/0/0/0               0/0/0/0
```

```
[edit]
lab@R1# run show bgp neighbor 172.30.5.41
Peer: 172.30.5.41+179 AS 54591 Local: 172.30.5.1+54736 AS 54591
  Type: Internal    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress AuthKey LogUpDown Refresh>
  Authentication key is configured
  Local Address: 172.30.5.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 172.30.5.41     Local ID: 172.30.5.1        Active Holdtime: 90
  Keepalive Interval: 30         Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  Peer does not support Receiver functionality
  Peer does not support 4 byte AS extension
  Peer does not support Addpath
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              0
    Received prefixes:            0
    Accepted prefixes:            0
    Suppressed due to damping:    0
    Advertised prefixes:          0
  Last traffic (seconds): Received 2    Sent 3    Checked 22
  Input messages:  Total 14     Updates 0       Refreshes 0     Octets 266
  Output messages: Total 16     Updates 0       Refreshes 0     Octets 344
  Output Queue[0]: 0
```

LogUpDown and AuthKey options are enabled.

2) Configure the EBGP sessions according to Table 28.

© 2020 Juniper Networks, Inc. All rights reserved.

According to Table 28 the routers R1, R2, R3, R5 and R8 have external BGP sessions to the Peers and Customers. Configure EBGP sessions using the parameters mentioned in Table 28.

Device	Peer Router	Peer AS	Peer IPv4 Address	Peer IPv6 Address
R1	P1-1	1679.12483	192.168.0.38	fc09:c0:ffee::6
R2	C3-1	64514	192.168.0.10	IPv4 compatible
R3	C2-1	64513	172.31.31.1	
R5	C1-1	64512	192.168.0.18	
R8	P1-2	1679.12483	192.168.0.30	fc09:c0:ffee::2
	P2-1	43208.365	192.168.0.6	

Table 28

a) R1

First you have to configure the interface for the external connectivity to peer P1-1

```
[edit]
lab@R1# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R1# set interfaces ge-0/0/5.310 vlan-id 310

[edit]
lab@R1# set interfaces ge-0/0/5.310 description "P1-1 connection"

[edit]
lab@R1# set interfaces ge-0/0/5.310 family inet address 192.168.0.37/30

[edit]
lab@R1# set interfaces ge-0/0/5.310 family inet6 address fc09:c0:ffee::5/126
```

Configure the IPv4 EBGP session.

```
[edit]
lab@R1# set protocols bgp group P1-1 peer-as 1679.12483
```

Remote AS is a 4-byte and is given in AS-DOT format.

```
[edit]
lab@R1# set protocols bgp group P1-1 neighbor 192.168.0.38
```

Remote direct IP address is configured as neighbor.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols bgp group P1-1 log-updown
```

Neighbor changes are logged to syslog.

```
[edit]
lab@R1# set protocols bgp group P1-1 hold-time 30
```

You have to lower the BGP keepalives from 30 to 10 seconds. To accomplish this requirement, you have to change the hold-time, which is three times the keepalive time. By default, the hold-time is 90 seconds, which means that the keepalives are sent each 30 seconds. Lowering the hold-time to 30 seconds lowers the keepalive to 10 seconds.

```
[edit]
lab@R1# commit
commit complete
```

Configure IPv6 in a separate EBGP group.

```
[edit]
lab@R1# set protocols bgp group P1-1-ipv6 hold-time 30

[edit]
lab@R1# set protocols bgp group P1-1-ipv6 peer-as 1679.12483

[edit]
lab@R1# set protocols bgp group P1-1-ipv6 neighbor fc09:c0:ffee::6

[edit]
lab@R1# set protocols bgp group P1-1-ipv6 log-updown

[edit]
lab@R1# commit
commit complete
```

Verify the EBGP operation.

```
[edit]
lab@R1# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     564        564          0          0          0          0
inet6.0
                      16         16          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.41            54591          30        1029          0          0        13:46
0/0/0/0                0/0/0/0
192.168.0.38       110047427         509          24          0          0         3:18
564/564/564/0          0/0/0/0
fc09:c0:ffee::6    110047427           3           3          0          0            3
Establ
  inet6.0: 16/16/16/0
```

Both sessions are Established. R1 is receiving routes from P1-1.

```
[edit]
lab@R1# run show bgp neighbor 192.168.0.38
Peer: 192.168.0.38+179 AS 110047427 Local: 192.168.0.37+50553 AS 54591
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference HoldTime LogUpDown PeerAS Refresh>
  Holdtime: 30 Preference: 170
  Number of flaps: 0
  Peer ID: 172.17.15.1     Local ID: 172.30.5.1        Active Holdtime: 30
  Keepalive Interval: 10          Peer index: 0
  BFD: disabled, down
  Local Interface: ge-0/0/5.310
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 110047427)
  Peer does not support Addpath
  Table inet.0 Bit: 10001
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              564
    Received prefixes:            564
    Accepted prefixes:            564
    Suppressed due to damping:    0
    Advertised prefixes:          0
  Last traffic (seconds): Received 1    Sent 2    Checked 2
  Input messages:  Total 487    Updates 486     Refreshes 0     Octets 34033
  Output messages: Total 3      Updates 0       Refreshes 0     Octets 120
  Output Queue[0]: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

The keepalive and hold-time are reflecting the configured values. You can see that lowering the hold-time leads to lowering three time the keepalive timer.

b) R2

R2 has a single EBGP session to customer C3-1. The IPv4 session is also used to carry the IPv6 routes simultaneously. IPv4-mapped IPv6 addresses are used. This is a newer approach compared to the IPv4 compatible IPv6 addresses used in Chapter 3. Please note that IPv4 compatible addresses are deprecated.

Configure the interface for connectivity towards customer C3-1.

```
[edit]
lab@R2# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R2# set interfaces ge-0/0/5.303 vlan-id 303

[edit]
lab@R2# set interfaces ge-0/0/5.303 description "C3-1 connection"

[edit]
lab@R2# set interfaces ge-0/0/5.303 family inet address 192.168.0.9/30

[edit]
lab@R2# set interfaces ge-0/0/5.303 family inet6 address ::ffff:192.168.0.9/126
```

The format of the IPv4-mapped IPv6 addresses is ::ffff:<IPv4 address> . The address above ::ffff:192.168.0.9/126 is represented as ::ffff:C0A8:9/126 in IPv6 format.

```
[edit]
lab@R2# set system allow-v4mapped-packets
```

By default, Junos does not process IPv4-mapped addresses because of security measures. To enable the processing of these packets you have to explicitly enable it under the **system** hierarchy with the **allow-v4mapped-packets** knob.

Configure the EBGP session with C3-1.

```
[edit]
lab@R2# set protocols bgp group C3-1 peer-as 64514

[edit]
lab@R2# set protocols bgp group C3-1 family inet unicast prefix-limit maximum 20

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# set protocols bgp group C3-1 family inet unicast prefix-limit teardown idle-
timeout 5
```

All EBGP sessions to the customers should not accept more than 20 prefixes. If these sessions receive more than 20 prefixes, they should be torn down for at least 5 minutes.

```
[edit]
lab@R2# set protocols bgp group C3-1 family inet6 unicast prefix-limit maximum 20

[edit]
lab@R2# set protocols bgp group C3-1 family inet6 unicast prefix-limit teardown idle-
timeout 5
```

The prefix limitation should be applied to the IPv6 routes carried over the IPv4 BGP session also.

```
[edit]
lab@R2# set protocols bgp group C3-1 neighbor 192.168.0.10

[edit]
lab@R2# set protocols bgp group C3-1 damping
```

The routes received from C3-1 should be damped using the default settings.

```
[edit]
lab@R2# set protocols bgp group C3-1 log-updown

[edit]
lab@R2# commit
commit complete
```

Verify the operation of the EBGP session on R2 towards customer C3-1.

```
[edit]
lab@R2# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                        16         16          0          0          0          0
inet6.0
                        16         16          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591         49         53       0       0       22:16
0/0/0/0               0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
192.168.0.10            64514          6          4       0       0          46
Establ
  inet.0: 16/16/16/0
  inet6.0: 16/16/16/0
```

Session is Established and routes are received in both families.

```
[edit]
lab@R2# run show route receive-protocol bgp 192.168.0.10

inet.0: 48 destinations, 49 routes (48 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.32.0/24          192.168.0.10                            64514 I
* 172.31.33.0/24          192.168.0.10                            64514 I
* 172.31.34.0/24          192.168.0.10                            64514 I
* 172.31.35.0/24          192.168.0.10                            64514 I
* 172.31.36.0/24          192.168.0.10                            64514 I
* 172.31.37.0/24          192.168.0.10                            64514 I
* 172.31.38.0/24          192.168.0.10                            64514 I
* 172.31.39.0/24          192.168.0.10                            64514 I
* 172.31.40.0/24          192.168.0.10                            64514 I
* 172.31.41.0/24          192.168.0.10                            64514 I
* 172.31.42.0/24          192.168.0.10                            64514 I
* 172.31.43.0/24          192.168.0.10                            64514 I
* 172.31.44.0/24          192.168.0.10                            64514 I
* 172.31.45.0/24          192.168.0.10                            64514 I
* 172.31.46.0/24          192.168.0.10                            64514 I
* 172.31.47.0/24          192.168.0.10                            64514 I

inet6.0: 33 destinations, 37 routes (33 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  fd18:cccc:dddd:2::/80
*                         ::ffff:192.168.0.10                     64514 I
  fd18:cccc:dddd:2:1::/80
*                         ::ffff:192.168.0.10                     64514 I
  fd18:cccc:dddd:2:2::/80
*                         ::ffff:192.168.0.10                     64514 I
  fd18:cccc:dddd:2:3::/80
*                         ::ffff:192.168.0.10                     64514 I
  fd18:cccc:dddd:2:4::/80
*                         ::ffff:192.168.0.10                     64514 I
  fd18:cccc:dddd:2:5::/80
*                         ::ffff:192.168.0.10                     64514 I
  fd18:cccc:dddd:2:6::/80
*                         ::ffff:192.168.0.10                     64514 I
  fd18:cccc:dddd:2:7::/80
*                         ::ffff:192.168.0.10                     64514 I
  fd18:cccc:dddd:2:8::/80
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
*                       ::ffff:192.168.0.10                    64514 I
  fd18:cccc:dddd:2:9::/80
*                       ::ffff:192.168.0.10                    64514 I
...
```

Both `inet-unicast` and `inet6-unicast` families are negotiated over a single session. PrefixLimit and Damping options are enabled for the session.

```
[edit]
lab@R2# run show bgp neighbor 192.168.0.10
Peer: 192.168.0.10+179 AS 64514 Local: 192.168.0.9+55829 AS 54591
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference Damping AddressFamily PeerAS PrefixLimit Refresh>
  Address families configured: inet-unicast inet6-unicast
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 172.31.47.1     Local ID: 172.30.5.2        Active Holdtime: 90
  Keepalive Interval: 30         Peer index: 0
  BFD: disabled, down
  Local Interface: ge-0/0/5.303
  NLRI for restart configured on peer: inet-unicast inet6-unicast
  NLRI advertised by peer: inet-unicast inet6-unicast
  NLRI for this session: inet-unicast inet6-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Peer supports 4 byte AS extension (peer-as 64514)
  Peer does not support Addpath
  Table inet.0 Bit: 10001
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              16
    Received prefixes:            16
    Accepted prefixes:            16
    Suppressed due to damping:    0
    Advertised prefixes:          0
  Table inet6.0 Bit: 20000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              16
    Received prefixes:            16
    Accepted prefixes:            16
    Suppressed due to damping:    0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Advertised prefixes:          0
Last traffic (seconds): Received 8    Sent 26   Checked 81
Input messages:  Total 41     Updates 4       Refreshes 0     Octets 1116
Output messages: Total 38     Updates 0       Refreshes 0     Octets 823
Output Queue[0]: 0
Output Queue[1]: 0
```

c) R3

R3 should have an External BGP session with customer C2-1. Two separate physical interfaces are available for connectivity to C2-1. You have to make sure that a failure of one of those interfaces will not cause the EBGP session to drop. The EBGP session should be established between the loopback interfaces.

Configure the interfaces for connectivity to C2-1.

```
[edit]
lab@R3# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R3# set interfaces ge-0/0/5.306 vlan-id 306

[edit]
lab@R3# set interfaces ge-0/0/5.306 description "C2-1 connection 1"

[edit]
lab@R3# set interfaces ge-0/0/5.306 family inet address 192.168.0.21/30

[edit]
lab@R3# set interfaces ge-0/0/5.307 vlan-id 307

[edit]
lab@R3# set interfaces ge-0/0/5.307 description "C2-1 connection 2"

[edit]
lab@R3# set interfaces ge-0/0/5.307 family inet address 192.168.0.25/30
```

A static route is allowed to be configured for C2-1 loopback reachability. Both direct interface addresses are used as the next-hops of the static route.

```
[edit]
lab@R3# set routing-options static route 172.31.31.1/32 next-hop [ 192.168.0.22
192.168.0.26 ]
```

Configure the EBGP session with neighbor C2-1.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols bgp group C2-1 neighbor 172.31.31.1

[edit]
lab@R3# set protocols bgp group C2-1 local-address 172.30.5.3

[edit]
lab@R3# set protocols bgp group C2-1 multihop
```

The multihop option is used for establishing External BGP sessions more than one hop away.

```
[edit]
lab@R3# set protocols bgp group C2-1 peer-as 64513

[edit]
lab@R3# set protocols bgp group C2-1 family inet unicast prefix-limit maximum 20

[edit]
lab@R3# set protocols bgp group C2-1 family inet unicast prefix-limit teardown idle-
timeout 5

[edit]
lab@R3# set protocols bgp group C2-1 damping

[edit]
lab@R3# set protocols bgp group C2-1 log-updown

[edit]
lab@R3# commit
commit complete
```

Verify if the EBGP session is established.

```
[edit]
lab@R3# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                        16         16          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41   54591         58         64       0       0       27:03 0/0/0/0
0/0/0/0
172.31.31.1   64513          5          6       0       0        1:18 16/16/16/0
0/0/0/0
```

d) R5

© 2020 Juniper Networks, Inc. All rights reserved.

Router R5 has external BGP session to customer C1-1 in neighboring AS 64512.

Configure the physical interface to the neighbor router.

```
[edit]
lab@R5# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R5# set interfaces ge-0/0/5.305 vlan-id 305

[edit]
lab@R5# set interfaces ge-0/0/5.305 description "C1-1 connection"

[edit]
lab@R5# set interfaces ge-0/0/5.305 family inet address 192.168.0.17/30
```

Configure the External BGP session to C1-1.

```
[edit]
lab@R5# set protocols bgp group C1-1 peer-as 64512

[edit]
lab@R5# set protocols bgp group C1-1 neighbor 192.168.0.18

[edit]
lab@R5# set protocols bgp group C1-1 family inet unicast prefix-limit maximum 20

[edit]
lab@R5# set protocols bgp group C1-1 family inet unicast prefix-limit teardown idle-
timeout 5

[edit]
lab@R5# set protocols bgp group C1-1 log-updown

[edit]
lab@R5# set protocols bgp group C1-1 damping
```

The prefixes received from customer C1-1 should be damped as well as from the other customer. However, this damping policy should be more aggressive. You can modify only one damping parameter to accomplish this Step. Modifying the damping parameters is done with damping policies applied to the BGP neighbor. By default, the suppress value is 3000, above this value the route is temporary removed from the routing table. If you lower the suppression to 2000, routes will be damped /temporary removed more aggressively.

```
[edit]
lab@R5# set policy-options damping aggressive suppress 2000
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement damp-aggressive term 1 then damping
aggressive

[edit]
lab@R5# set protocols bgp group C1-1 import damp-aggressive

[edit]
lab@R5# commit
commit complete
```

Verify if the external BGP session to C1-1 is established.

```
[edit]
lab@R5# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      16         16          0          0          0          0
Peer                  AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41    54591         64         72       0       0       29:51 0/0/0/0
0/0/0/0
192.168.0.18   64512          5          3       0       0          11 16/16/16/0
0/0/0/0
```

e) R8

Configure the EBGP session for router R8. It has three external sessions, one IPv4 session to customer P2-1, one IPv4 EBGP session to P1-2 and an IPv6 session to P1-2. Both IPv4 and IPv6 sessions to P1-2 are sharing the same physical interface.

Configure the physical interfaces on R8.

```
[edit]
lab@R8# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R8# set interfaces ge-0/0/5.302 vlan-id 302

[edit]
lab@R8# set interfaces ge-0/0/5.302 description "P2-1 connection"

[edit]
lab@R8# set interfaces ge-0/0/5.302 family inet address 192.168.0.5/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set interfaces ge-0/0/5.308 vlan-id 308

[edit]
lab@R8# set interfaces ge-0/0/5.308 description "P1-2 connection"

[edit]
lab@R8# set interfaces ge-0/0/5.308 family inet address 192.168.0.29/30

[edit]
lab@R8# set interfaces ge-0/0/5.308 family inet6 address fc09:c0:ffee::1/126
```

Configure the External BGP sessions on R8.

```
[edit]
lab@R8# set protocols bgp group P1-2 log-updown

[edit]
lab@R8# set protocols bgp group P1-2 hold-time 30

[edit]
lab@R8# set protocols bgp group P1-2 neighbor 192.168.0.30

[edit]
lab@R8# set protocols bgp group P1-2 peer-as 1679.12483

[edit]
lab@R8# set protocols bgp group P1-2-ipv6 log-updown

[edit]
lab@R8# set protocols bgp group P1-2-ipv6 hold-time 30

[edit]
lab@R8# set protocols bgp group P1-2-ipv6 peer-as 1679.12483

[edit]
lab@R8# set protocols bgp group P1-2-ipv6 neighbor fc09:c0:ffee::2

[edit]
lab@R8# set protocols bgp group P2-1 log-updown

[edit]
lab@R8# set protocols bgp group P2-1 hold-time 30

[edit]
lab@R8# set protocols bgp group P2-1 peer-as 43208.365

[edit]
lab@R8# set protocols bgp group P2-1 neighbor 192.168.0.6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# commit
commit complete
```

Verify the BGP summary table.

```
[edit]
lab@R8# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       680        669          0          0          0          0
inet6.0
                        16         16          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591         78        645       0       0       36:26 0/0/0/0
0/0/0/0
192.168.0.6      2831679853         96        490       0       0        1:15
113/116/116/0       0/0/0/0
192.168.0.30      110047427        487         91       0       0           1
556/564/564/0       0/0/0/0
fc09:c0:ffee::2   110047427         11         11       0       0        1:19 Establ
  inet6.0: 16/16/16/0
```

3) Apply routing policies to control route/ prefix exchange based on the requirements.

The routes exchanged with the external peers and customers should be controlled with routing policies. Those routing policies are based on the task requirements.

Several communities are used to mark the routes received from the external neighbors based on the neighbor type. You can define the communities on each router in the network for consistency and easy manipulation of the routes.

Define the same communities on all 8 routers.

```
[edit]
lab@R1# set policy-options community customer members 54591:200

[edit]
lab@R1# set policy-options community peer members 54591:100

[edit]
lab@R1# set policy-options community customer members 54591:200
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set policy-options community wildcard members *:*
```

Wildcard community matches any community attached to the routes. It will be used to delete any currently attached communities.

```
[edit]
lab@R1# set policy-options community rtbh members 6451.:666

[edit]
lab@R1# commit
commit complete
```

a) Configure routing policies on R1.

The private AS numbers should not be advertised to any external peer. The remove-private knob deletes the first occurrence of private AS numbers starting from the left of the AS-PATH string. The private AS range is predefined from 64512 to 65535.

Note: The examples below are for demonstration purposes and to show the reason why you must include the `remove-private` configuration option. You will not see this behavior until later in the lab once you have configured the next-hop self policy on another router to ensure you have routes to export to the neighbor.

Verify if routes with private AS numbers are advertised to the external peers.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.38 aspath-regex "64514
*"

inet.0: 751 destinations, 759 routes (751 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.32.0/24          Self                                    64514 I
* 172.31.33.0/24          Self                                    64514 I
* 172.31.34.0/24          Self                                    64514 I
* 172.31.35.0/24          Self                                    64514 I
* 172.31.36.0/24          Self                                    64514 I
* 172.31.37.0/24          Self                                    64514 I
* 172.31.38.0/24          Self                                    64514 I
* 172.31.39.0/24          Self                                    64514 I
* 172.31.40.0/24          Self                                    64514 I
--- (more) ---
```

Applying the **`remove-private`** knob.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols bgp group P1-1 remove-private

[edit]
lab@R1# set protocols bgp group P1-1-ipv6 remove-private

[edit]
lab@R1# commit
commit complete
```

Verify if the private AS numbers are removed from advertisements to P1-1.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.38 aspath-regex "64514
*"

inet.0: 751 destinations, 759 routes (751 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24           Self                                    I
* 172.31.1.0/24           Self                                    I
* 172.31.2.0/24           Self                                    I
* 172.31.3.0/24           Self                                    I
* 172.31.4.0/24           Self                                    I
* 172.31.5.0/24           Self                                    I
* 172.31.6.0/24           Self                                    I
--- (more) ---
```

On all BGP border routers you have to apply export policy to the Internal BGP session to changing the next-hop of the external routes. This way the other routers will be able to resolve the next-hop and will install the external routes in the routing table.

```
[edit]
lab@R1# set policy-options policy-statement nhs term 1 from protocol bgp

[edit]
lab@R1# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R1# set protocols bgp group ibgp export nhs

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

After applying the policy, the router advertises the external BGP routes with next-hop self.

```
[edit]
lab@R1# run show route advertising-protocol bgp 172.30.5.41

inet.0: 751 destinations, 759 routes (751 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/4               Self                         100        110047427 1620
61671 I
* 0.0.0.0/5               Self                         100        110047427 1620
61671 27075 I
* 1.64.0.0/10             Self                         100        110047427 1620
61671 I
* 1.84.160.0/20           Self                         100        110047427 1620
33112 I
* 1.96.0.0/11             Self                         100        110047427 1620
33112 63164 40776 51777 I
...
```

Configure policy to delete any existing communities to routes advertised to neighbor P1-1.

```
[edit]
lab@R1# set policy-option policy-statement delete-communities term 1 from protocol bgp

[edit]
lab@R1# set policy-option policy-statement delete-communities term 1 then community
delete wildcard
```

Configure a filter to the import routes received from the external BGP IPv4 session. The route 0.0.0.0 with any mask should not be accepted. Also routes with mask shorter than /8 and longer /24 should not be accepted.

```
[edit]
lab@R1# set policy-option policy-statement ebgp-import-filter term 1 from route-filter
0/0 upto /7

[edit]
lab@R1# set policy-option policy-statement ebgp-import-filter term 1 then reject

[edit]
lab@R1# set policy-option policy-statement ebgp-import-filter term 2 from route-filter
0/0 prefix-length-range /25-/32

[edit]
lab@R1# set policy-option policy-statement ebgp-import-filter term 2 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set policy-option policy-statement ebgp-import-filter term 3 from route-filter
0/0 through 0/32

[edit]
lab@R1# set policy-option policy-statement ebgp-import-filter term 3 then reject
```

The IPv6 routes accepted from P1-1 should be originated from their AS.

```
[edit]
lab@R1# set policy-options policy-statement ebgp-ipv6-import-filter term 1 from as-path
p1-ipv6-foreign

[edit]
lab@R1# set policy-options policy-statement ebgp-ipv6-import-filter term 1 then reject

[edit]
lab@R1# set policy-options as-path p1-ipv6-foreign ".{2,}"
```

The regular expression as-path is matching an AS path length of two or more AS numbers. Routes that match are not accepted.

The community value of 54591:100 should be set to the routes received from P1-1.

```
[edit]
lab@R1# set policy-options policy-statement peer-routes term 1 then community set peer
```

You must make sure that R1 is the preferred gateway for inbound and outbound IPv4 traffic to P1 AS. Increasing the local-preference ensures that R1 is preferred for outbound traffic to P1.

```
[edit]
lab@R1# set policy-options policy-statement p1-preference term 1 then local-preference
150
```

Routes received from peer P1 should not be advertised to peer P2 and vice versa. You cannot use community here because all routes from both peers are marked with the same community. You can use the as-path to match routes originated from P2 AS.

```
[edit]
lab@R1# set policy-options policy-statement no-p2-routes-export term 1 from protocol
bgp

[edit]
lab@R1# set policy-options policy-statement no-p2-routes-export term 1 from as-path p2-
neighbor
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set policy-options policy-statement no-p2-routes-export term 1 then reject

[edit]
lab@R1# set policy-options as-path p2-neighbor "2831679853 .*"
```

The local summary range, including the RIP routes should be advertised to all External BGP neighbors. The summary route will be exported in OSPF on R3 and R6 in a later step. R3 and R6 are chosen because they have all more specific routes in the routing table. You can use that summary route for BGP advertisements to other neighbors.

```
[edit]
lab@R1# set policy-options policy-statement local-range term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R1# set policy-options policy-statement local-range term 1 then accept
```

IPv6 routes advertised to P1 should not be advertised outside P1 AS. The well-known no-export community should be set to the IPv6 routes. This way P1 will not advertise these routes outside the AS.

```
[edit]
lab@R1# set policy-options policy-statement no-export-routes term 1 from protocol bgp

[edit]
lab@R1# set policy-options policy-statement no-export-routes term 1 then community add
no-export

[edit]
lab@R1# set policy-options community no-export members no-export
```

Any customer route marked with the community *:666 should be discarded. The policy below should be configured on all routers regardless if they have External peers.

```
[edit]
lab@R1# set policy-options policy-statement black-hole term 1 from protocol bgp

[edit]
lab@R1# set policy-options policy-statement black-hole term 1 from community rtbh

[edit]
lab@R1# set policy-options policy-statement black-hole term 1 then next-hop discard
```

Now that all policies are created, you must apply them on the neighbors in the right direction.

© 2020 Juniper Networks, Inc. All rights reserved.

The black-hole policy is applied in the import direction on all IBGP sessions with the route reflector.

```
[edit]
lab@R1# set protocols bgp group ibgp import black-hole
```

The ebgp-import-filter, peer-routes and p1-preference policies are applied in import direction to P1-1 IPv4 session.

```
[edit]
lab@R1# set protocols bgp group P1-1 import [ ebgp-import-filter peer-routes p1-
preference ]
```

The no-p2-routes-export, local-range and delete-communities is applied in export direction.

```
[edit]
lab@R1# set protocols bgp group P1-1 export [ no-p2-routes-export local-range delete-
communities ]
```

The following import policies must be applied to the IPV6 BGP session with neighbor P1-1: ebgp-ipv6-import and peer routes.

```
[edit]
lab@R1# set protocols bgp group P1-1-ipv6 import [ ebgp-ipv6-import-filter peer-routes
]
```

Also the following export policies must be applied: delete-communities and no-export-routes.

```
[edit]
lab@R1# set protocols bgp group P1-1-ipv6 export [ delete-communities no-export-routes
]
```

```
[edit]
lab@R1# commit
commit complete
```

b) Configure routing policies on R2.

Router R2 has External BGP neighborship with the customer C3-1 router.

The default community list is configured on R2 together with policies. Also, the remove-private knob is used.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set protocols bgp group C3-1 remove-private

[edit]
lab@R2# set policy-options policy-statement nhs term 1 from protocol bgp

[edit]
lab@R2# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R2# set policy-options policy-statement delete-communities term 1 from protocol bgp

[edit]
lab@R2# set policy-options policy-statement delete-communities term 1 then community
delete wildcard

[edit]
lab@R2# set policy-options policy-statement ebgp-import-filter term 1 from route-filter
0/0 upto /7

[edit]
lab@R2# set policy-options policy-statement ebgp-import-filter term 1 then reject

[edit]
lab@R2# set policy-options policy-statement ebgp-import-filter term 2 from route-filter
0/0 prefix-length-range /25-/32

[edit]
lab@R2# set policy-options policy-statement ebgp-import-filter term 2 then reject

[edit]
lab@R2# set policy-options policy-statement ebgp-import-filter term 3 from route-filter
0/0 through 0/32

[edit]
lab@R2# set policy-options policy-statement ebgp-import-filter term 3 then reject

[edit]
lab@R2# set policy-options policy-statement local-range term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R2# set policy-options policy-statement local-range term 1 then accept

[edit]
lab@R2# set policy-options policy-statement black-hole term 1 from protocol bgp

[edit]
lab@R2# set policy-options policy-statement black-hole term 1 from community rtbh
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set policy-options policy-statement black-hole term 1 then next-hop discard
```

Specific customer C3 policies on R2 should be configured according to the task requirements.

All customer routes should be marked with the community 54591:200.

```
[edit]
lab@R2# set policy-options policy-statement customer-routes term 1 then community add
customer
```

All direct routes received from customer should be preferred over the same routes received from other peers.

```
[edit]
lab@R2# set policy-options policy-statement customer-preferred term 1 then local-
preference 200
```

Applying the policies to the respected sessions.

```
[edit]
lab@R2# set protocols bgp group ibgp import black-hole

[edit]
lab@R2# set protocols bgp group ibgp export nhs
```

The `ebgp-import-filter`, `customer-routes` and `customer-preferred` policies are applied in the import direction.

```
[edit]
lab@R2# set protocols bgp group C3-1 import [ ebgp-import-filter customer-routes
customer-preferred ]
```

The `local-range` and `delete-communities` policies are both applied in the export direction.

```
[edit]
lab@R2# set protocols bgp group C3-1 export [ local-range delete-communities ]

[edit]
lab@R2# commit
commit complete
```

c) Configure routing policies on R3.

© 2020 Juniper Networks, Inc. All rights reserved.

Router R3 has one External BGP session with customer C2-1. R3 is the OSPF ABR that advertises the aggregate route to the other routers in the network.

Configure an export policy to advertise the aggregate in OSPF.

```
[edit]
lab@R3# set routing-options aggregate route 172.30.0.0/16
```

First, you must create a summary aggregate route in the routing table.

```
[edit]
lab@R3# set policy-options policy-statement local-range term 1 from protocol aggregate

[edit]
lab@R3# set policy-options policy-statement local-range term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R3# set policy-options policy-statement local-range term 1 then accept
```

And you must apply the policy to the OSPF process.

```
[edit]
lab@R3# set protocols ospf export local-range

[edit]
lab@R3# commit
commit complete
```

Configure the common policies on R3.

```
[edit]
lab@R3# set protocols bgp group C2-1 remove-private

[edit]
lab@R3# set policy-options policy-statement nhs term 1 from protocol bgp

[edit]
lab@R3# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R3# set policy-option policy-statement delete-communities term 1 from protocol bgp

[edit]
lab@R3# set policy-option policy-statement delete-communities term 1 then community
delete wildcard
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set policy-option policy-statement ebgp-import-filter term 1 from route-filter
0/0 upto /7

[edit]
lab@R3# set policy-option policy-statement ebgp-import-filter term 1 then reject

[edit]
lab@R3# set policy-option policy-statement ebgp-import-filter term 2 from route-filter
0/0 prefix-length-range /25-/32

[edit]
lab@R3# set policy-option policy-statement ebgp-import-filter term 2 then reject

[edit]
lab@R3# set policy-option policy-statement ebgp-import-filter term 3 from route-filter
0/0 through 0/32

[edit]
lab@R3# set policy-option policy-statement ebgp-import-filter term 3 then reject

[edit]
lab@R3# set policy-options policy-statement local-range term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R3# set policy-options policy-statement local-range term 1 then accept

[edit]
lab@R3# set policy-options policy-statement black-hole term 1 from protocol bgp

[edit]
lab@R3# set policy-options policy-statement black-hole term 1 from community rtbh

[edit]
lab@R3# set policy-options policy-statement black-hole term 1 then next-hop discard
```

You can configure and apply the same policies for customer C2-1 on R3.

```
[edit]
lab@R3# set policy-options policy-statement customer-routes term 1 then community add
customer

[edit]
lab@R3# set policy-options policy-statement customer-preferred term 1 then local-
preference 200
```

Apply the policies on R3 BGP sessions.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols bgp group ibgp import black-hole

[edit]
lab@R3# set protocols bgp group ibgp export nhs

[edit]
lab@R3# set protocols bgp group C2-1 import [ ebgp-import-filter customer-routes
customer-preferred ]

[edit]
lab@R3# set protocols bgp group C2-1 export [ local-range delete-communities ]

[edit]
lab@R3# commit
commit complete
```

d) Configure the routing policies on R4.

Router R4 has no External BGP sessions. However, you still have to define the common communities and apply the remote triggered black hole policy to discard traffic marked with the rtbh community.

```
[edit]
lab@R4# set policy-options policy-statement black-hole term 1 from protocol bgp

[edit]
lab@R4# set policy-options policy-statement black-hole term 1 from community rtbh

[edit]
lab@R4# set policy-options policy-statement black-hole term 1 then next-hop discard
```

Apply the policy in the import direction on the IBGP session with the route reflector.

```
[edit]
lab@R4# set protocols bgp group ibgp import black-hole

[edit]
lab@R4# commit
commit complete
```

e) Configure routing policies on R5.

Router R5 has an External BGP session with customer C1-1.

Define the list with communities and common routing policies used for all customer BGP sessions.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols bgp group C1-1 remove-private

[edit]
lab@R5# set policy-options policy-statement nhs term 1 from protocol bgp

[edit]
lab@R5# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R5# set policy-option policy-statement ebgp-import-filter term 1 from route-filter
0/0 upto /7

[edit]
lab@R5# set policy-option policy-statement ebgp-import-filter term 1 then reject

[edit]
lab@R5# set policy-option policy-statement ebgp-import-filter term 2 from route-filter
0/0 prefix-length-range /25-/32

[edit]
lab@R5# set policy-option policy-statement ebgp-import-filter term 2 then reject

[edit]
lab@R5# set policy-option policy-statement ebgp-import-filter term 3 from route-filter
0/0 through 0/32

[edit]
lab@R5# set policy-option policy-statement ebgp-import-filter term 3 then reject

[edit]
lab@R5# set policy-options policy-statement black-hole term 1 from protocol bgp

[edit]
lab@R5# set policy-options policy-statement black-hole term 1 from community rtbh

[edit]
lab@R5# set policy-options policy-statement black-hole term 1 then next-hop discard
```

For C1-1 you have requirement to advertise only a default route.

You can create single policy to advertise a default route and the summary routes only to customer C1-1. The default route is already present in the routing table received from both ABRs.

```
[edit]
lab@R5# set policy-options policy-statement default-and-local term 1 from route-filter
0/0 exact
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement default-and-local term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R5# set policy-options policy-statement default-and-local term 1 then accept

[edit]
lab@R5# set policy-options policy-statement default-and-local term 2 then reject

[edit]
lab@R5# set policy-options policy-statement customer-routes term 1 then community add
customer

[edit]
lab@R5# set policy-options policy-statement customer-preferred term 1 then local-
preference 200
```

A second term rejects everything else.

Applying the policies to the External BGP neighbor.

```
[edit]
lab@R5# set protocols bgp group ibgp import black-hole

[edit]
lab@R5# set protocols bgp group ibgp export nhs

[edit]
lab@R5# set protocols bgp group C1-1 import [damp-aggressive ebgp-import-filter
customer-routes customer-preferred ]

[edit]
lab@R5# set protocols bgp group C1-1 export default-and-local

[edit]
lab@R5# commit
commit complete
```

f) Configure routing policies on R6.

Router R6 has no External BGP neighbors. Only the black-hole policy should be configured and applied. Because R6 is the "backup" OSPF ABR, you also have to configure the OSPF export policy for the aggregate route.

```
[edit]
lab@R6# set routing-options aggregate route 172.30.0.0/16
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set policy-options policy-statement local-range term 1 from protocol aggregate

[edit]
lab@R6# set policy-options policy-statement local-range term 1 from route-filter
172.30.0.0/16 exact

[edit]
lab@R6# set policy-options policy-statement local-range term 1 then accept

[edit]
lab@R6# set protocols ospf export local-range
```

Define the common RTBH policy.

```
[edit]
lab@R6# set policy-options policy-statement black-hole term 1 from protocol bgp

[edit]
lab@R6# set policy-options policy-statement black-hole term 1 from community rtbh

[edit]
lab@R6# set policy-options policy-statement black-hole term 1 then next-hop discard
```

Apply the black-hole policy in the import direction to the IBGP session.

```
[edit]
lab@R6# set protocols bgp group ibgp import black-hole

[edit]
lab@R6# commit
commit complete
```

g) Configure routing policies on R7.

Router R7 has no External BGP sessions. The policy configuration of R4 can also be used on R7.

```
[edit]
lab@R7# set policy-options policy-statement black-hole term 1 from protocol bgp

[edit]
lab@R7# set policy-options policy-statement black-hole term 1 from community rtbh

[edit]
lab@R7# set policy-options policy-statement black-hole term 1 then next-hop discard
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols bgp group ibgp import black-hole

[edit]
lab@R7# commit
commit complete
```

h) Configure routing policies on R8.

Router R8 has two External BGP neighbors with P2-1 and P1-2 routers. The policies used on R1 for the P1-1 neighbor can also be used on R8, together with additional specific policies to meet the requirements for this router.

```
[edit]
lab@R8# set protocols bgp group P1-2 remove-private

[edit]
lab@R8# set protocols bgp group P1-2-ipv6 remove-private

[edit]
lab@R8# set protocols bgp group P2-1 remove-private

[edit]
lab@R8# set policy-options community no-export members no-export
```

The well-known community no-export is attached to routes advertisements to P1-2.

```
[edit]
lab@R8# set policy-options policy-statement delete-communities term 1 from protocol bgp

[edit]
lab@R8# set policy-options policy-statement delete-communities term 1 then community
delete wildcard

[edit]
lab@R8# set policy-options policy-statement nhs term 1 from protocol bgp

[edit]
lab@R8# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R8# set policy-option policy-statement ebgp-import-filter term 1 from route-filter
0/0 upto /7

[edit]
lab@R8# set policy-option policy-statement ebgp-import-filter term 1 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set policy-option policy-statement ebgp-import-filter term 2 from route-filter
0/0 prefix-length-range /25-/32

[edit]
lab@R8# set policy-option policy-statement ebgp-import-filter term 2 then reject

[edit]
lab@R8# set policy-option policy-statement ebgp-import-filter term 3 from route-filter
0/0 through 0/32

[edit]
lab@R8# set policy-option policy-statement ebgp-import-filter term 3 then reject

[edit]
lab@R8# set policy-option policy-statement ebgp-ipv6-import-filter term 1 from as-path
p1-ipv6-foreign

[edit]
lab@R8# set policy-option policy-statement ebgp-ipv6-import-filter term 1 then reject

[edit]
lab@R8# set policy-options as-path p1-ipv6-foreign ".{2,}"

[edit]
lab@R8# set policy-option policy-statement peer-routes term 1 then community set peer

[edit]
lab@R8# set policy-option policy-statement no-export-routes term 1 from protocol bgp

[edit]
lab@R8# set policy-option policy-statement no-export-routes term 1 then community add
no-export

[edit]
lab@R8# set policy-option policy-statement no-p2-routes-export term 1 from protocol bgp

[edit]
lab@R8# set policy-option policy-statement no-p2-routes-export term 1 from as-path p2-
neighbor

[edit]
lab@R8# set policy-options policy-statement no-p2-routes-export term 1 then reject

[edit]
lab@R8# set policy-options policy-statement local-range term 1 from route-filter
172.30.0.0/16 exact
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set policy-options policy-statement local-range term 1 then accept

[edit]
lab@R8# set policy-options as-path p2-neighbor "2831679853 .*"

[edit]
lab@R8# set policy-options policy-statement black-hole term 1 from protocol bgp

[edit]
lab@R8# set policy-options policy-statement black-hole term 1 from community rtbh

[edit]
lab@R8# set policy-options policy-statement black-hole term 1 then next-hop discard
```

You must make sure that router R1 is the preferred entry point for P1. You can accomplish that with advertising the routes from R8 to P1-2 with a longer AS path.

```
[edit]
lab@R8# set policy-option policy-statement long-as-path term 1 from protocol bgp

[edit]
lab@R8# set policy-option policy-statement long-as-path term 1 then as-path-prepend
"54591 54591 54591"
```

Routes with a /32 subnet mask can be accepted from neighbor P2-1.

```
[edit]
lab@R8# set policy-option policy-statement allow-p2-loopback term 1 from as-path p2-
native

[edit]
lab@R8# set policy-option policy-statement allow-p2-loopback term 1 from route-filter
0/0 prefix-length-range /32-/32

[edit]
lab@R8# set policy-option policy-statement allow-p2-loopback term 1 then community set
peer
```

Note: Do not forget to set the peer community to the /32 routes also.

```
[edit]
lab@R8# set policy-option policy-statement allow-p2-loopback term 1 then accept

[edit]
lab@R8# set policy-options as-path p2-native 2831679853
```

IPv4 prefixes with an AS path longer than 5 should not be accepted.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set policy-option policy-statement p2-long-path-filter term 1 from as-path p2-
long-path

[edit]
lab@R8# set policy-option policy-statement p2-long-path-filter term 1 then reject

[edit]
lab@R8# set policy-options as-path p2-long-path ".{6,}"
```

Routes between P1 and P2 AS should not be reachable through your local AS (transit). You already created a policy to not advertise the P2 routes to P1.

```
[edit]
lab@R8# set policy-option policy-statement no-p1-routes-export term 1 from protocol bgp

[edit]
lab@R8# set policy-option policy-statement no-p1-routes-export term 1 from as-path p1-
neighbor

[edit]
lab@R8# set policy-option policy-statement no-p1-routes-export term 1 then reject

[edit]
lab@R8# set policy-options as-path p1-neighbor "110047427 .*"
```

Apply the configured policies to the BGP sessions.

```
[edit]
lab@R8# set protocols bgp group ibgp import black-hole

[edit]
lab@R8# set protocols bgp group ibgp export nhs

[edit]
lab@R8# set protocols bgp group P1-2 import [ ebgp-import-filter peer-routes ]

[edit]
lab@R8# set protocols bgp group P1-2 export [ no-p2-routes-export local-range delete-
communities long-as-path]

[edit]
lab@R8# set protocols bgp group P1-2-ipv6 import [ ebgp-ipv6-import-filter peer-routes
]

[edit]
lab@R8# set protocols bgp group P1-2-ipv6 export [ delete-communities no-export-routes
]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set protocols bgp group P2-1 import [ allow-p2-loopback ebgp-import-filter p2-
long-path-filter peer-routes ]
```

Note: The policy allow-p2-loopback should be applied before the ebgp-import-filter policy, otherwise the routes with the /32 mask will be rejected.

```
[edit]
lab@R8# set protocols bgp group P2-1 export [ no-p1-routes-export local-range delete-
communities ]
```

```
[edit]
lab@R8# commit
commit complete
```

4) Verify the operation of the routing policies according to the requirements.

You must make sure that only the required prefixes are received and advertised.

a) Verify policy operation on R1.

Verify if unresolved next-hops are present in the next-hop resolution table.

```
[edit]
lab@R1# run show route resolution unresolved
Tree Index 1
```

Verify if routes with private AS are advertised.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.38 aspath-regex
"[64512-65535]"

inet.0: 752 destinations, 914 routes (752 active, 0 holddown, 162 hidden)
```

Verify if the received routes from peer P1-1 are advertised to the route reflector with the next hop changed to self and if the correct community is set.

```
[edit]
lab@R1# run show route advertising-protocol bgp 172.30.5.41 terse community-name
peer

inet.0: 752 destinations, 914 routes (752 active, 0 holddown, 162 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 1.64.0.0/10               Self                          150        110047427 1620
61671 I
* 1.84.160.0/20             Self                          150        110047427 1620
33112 I
* 1.96.0.0/11               Self                          150        110047427 1620
33112 63164 40776 51777 I
* 1.161.192.0/21            Self                          150        110047427 1620
33112 30404 32138 45045 I
* 1.166.89.0/24             Self                          150        110047427 56422
47123 34908 57085 3090 19128 21300 I
* 1.176.0.0/12              Self                          150        110047427 1620
33112 49129 16320 52954 I
* 3.49.128.0/19             Self                          150        110047427 56422
7923 42335 17359 9293 38097 38522 61520 I
...
```

Verify if the customer routes are advertised to the P1-1 neighbor.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.38 | match 172.31
* 172.31.0.0/24             Self                                       I
* 172.31.1.0/24             Self                                       I
* 172.31.2.0/24             Self                                       I
* 172.31.3.0/24             Self                                       I
* 172.31.4.0/24             Self                                       I
* 172.31.5.0/24             Self                                       I
* 172.31.6.0/24             Self                                       I
* 172.31.7.0/24             Self                                       I
* 172.31.8.0/24             Self                                       I
* 172.31.9.0/24             Self                                       I
* 172.31.10.0/24            Self                                       I
* 172.31.11.0/24            Self                                       I
* 172.31.12.0/24            Self                                       I
* 172.31.13.0/24            Self                                       I
* 172.31.14.0/24            Self                                       I
* 172.31.15.0/24            Self                                       I
* 172.31.16.0/24            Self                                       I
* 172.31.17.0/24            Self                                       I
* 172.31.18.0/24            Self                                       I
* 172.31.19.0/24            Self                                       I
* 172.31.20.0/24            Self                                       I
* 172.31.21.0/24            Self                                       I
* 172.31.22.0/24            Self                                       I
* 172.31.23.0/24            Self                                       I
* 172.31.24.0/24            Self                                       I
* 172.31.25.0/24            Self                                       I
* 172.31.26.0/24            Self                                       I
* 172.31.27.0/24            Self                                       I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.28.0/24          Self                                I
* 172.31.29.0/24          Self                                I
* 172.31.30.0/24          Self                                I
* 172.31.31.0/24          Self                                I
* 172.31.32.0/24          Self                                I
* 172.31.33.0/24          Self                                I
* 172.31.34.0/24          Self                                I
* 172.31.35.0/24          Self                                I
* 172.31.36.0/24          Self                                I
* 172.31.37.0/24          Self                                I
* 172.31.38.0/24          Self                                I
* 172.31.39.0/24          Self                                I
* 172.31.40.0/24          Self                                I
* 172.31.41.0/24          Self                                I
* 172.31.42.0/24          Self                                I
* 172.31.43.0/24          Self                                I
* 172.31.44.0/24          Self                                I
* 172.31.45.0/24          Self                                I
* 172.31.46.0/24          Self                                I
* 172.31.47.0/24          Self                                I
```

Verify if routes with any community set are advertised to neighbor P1-1.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.38 community-name peer

inet.0: 685 destinations, 685 routes (523 active, 0 holddown, 162 hidden)
```

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.38 community-name
customer

inet.0: 685 destinations, 685 routes (523 active, 0 holddown, 162 hidden)
```

Verify if routes 0.0.0.0 with any mask are filtered.

```
[edit]
lab@R1# run show route protocol bgp terse | match "0.0.0.0"
```

Verify if routes with a mask shorter than /8 and longer than /24 are filtered.

```
[edit]
lab@R1# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32     B 170        100              >172.30.0.10      2831679853 I
* 172.17.47.3/32     B 170        100              >172.30.0.10      2831679853 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Only two routes with mask /32 are received from AS 43208.365, this is allowed by the requirements.

Verify if the routes from peer P2 AS are not advertised to P1 AS.

```
[edit]
lab@R1# run show route protocol bgp 192.168.0.38 aspath-regex "2831679853 .*"

inet.0: 701 destinations, 701 routes (539 active, 0 holddown, 162 hidden)
```

Verify if routes marked with rtbh community are discarded.

```
[edit]
lab@R1# run show route protocol bgp terse community-name rtbh

inet.0: 701 destinations, 701 routes (539 active, 0 holddown, 162 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop         AS path
* 172.31.1.0/24      B 170        200             Discard          64512 I

inet6.0: 33 destinations, 37 routes (33 active, 0 holddown, 0 hidden)
```

Verify if the local summary route is advertised to the External BGP neighbor.

```
[edit]
lab@R1# run show route advertising-protocol bgp 192.168.0.38 | match 172.30.
* 172.30.0.0/16          Self                 0                    I
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify if the IPv6 routes are received from neighbor P1-1.

```
[edit]
lab@R1# run show route table inet6.0

inet6.0: 33 destinations, 37 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fc09:c0:ffee::4/126*[Direct/0] 03:11:57
                    > via ge-0/0/5.310
fc09:c0:ffee::5/128*[Local/0] 03:12:10
                      Local via ge-0/0/5.310
fd01:aaaa:bbbb::/80*[BGP/170] 03:11:46, localpref 100
                      AS path: 110047427 I
                     > to fc09:c0:ffee::6 via ge-0/0/5.310
fd01:aaaa:bbbb:0:1::/80
                    *[BGP/170] 03:11:46, localpref 100
                      AS path: 110047427 I
                     > to fc09:c0:ffee::6 via ge-0/0/5.310
fd01:aaaa:bbbb:0:2::/80
                    *[BGP/170] 03:11:46, localpref 100
                      AS path: 110047427 I
                     > to fc09:c0:ffee::6 via ge-0/0/5.310
...
```

Verify if routes with an AS path longer than 2 are received.

```
[edit]
lab@R1# run show route table inet6.0 aspath-regex ".{2,}"

inet6.0: 33 destinations, 37 routes (33 active, 0 holddown, 0 hidden)
```

No IPv6 routes are advertised from R1. Family IPv6 unicast is not signaled with the route reflector, those routes will be transported over LSPs in the next chapters.

```
[edit]
lab@R1# run show route advertising-protocol bgp fc09:c0:ffee::6
```

b) Verify policy operation on router **R2**.

```
[edit]
lab@R2# run show route resolution unresolved
Tree Index 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.0.10 aspath-regex
"[64512-65535]"

inet.0: 750 destinations, 751 routes (750 active, 0 holddown, 0 hidden)
```

Routes received from customer C3-1 have a higher local preference as they should be preferred over the same routes received from other neighbors.

```
[edit]
lab@R2# run show route advertising-protocol bgp 172.30.5.41 terse community-name
customer

inet.0: 750 destinations, 751 routes (750 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.32.0/24          Self                         200        64514 I
* 172.31.33.0/24          Self                         200        64514 I
* 172.31.34.0/24          Self                         200        64514 I
* 172.31.35.0/24          Self                         200        64514 I
* 172.31.36.0/24          Self                         200        64514 I
* 172.31.37.0/24          Self                         200        64514 I
* 172.31.38.0/24          Self                         200        64514 I
* 172.31.39.0/24          Self                         200        64514 I
* 172.31.40.0/24          Self                         200        64514 I
* 172.31.41.0/24          Self                         200        64514 I
* 172.31.42.0/24          Self                         200        64514 I
* 172.31.43.0/24          Self                         200        64514 I
* 172.31.44.0/24          Self                         200        64514 I
* 172.31.45.0/24          Self                         200        64514 I
* 172.31.46.0/24          Self                         200        64514 I
* 172.31.47.0/24          Self                         200        64514 I
```

```
[edit]
lab@R2# run show route protocol bgp terse | match " 0.0.0.0"
```

```
[edit]
lab@R2# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32     B 170        100             >172.30.0.1       2831679853 I
* 172.17.47.3/32     B 170        100             >172.30.0.1       2831679853 I
```

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.0.10 community-name peer
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 521 destinations, 522 routes (521 active, 0 holddown, 0 hidden)
```

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.0.10 community-name
customer

inet.0: 521 destinations, 522 routes (521 active, 0 holddown, 0 hidden)
```

```
[edit]
lab@R2# run show route advertising-protocol bgp 192.168.0.10 | match 172.30.
* 172.30.0.0/16            Self                  0                      I
```

```
[edit]
lab@R2# run show route protocol bgp terse community-name rtbh

inet.0: 537 destinations, 538 routes (537 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination         P Prf   Metric 1   Metric 2  Next hop         AS path
* 172.31.1.0/24       B 170        200             Discard          64512 I

inet6.0: 33 destinations, 37 routes (33 active, 0 holddown, 0 hidden)
```

c) Verify policy operation on router **R3**.

```
[edit]
lab@R3# run show route resolution unresolved
Tree Index 1
```

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.31.31.1 aspath-regex "[64512-
65535]"

inet.0: 774 destinations, 776 routes (774 active, 0 holddown, 0 hidden)
```

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.5.41 terse community-name
customer

inet.0: 774 destinations, 776 routes (774 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.16.0/24          Self                         200        64513 I
* 172.31.17.0/24          Self                         200        64513 I
* 172.31.18.0/24          Self                         200        64513 I
* 172.31.19.0/24          Self                         200        64513 I
* 172.31.20.0/24          Self                         200        64513 I
* 172.31.21.0/24          Self                         200        64513 I
* 172.31.22.0/24          Self                         200        64513 I
* 172.31.23.0/24          Self                         200        64513 I
* 172.31.24.0/24          Self                         200        64513 I
* 172.31.25.0/24          Self                         200        64513 I
* 172.31.26.0/24          Self                         200        64513 I
* 172.31.27.0/24          Self                         200        64513 I
* 172.31.28.0/24          Self                         200        64513 I
* 172.31.29.0/24          Self                         200        64513 I
* 172.31.30.0/24          Self                         200        64513 I
* 172.31.31.0/24          Self                         200        64513 I
```

```
[edit]
lab@R3# run show route protocol bgp terse | match " 0.0.0.0"
```

```
[edit]
lab@R3# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32     B 170        100             >172.30.0.34      2831679853 I
* 172.17.47.3/32     B 170        100             >172.30.0.34      2831679853 I
```

```
[edit]
lab@R3# run show route protocol bgp terse community-name rtbh

inet.0: 561 destinations, 563 routes (561 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination         P Prf   Metric 1   Metric 2  Next hop         AS path
* 172.31.1.0/24       B 170        200             Discard          64512 I

inet6.0: 16 destinations, 21 routes (16 active, 0 holddown, 0 hidden)
```

d) Verify policy operation on router R4.

```
[edit]
lab@R4# run show route resolution unresolved
Tree Index 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show route protocol bgp terse | match " 0.0.0.0"
```

```
[edit]
lab@R4# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32     B 170        100             >172.30.0.81      2831679853 I
* 172.17.47.3/32     B 170        100             >172.30.0.81      2831679853 I
```

```
[edit]
lab@R4# run show route protocol bgp terse community-name rtbh

inet.0: 555 destinations, 555 routes (555 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop         AS path
* 172.31.1.0/24      B 170        200             Discard          64512 I

inet6.0: 14 destinations, 15 routes (14 active, 0 holddown, 0 hidden)
```

e) Verify policy operation on router R5.

```
[edit]
lab@R5# run show route resolution unresolved
Tree Index 1
```

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.18 aspath-regex
"[64512-65535]"

inet.0: 771 destinations, 771 routes (771 active, 0 holddown, 0 hidden)
```

```
[edit]
lab@R5# run show route advertising-protocol bgp 172.30.5.41 terse community-name
customer

inet.0: 771 destinations, 771 routes (771 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.0.0/24           Self                         200        64512 I
* 172.31.1.0/24           Self                         200        64512 I
* 172.31.2.0/24           Self                         200        64512 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.3.0/24          Self                  200        64512 I
* 172.31.4.0/24          Self                  200        64512 I
* 172.31.5.0/24          Self                  200        64512 I
* 172.31.6.0/24          Self                  200        64512 I
* 172.31.7.0/24          Self                  200        64512 I
* 172.31.8.0/24          Self                  200        64512 I
* 172.31.9.0/24          Self                  200        64512 I
* 172.31.10.0/24         Self                  200        64512 I
* 172.31.11.0/24         Self                  200        64512 I
* 172.31.12.0/24         Self                  200        64512 I
* 172.31.13.0/24         Self                  200        64512 I
* 172.31.14.0/24         Self                  200        64512 I
* 172.31.15.0/24         Self                  200        64512 I
```

```
[edit]
lab@R5# run show route protocol bgp terse | match " 0.0.0.0"
```

```
[edit]
lab@R5# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32     B 170        100             >172.30.0.94      2831679853 I
* 172.17.47.3/32     B 170        100             >172.30.0.94      2831679853 I
```

The default route is advertised to customer C1-1 as required by the task.

```
[edit]
lab@R5# run show route advertising-protocol bgp 192.168.0.18

inet.0: 558 destinations, 558 routes (558 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/0               Self                 15                 I
* 172.30.0.0/16           Self                 0                  I
```

```
[edit]
lab@R5# run show route protocol bgp terse community-name rtbh

inet.0: 558 destinations, 558 routes (558 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop         AS path
* 172.31.1.0/24      B 170        200             >192.168.0.18     64512 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet6.0: 14 destinations, 15 routes (14 active, 0 holddown, 0 hidden)
```

f) Verify policy operation on router R6.

```
[edit]
lab@R6# run show route resolution unresolved
Tree Index 1
```

```
[edit]
lab@R6# run show route protocol bgp terse | match " 0.0.0.0"
```

```
[edit]
lab@R6# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32     B 170        100             >172.30.0.22      2831679853 I
* 172.17.47.3/32     B 170        100             >172.30.0.22      2831679853 I
```

```
[edit]
lab@R6# run show route protocol bgp terse community-name rtbh

inet.0: 562 destinations, 580 routes (562 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop         AS path
* 172.31.1.0/24      B 170        200             Discard          64512 I

inet6.0: 16 destinations, 20 routes (16 active, 0 holddown, 0 hidden)
```

g) Verify policy operation on router R7.

```
[edit]
lab@R7# run show route resolution unresolved
Tree Index 1
```

```
[edit]
lab@R7# run show route protocol bgp terse | match " 0.0.0.0"
```

```
[edit]
lab@R7# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.17.47.2/32     B 170        100              >172.30.0.5          2831679853 I
* 172.17.47.3/32     B 170        100              >172.30.0.5          2831679853 I
```

```
[edit]
lab@R7# run show route protocol bgp terse community-name rtbh

inet.0: 540 destinations, 540 routes (540 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop         AS path
* 172.31.1.0/24      B 170        200             Discard          64512 I

inet6.0: 14 destinations, 16 routes (14 active, 0 holddown, 0 hidden)
```

h) Verify policy operation on router R8.

```
[edit]
lab@R8# run show route resolution unresolved
Tree Index 1
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.6 aspath-regex "[64512-
65535]"

inet.0: 751 destinations, 1164 routes (538 active, 0 holddown, 222 hidden)
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.30 aspath-regex
"[64512-65535]"

inet.0: 751 destinations, 1164 routes (538 active, 0 holddown, 222 hidden)
```

```
[edit]
lab@R8# run show route advertising-protocol bgp fc09:c0:ffee::2 aspath-regex
"[64512-65535]"
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.30.5.41 terse community-name
peer

inet.0: 751 destinations, 1164 routes (538 active, 0 holddown, 222 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Prefix                  Nexthop              MED     Lclpref    AS path
* 17.205.88.0/21          Self                         100        2831679853 9726
19412 16490 6122 I
* 21.0.0.0/8              Self                         100        2831679853
26697 56251 10277 21151 I
* 45.48.0.0/12            Self                         100        2831679853
26697 38425 19086 12180 I
* 53.188.0.0/14           Self                         100        2831679853 9726
63776 34040 33332 ?
* 62.96.0.0/11            Self                         100        2831679853
26697 63776 24878 36176 I
* 64.0.0.0/8              Self                         100        2831679853
26697 56251 I
* 72.52.208.0/20          Self                         100        2831679853 9726
4341 50484 ?
...
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.6 | match 172.31
* 172.31.0.0/24           Self                                          I
* 172.31.1.0/24           Self                                          I
* 172.31.2.0/24           Self                                          I
* 172.31.3.0/24           Self                                          I
* 172.31.4.0/24           Self                                          I
* 172.31.5.0/24           Self                                          I
* 172.31.6.0/24           Self                                          I
* 172.31.7.0/24           Self                                          I
* 172.31.8.0/24           Self                                          I
* 172.31.9.0/24           Self                                          I
* 172.31.10.0/24          Self                                          I
* 172.31.11.0/24          Self                                          I
* 172.31.12.0/24          Self                                          I
* 172.31.13.0/24          Self                                          I
* 172.31.14.0/24          Self                                          I
* 172.31.15.0/24          Self                                          I
...
```

The local AS path is prepended to the routes advertised to P1-2, making R1 the preferred entry point.

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.30 | match 172.31
* 172.31.0.0/24           Self                                    54591 54591
54591 [54591] I
* 172.31.1.0/24           Self                                    54591 54591
54591 [54591] I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.31.2.0/24          Self                                  54591 54591
54591 [54591] I
* 172.31.3.0/24          Self                                  54591 54591
54591 [54591] I
* 172.31.4.0/24          Self                                  54591 54591
54591 [54591] I
* 172.31.5.0/24          Self                                  54591 54591
54591 [54591] I
* 172.31.6.0/24          Self                                  54591 54591
54591 [54591] I
...
```

Routes from P1 AS are not advertised to P2 AS and vice versa.

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.6 aspath-regex
"2831679853 .*"

inet.0: 751 destinations, 1164 routes (538 active, 0 holddown, 222 hidden)
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.30 aspath-regex
"110047427 .*"

inet.0: 751 destinations, 1164 routes (538 active, 0 holddown, 222 hidden)
```

Routes with an AS path longer than 5 are not accepted from neighbor P2-1.

```
lab@R8# run show route receive-protocol bgp 192.168.0.6 aspath-regex ".{6,}"

inet.0: 751 destinations, 1164 routes (538 active, 0 holddown, 222 hidden)

inet6.0: 33 destinations, 36 routes (33 active, 0 holddown, 0 hidden)
```

```
[edit]
lab@R8# run show route protocol bgp terse | match " 0.0.0.0"
```

```
[edit]
lab@R8# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9] )|(/3[0-2]
)"
* 172.17.47.2/32     B 170        100             >192.168.0.6      2831679853 I
* 172.17.47.3/32     B 170        100             >192.168.0.6      2831679853 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.6 | match 172.30.
* 172.30.0.0/16             Self                    0                    I
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.6 community-name peer

inet.0: 751 destinations, 1164 routes (538 active, 0 holddown, 222 hidden)
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 192.168.0.6 community-name
customer

inet.0: 751 destinations, 1164 routes (538 active, 0 holddown, 222 hidden)
```

```
[edit]
lab@R8# run show route protocol bgp terse community-name rtbh

inet.0: 751 destinations, 1164 routes (538 active, 0 holddown, 222 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop         AS path
* 172.31.1.0/24      B 170        200             Discard          64512 I

inet6.0: 33 destinations, 36 routes (33 active, 0 holddown, 0 hidden)
```

IPv6 routes with an AS path longer than 2 are not accepted from P1-2 neighbor.

```
[edit]
lab@R8# run show route table inet6.0 aspath-regex ".{2,}"

inet6.0: 33 destinations, 36 routes (33 active, 0 holddown, 0 hidden)
```

Task 5: MPLS Configuration

In this chapter the network will be configured with various LDP and RSVP transport LSPs and traffic protection mechanisms must be implemented.

1) Configure LDP on the interfaces shown in Table 29. Configure LDP on the route reflector. Configure MD5 authentication for all LDP sessions. Configure

© 2020 Juniper Networks, Inc. All rights reserved.

OSPF to track the status of LDP on the interfaces and Configure the LDP routes to use the IGP metric.

Router	Interface
R1	ge-0/0/4.206
R2	ge-0/0/4.207
R3	ge-0/0/4.135
	ae0.0
R4	ge-0/0/4.146
	ae0.0
R5	ge-0/0/4.135
	ae0.0
R6	ge-0/0/4.146
	ae0.0

Table 29

The MPLS network design uses RSVP signaled LSPs in Area 0. LDP is used towards the route reflector and in Area 1. LDP tunneling is used to connect both LDP islands. The route reflector is enabled for MPLS forwarding with LDP interfaces towards R1 and R2.

a) Configure router R1.

The **family mpls** configuration should be applied on all interfaces forwarding labeled traffic. Additionally, the interfaces should be added to the **protocols mpls** hierarchy.

```
[edit]
lab@R1# set interfaces ae0.0 family mpls

[edit]
lab@R1# set interfaces ge-0/0/4.117 family mpls

[edit]
lab@R1# set interfaces ge-0/0/4.118 family mpls

[edit]
lab@R1# set interfaces ge-0/0/4.206 family mpls

[edit]
lab@R1# set protocols mpls interface ae0.0

[edit]
lab@R1# set protocols mpls interface ge-0/0/4.117
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols mpls interface ge-0/0/4.118

[edit]
lab@R1# set protocols mpls interface ge-0/0/4.206
```

Adding interface to the route reflector to the LDP process and configuring authentication for the session.

```
[edit]
lab@R1# set protocols ldp interface ge-0/0/4.206

[edit]
lab@R1# set protocols ldp track-igp-metric

[edit]
lab@R1# set protocols ldp session 172.30.5.41 authentication-key workbook
```

Configure LDP synchronization for the OSPF interfaces. Be careful to enable `ldp-synchronization` only under interfaces enabled for LDP. If an LDP session is not established, OSPF will announce a maximum metric on that interface and this will influence IGP path selection.

```
[edit]
lab@R1# set protocols ospf area 0 interface ge-0/0/4.206 ldp-synchronization

[edit]
lab@R1# commit
commit complete
```

b) Configure the route reflector.

```
[edit]
lab@route-reflector# set interfaces ge-0/0/1.206 family mpls

[edit]
lab@route-reflector# set interfaces ge-0/0/1.207 family mpls

[edit]
lab@route-reflector# set protocols mpls interface ge-0/0/1.206

[edit]
lab@route-reflector# set protocols mpls interface ge-0/0/1.207

[edit]
lab@route-reflector# set protocols ldp track-igp-metric
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# set protocols ldp interface ge-0/0/1.206

[edit]
lab@route-reflector# set protocols ldp interface ge-0/0/1.207

[edit]
lab@route-reflector# set protocols ldp interface lo0.0

[edit]
lab@route-reflector# set protocols ldp session 172.30.5.1 authentication-key workbook

[edit]
lab@route-reflector# set protocols ldp session 172.30.5.2 authentication-key workbook

[edit]
lab@route-reflector# commit
commit complete
```

c) Configure R2.

The configuration on R2 is identical as on R1.

```
[edit]
lab@R2# set interfaces ae0.0 family mpls

[edit]
lab@R2# set interfaces ge-0/0/4.123 family mpls

[edit]
lab@R2# set interfaces ge-0/0/4.126 family mpls

[edit]
lab@R2# set interfaces ge-0/0/4.207 family mpls

[edit]
lab@R2# set protocols mpls interface ae0.0

[edit]
lab@R2# set protocols mpls interface ge-0/0/4.123

[edit]
lab@R2# set protocols mpls interface ge-0/0/4.126

[edit]
lab@R2# set protocols mpls interface ge-0/0/4.207
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set protocols ldp interface ge-0/0/4.207

[edit]
lab@R2# set protocols ldp track-igp-metric

[edit]
lab@R2# set protocols ldp session 172.30.5.41 authentication-key workbook

[edit]
lab@R2# set protocols ospf area 0 interface ge-0/0/4.207 ldp-synchronization

[edit]
lab@R2# commit
commit complete
```

Verify the operation of LDP on the route reflector.

```
[edit]
lab@route-reflector# run show ldp neighbor
Address              Interface          Label space ID         Hold time
172.30.0.65          ge-0/0/1.206       172.30.5.1:0             14
172.30.0.69          ge-0/0/1.207       172.30.5.2:0             11
```

```
[edit]
lab@route-reflector# run show ldp database
Input label database, 172.30.5.41:0--172.30.5.1:0
  Label     Prefix
      3     172.30.5.1/32
 299792     172.30.5.2/32
 299776     172.30.5.41/32

Output label database, 172.30.5.41:0--172.30.5.1:0
  Label     Prefix
 100000     172.30.5.1/32
 100016     172.30.5.2/32
      3     172.30.5.41/32

Input label database, 172.30.5.41:0--172.30.5.2:0
  Label     Prefix
 299792     172.30.5.1/32
      3     172.30.5.2/32
 299776     172.30.5.41/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Output label database, 172.30.5.41:0--172.30.5.2:0
  Label     Prefix
 100000     172.30.5.1/32
 100016     172.30.5.2/32
      3     172.30.5.41/32
```

Verify LDP operation on R1.

```
[edit]
lab@R1# run show ldp neighbor extensive
Address                 Interface              Label space ID         Hold time
172.30.0.66             ge-0/0/4.206           172.30.5.41:0            14
  Transport address: 172.30.5.41, Configuration sequence: 1
  Up for 00:20:59
  Reference count: 1
  Hold time: 15, Proposed local/peer: 15/15
  Hello flags: none
  Neighbor types: discovered
```

```
[edit]
lab@R1# run show ldp session extensive
Address: 172.30.5.41, State: Operational, Connection: Open, Hold time: 24
  Session ID: 172.30.5.1:0--172.30.5.41:0
  Next keepalive in 4 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 172.30.5.1, Remote address: 172.30.5.41
  Up for 00:21:15
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Authentication type: MD5
  Local - Restart: disabled, Helper mode: enabled
  Remote - Restart: disabled, Helper mode: enabled
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    172.30.0.66
    172.30.0.70
    172.30.5.41
  Queue depth: 0
Message type              Total                 Last 5 seconds
                      Sent      Received          Sent      Received
Initialization           1             1             0             0
Keepalive              127           127             1             1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Notification                0          0          0          0
Address                     1          1          0          0
Address withdraw            0          0          0          0
Label mapping               3          3          0          0
Label request               0          0          0          0
Label withdraw              0          0          0          0
Label release               0          0          0          0
Label abort                 0          0          0          0
```

```
[edit]
lab@R1# run show ospf interface ge-0/0/4.206 extensive

ge-0/0/4.206        PtToPt  0.0.0.0         0.0.0.0         0.0.0.0             1
  Type: P2P, Address: 172.30.0.65, Mask: 255.255.255.252, MTU: 1500, Cost: 10
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 10
  LDP sync state: in sync, for: 00:21:55, reason: LDP session up
            config holdtime: infinity
```

LDP synchronization is enabled and the session is established.

Verify the operation of LDP on R2.

```
[edit]
lab@R2# run show ldp session
  Address           State        Connection     Hold time
172.30.5.41         Operational  Open             26
```

```
[edit]
lab@R2# run show route table inet.3

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       *[LDP/9] 00:20:57, metric 11
                     > to 172.30.0.70 via ge-0/0/4.207, Push 100000
172.30.5.41/32      *[LDP/9] 00:21:07, metric 10
                     > to 172.30.0.70 via ge-0/0/4.207
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show ldp database
Input label database, 172.30.5.2:0--172.30.5.41:0
  Label      Prefix
 100000      172.30.5.1/32
 100016      172.30.5.2/32
      3      172.30.5.41/32

Output label database, 172.30.5.2:0--172.30.5.41:0
  Label      Prefix
 299792      172.30.5.1/32
      3      172.30.5.2/32
 299776      172.30.5.41/32
```

d) Configure R3.

On R3 both interfaces to R4 and R5 must be enabled for LDP.

```
[edit]
lab@R3# set interfaces ae0.0 family mpls

[edit]
lab@R3# set interfaces ge-0/0/4.123 family mpls

[edit]
lab@R3# set interfaces ge-0/0/4.135 family mpls

[edit]
lab@R3# set interfaces ge-0/0/4.137 family mpls

[edit]
lab@R3# set interfaces ge-0/0/4.138 family mpls

[edit]
lab@R3# set protocols mpls interface ae0.0

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.123

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.135

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.137


[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set protocols mpls interface ge-0/0/4.138

[edit]
lab@R3# set protocols ldp interface ge-0/0/4.135

[edit]
lab@R3# set protocols ldp interface ae0.0

[edit]
lab@R3# set protocols ldp session 172.30.5.4 authentication-key workbook

[edit]
lab@R3# set protocols ldp session 172.30.5.5 authentication-key workbook

[edit]
lab@R3# set protocols ospf area 0.0.0.1 interface ae0.0 ldp-synchronization

[edit]
lab@R3# set protocols ospf area 0.0.0.1 interface ge-0/0/4.135 ldp-synchronization

[edit]
lab@R3# commit
commit complete
```

a) Configure router R4.

```
[edit]
lab@R4# set interfaces ae0.0 family mpls

[edit]
lab@R4# set interfaces ge-0/0/4.146 family mpls

[edit]
lab@R4# set protocols mpls interface ae0.0

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.146

[edit]
lab@R4# set protocols ldp interface ae0.0

[edit]
lab@R4# set protocols ldp interface ge-0/0/4.146

[edit]
lab@R4# set protocols ldp track-igp-metric

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# set protocols ldp session 172.30.5.3 authentication-key workbook

[edit]
lab@R4# set protocols ldp session 172.30.5.6 authentication-key workbook

[edit]
lab@R4# set protocols ospf area 0.0.0.1 interface ae0.0 ldp-synchronization

[edit]
lab@R4# set protocols ospf area 0.0.0.1 interface ge-0/0/4.146 ldp-synchronization

[edit]
lab@R3# commit
commit complete
```

a) Configure router R5.

```
[edit]
lab@R5# set interfaces ae0.0 family mpls

[edit]
lab@R5# set interfaces ge-0/0/4.135 family mpls

[edit]
lab@R5# set protocols mpls interface ae0.0

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.135

[edit]
lab@R5# set protocols ldp interface ae0.0

[edit]
lab@R5# set protocols ldp interface ge-0/0/4.135

[edit]
lab@R5# set protocols ldp track-igp-metric

[edit]
lab@R5# set protocols ldp session 172.30.5.3 authentication-key workbook

[edit]
lab@R5# set protocols ldp session 172.30.5.6 authentication-key workbook

[edit]
lab@R5# set protocols ospf area 0.0.0.1 interface ae0.0 ldp-synchronization

[edit]
lab@R5# set protocols ospf area 0.0.0.1 interface ge-0/0/4.135 ldp-synchronization
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# commit
commit complete
```

a) Configure router R6.

```
[edit]
lab@R6# set interfaces ae0.0 family mpls

[edit]
lab@R6# set interfaces ge-0/0/4.126 family mpls

[edit]
lab@R6# set interfaces ge-0/0/4.146 family mpls

[edit]
lab@R6# set interfaces ge-0/0/4.167 family mpls

[edit]
lab@R6# set interfaces ge-0/0/4.168 family mpls

[edit]
lab@R6# set protocols mpls interface ae0.0

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.126

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.146

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.167

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.168

[edit]
lab@R6# set protocols ldp interface ge-0/0/4.146

[edit]
lab@R6# set protocols ldp interface ae0.0

[edit]
lab@R6# set protocols ldp track-igp-metric

[edit]
lab@R6# set protocols ldp session 172.30.5.4 authentication-key workbook
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols ldp session 172.30.5.5 authentication-key workbook

[edit]
lab@R6# set protocols ospf area 0.0.0.1 interface ae0.0 ldp-synchronization

[edit]
lab@R6# set protocols ospf area 0.0.0.1 interface ge-0/0/4.146 ldp-synchronization

[edit]
lab@R6# commit
commit complete
```

Verify the operation of LDP on R3.

```
[edit]
lab@R3# run show ldp session
  Address           State        Connection     Hold time
172.30.5.4          Operational  Open             24
172.30.5.5          Operational  Open             24
```

```
[edit]
lab@R3# run show ldp neighbor
Address            Interface          Label space ID         Hold time
172.30.0.82        ae0.0              172.30.5.4:0             12
172.30.0.86        ge-0/0/4.135       172.30.5.5:0             10
```

Verify LDP operation on router R4.

```
[edit]
lab@R4# run show ldp session
  Address           State        Connection     Hold time
172.30.5.3          Operational  Open             23
172.30.5.6          Operational  Open             27
```

```
[edit]
lab@R4# run show ldp neighbor
Address            Interface          Label space ID         Hold time
172.30.0.81        ae0.0              172.30.5.3:0             11
172.30.0.90        ge-0/0/4.146       172.30.5.6:0             10
```

```
[edit]
lab@R4# run show route table inet.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.3/32      *[LDP/9] 00:10:40, metric 5
                    > to 172.30.0.81 via ae0.0
172.30.5.5/32      *[LDP/9] 00:04:18, metric 15
                       to 172.30.0.81 via ae0.0, Push 299792
                    > to 172.30.0.90 via ge-0/0/4.146, Push 299776
172.30.5.6/32      *[LDP/9] 00:04:18, metric 10
                    > to 172.30.0.90 via ge-0/0/4.146
```

Verify LDP operation of the R5.

```
[edit]
lab@R5# run show ldp session
  Address           State        Connection     Hold time
172.30.5.3          Operational  Open             22
172.30.5.6          Operational  Open             27
```

```
[edit]
lab@R5# run show ldp neighbor
Address            Interface          Label space ID         Hold time
172.30.0.94        ae0.0              172.30.5.6:0             12
172.30.0.85        ge-0/0/4.135       172.30.5.3:0             14
```

```
[edit]
lab@R5# run show route table inet.3

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.3/32      *[LDP/9] 00:06:38, metric 10
                    > to 172.30.0.85 via ge-0/0/4.135
172.30.5.4/32      *[LDP/9] 00:05:04, metric 15
                    > to 172.30.0.94 via ae0.0, Push 299792
                       to 172.30.0.85 via ge-0/0/4.135, Push 299776
172.30.5.6/32      *[LDP/9] 00:05:04, metric 5
                    > to 172.30.0.94 via ae0.0
```

```
[edit]
lab@R5# run show ldp database
Input label database, 172.30.5.5:0--172.30.5.3:0
  Label     Prefix
      3     172.30.5.3/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 299776     172.30.5.4/32
 299792     172.30.5.5/32

Output label database, 172.30.5.5:0--172.30.5.3:0
  Label     Prefix
 299776     172.30.5.3/32
 299792     172.30.5.4/32
      3     172.30.5.5/32
 299808     172.30.5.6/32

Input label database, 172.30.5.5:0--172.30.5.6:0
  Label     Prefix
 299792     172.30.5.4/32
 299776     172.30.5.5/32
      3     172.30.5.6/32

Output label database, 172.30.5.5:0--172.30.5.6:0
  Label     Prefix
 299776     172.30.5.3/32
 299792     172.30.5.4/32
      3     172.30.5.5/32
 299808     172.30.5.6/32
```

Verify the operation of LDP on R6.

```
[edit]
lab@R6# run show ldp session
  Address            State        Connection     Hold time
172.30.5.4           Operational  Open             29
172.30.5.5           Operational  Open             29
```

```
[edit]
lab@R6# run show ldp neighbor
Address            Interface          Label space ID         Hold time
172.30.0.93        ae0.0              172.30.5.5:0             12
172.30.0.89        ge-0/0/4.146       172.30.5.4:0             12
```

```
[edit]
lab@R6# run show route table inet.3

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.4/32      *[LDP/9] 00:05:21, metric 10
                    > to 172.30.0.89 via ge-0/0/4.146
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.5/32      *[LDP/9] 00:05:21, metric 5
                    > to 172.30.0.93 via ae0.0
```

R3's loopback IP address is missing in **inet.3**.

```
[edit]
lab@R6# run show ldp database
Input label database, 172.30.5.6:0--172.30.5.4:0
  Label     Prefix
 299776     172.30.5.3/32
      3     172.30.5.4/32
 299792     172.30.5.5/32
 299808     172.30.5.6/32

Output label database, 172.30.5.6:0--172.30.5.4:0
  Label     Prefix
 299792     172.30.5.4/32
 299776     172.30.5.5/32
      3     172.30.5.6/32

Input label database, 172.30.5.6:0--172.30.5.5:0
  Label     Prefix
 299776     172.30.5.3/32
 299792     172.30.5.4/32
      3     172.30.5.5/32
 299808     172.30.5.6/32

Output label database, 172.30.5.6:0--172.30.5.5:0
  Label     Prefix
 299792     172.30.5.4/32
 299776     172.30.5.5/32
      3     172.30.5.6/32
```

Router R4 and R5 advertise a label to R6 for R3's loopback IP address (FEC). Let's look at the `inet.0` routing table on R6 for R3s loopback IP address.

```
[edit]
lab@R6# run show route 172.30.5.3/32

inet.0: 562 destinations, 580 routes (562 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.3/32      *[OSPF/10] 00:24:06, metric 20
                      to 172.30.0.46 via ge-0/0/4.167
                    > to 172.30.0.22 via ge-0/0/4.168
```


© 2020 Juniper Networks, Inc. All rights reserved.

The reason why the route is not reachable via LDP is because the IPv4 route for that destination is active in the `inet.0` table through a different interface. An LSR should not use a label if the exact same route is not present in the routing table.

The output below is taken to verify R4's loopback IP address. You can see that the LDP next-hop follows the IGP route in the **`inet.0`** routing table.

```
[edit]
lab@R6# run show route 172.30.5.4/32

inet.0: 562 destinations, 580 routes (562 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.4/32       *[OSPF/10] 03:56:48, metric 10
                     > to 172.30.0.89 via ge-0/0/4.146

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.4/32       *[LDP/9] 00:07:07, metric 10
                     > to 172.30.0.89 via ge-0/0/4.146
```

2) Configure the RSVP interfaces and coloring as described in Table 30. Enable RSVP message aggregation and path MTU discovery. Configure MD5 authentication for all RSVP sessions.

Router	**Interface**	**Color**
R1	ae0.0	green
	ge-0/0/4.117	blue
	ge-0/0/4.118	purple
R2	ae0.0	green
	ge-0/0/4.126	blue
	ge-0/0/4.123	purple
R3	ge-0/0/4.123	purple
	ge-0/0/4.138	blue
	ge-0/0/4.137	green
R6	ge-0/0/4.126	blue
	ge-0/0/4.167	purple
	ge-0/0/4.168	green
R7	ge-0/0/4.117	blue
	ge-0/0/4.137	green
	ge-0/0/4.167	purple

© 2020 Juniper Networks, Inc. All rights reserved.

R8	ge-0/0/4.118	purple
	ge-0/0/4.138	blue
	ge-0/0/4.168	green

Table 3040

The RSVP LSPs will be used in Area 0 of the backbone of the network. Constrained Shortest Path First is used to calculate the LSP path through the network. The `family mpls` should be configured on all interfaces which need to support MPLS. Additionally, the interfaces should be added to `protocols mpls` hierarchy. This configuration was performed in the previous step.

a) Configure RSVP on R1.

Configure MTU signaling for RSVP to discover the Maximum Transmission Unit size for the LSP.

```
[edit]
lab@R1# set protocols mpls path-mtu rsvp mtu-signaling
```

Configure the admin-groups.

```
[edit]
lab@R1# set protocols mpls admin-groups green 0

[edit]
lab@R1# set protocols mpls admin-groups purple 1

[edit]
lab@R1# set protocols mpls admin-groups blue 2
```

Assigning the interfaces to the respected color admin-groups according to table 42.

```
[edit]
lab@R1# set protocols mpls interface ae0.0 admin-group green

[edit]
lab@R1# set protocols mpls interface ge-0/0/4.117 admin-group blue

[edit]
lab@R1# set protocols mpls interface ge-0/0/4.118 admin-group purple
```

Configure 20% oversubscription on all RSVP interfaces, except for the Aggregated Ethernet interfaces.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set protocols rsvp interface ge-0/0/4.117 subscription 120
```

Configure MD5 authentication for all RSVP interfaces.

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.117 authentication-key workbook
```

RSVP aggregation is packing multiple RSVP message into a single message.

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.117 aggregate

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.118 subscription 120

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.118 authentication-key workbook

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.118 aggregate

[edit]
lab@R1# set protocols rsvp interface ae0.0 authentication-key workbook

[edit]
lab@R1# set protocols rsvp interface ae0.0 aggregate
```

Unlike the IS-IS protocol, OSPF by default does not exchange opaque LSA Type 10 used for traffic engineering. You have to explicitly enable **traffic-engineering** under OSPF to populate the Traffic engineering Database (TED) on each router.

```
[edit]
lab@R1# set protocols ospf traffic-engineering

[edit]
lab@R1# commit
commit complete
```

Verify the operation of the RSVP protocol on R1.

```
[edit]
lab@R1# run show mpls interface
Interface         State       Administrative groups (x: extended)
ae0.0             Up          green
ge-0/0/4.117      Up          blue
ge-0/0/4.118      Up          purple
ge-0/0/4.206      Up          <none>
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

RSVP is soft state protocol therefore sessions will only be established when a request for an LSP is initiated.

```
[edit]
lab@R1# run show rsvp interface extensive
RSVP interface: 3 active
ae0.0 Index 69, State Ena/Up
  Authentication, Aggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.30.0.1
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 2Gbps
  ct0: StaticBW 2Gbps, AvailableBW 2Gbps
    MaxAvailableBW 2Gbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off

ge-0/0/4.117 Index 73, State Ena/Up
  Authentication, Aggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.30.0.5
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 120%,
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 1.2Gbps
    MaxAvailableBW 1.2Gbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  ...
```

b) Configure RSVP on R2.

```
[edit]
lab@R2# set protocols mpls path-mtu rsvp mtu-signaling
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set protocols mpls admin-groups green 0

[edit]
lab@R2# set protocols mpls admin-groups purple 1

[edit]
lab@R2# set protocols mpls admin-groups blue 2

[edit]
lab@R2# set protocols mpls interface ae0.0 admin-group green

[edit]
lab@R2# set protocols mpls interface ge-0/0/4.123 admin-group purple

[edit]
lab@R2# set protocols mpls interface ge-0/0/4.126 admin-group blue

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.123 subscription 120

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.123 authentication-key workbook

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.123 aggregate

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.126 subscription 120

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.126 authentication-key workbook

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.126 aggregate

[edit]
lab@R2# set protocols rsvp interface ae0.0 authentication-key workbook

[edit]
lab@R2# set protocols rsvp interface ae0.0 aggregate

[edit]
lab@R2# set protocols ospf traffic-engineering

[edit]
lab@R2# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify the operation of RSVP on R2.

```
[edit]
lab@R2# run show rsvp interface
RSVP interface: 3 active
                    Active Subscr- Static      Available   Reserved    Highwater
Interface   State   resv   iption  BW          BW          BW          mark
ae0.0       Up         0   100%    2Gbps       2Gbps       0bps        0bps
ge-0/0/4.123    Up          0   120%   1000Mbps    1.2Gbps     0bps        0bps
ge-0/0/4.126    Up          0   120%   1000Mbps    1.2Gbps     0bps        0bps
```

```
[edit]
lab@R2# run show mpls interface
Interface           State       Administrative groups (x: extended)
ae0.0               Up          green
ge-0/0/4.123        Up          purple
ge-0/0/4.126        Up          blue
ge-0/0/4.207        Up          <none>
```

c) Configure RSVP on R3.

```
[edit]
lab@R3# set protocols mpls path-mtu rsvp mtu-signaling

[edit]
lab@R3# set protocols mpls admin-groups green 0

[edit]
lab@R3# set protocols mpls admin-groups purple 1

[edit]
lab@R3# set protocols mpls admin-groups blue 2

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.123 admin-group purple

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.137 admin-group green

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.138 admin-group blue

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.123 subscription 120

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set protocols rsvp interface ge-0/0/4.123 authentication-key workbook

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.123 aggregate

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.137 subscription 120

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.137 authentication-key workbook

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.137 aggregate

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.138 subscription 120

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.138 authentication-key workbook

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.138 aggregate

[edit]
lab@R3# set protocols ospf traffic-engineering

[edit]
lab@R3# commit
commit complete
```

Verify the operation of the RSVP protocol on R3.

```
[edit]
lab@R3# run show mpls interface
Interface         State       Administrative groups (x: extended)
ae0.0             Up          <none>
ge-0/0/4.123      Up           purple
ge-0/0/4.135      Up          <none>
ge-0/0/4.137      Up           green
ge-0/0/4.138      Up           blue
```

d) Configure RSVP on R6.

```
[edit]
lab@R6# set protocols mpls path-mtu rsvp mtu-signaling

[edit]
lab@R6# set protocols mpls admin-groups green 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols mpls admin-groups purple 1

[edit]
lab@R6# set protocols mpls admin-groups blue 2

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.126 admin-group blue

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.167 admin-group purple

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.168 admin-group green

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.126 subscription 120

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.126 authentication-key workbook

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.126 aggregate

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.167 authentication-key workbook

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.167 subscription 120

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.167 aggregate

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.168 authentication-key workbook

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.168 subscription 120

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.168 aggregate

[edit]
lab@R6# set protocols ospf traffic-engineering

[edit]
lab@R6# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify the operation of the RSVP protocol on R6.

```
[edit]
lab@R6# run show mpls interface
Interface          State       Administrative groups (x: extended)
ae0.0              Up          <none>
ge-0/0/4.126       Up           blue
ge-0/0/4.146       Up          <none>
ge-0/0/4.167       Up           purple
ge-0/0/4.168       Up           green
```

e) Configure RSVP on R7.

On R7 no LDP interfaces are present. Only RSVP LSPs are signaled to the other PE routers. The core facing interfaces should be added to the **protocols mpls** hierarchy.

```
[edit]
lab@R7# set interfaces ge-0/0/4.117 family mpls

[edit]
lab@R7# set interfaces ge-0/0/4.137 family mpls

[edit]
lab@R7# set interfaces ge-0/0/4.167 family mpls

[edit]
lab@R7# set protocols mpls path-mtu rsvp mtu-signaling

[edit]
lab@R7# set protocols mpls admin-groups green 0

[edit]
lab@R7# set protocols mpls admin-groups purple 1

[edit]
lab@R7# set protocols mpls admin-groups blue 2

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.117 admin-group blue

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.137 admin-group green

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.167 admin-group purple

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.117 subscription 120
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.117 authentication-key workbook

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.117 aggregate

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.137 authentication-key workbook

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.137 subscription 120

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.137 aggregate

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.167 authentication-key workbook

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.167 subscription 120

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.167 aggregate

[edit]
lab@R7# set protocols ospf traffic-engineering

[edit]
lab@R7# commit
commit complete
```

Verify the operation of the RSVP protocol on R7.

```
[edit]
lab@R7# run show rsvp interface
RSVP interface: 3 active
                  Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
ge-0/0/4.117Up          0   120%  1000Mbps    1.2Gbps     0bps        0bps
ge-0/0/4.137Up          0   120%  1000Mbps    1.2Gbps     0bps        0bps
ge-0/0/4.167Up          0   120%  1000Mbps    1.2Gbps     0bps        0bps
```

```
[edit]
lab@R7# run show mpls interface
Interface         State       Administrative groups (x: extended)
ge-0/0/4.117      Up          blue
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ge-0/0/4.137      Up            green
ge-0/0/4.167      Up            purple
```

f) Configure RSVP on R8.

```
[edit]
lab@R8# set interfaces ge-0/0/4.118 family mpls

[edit]
lab@R8# set interfaces ge-0/0/4.138 family mpls

[edit]
lab@R8# set interfaces ge-0/0/4.168 family mpls

[edit]
lab@R8# set protocols mpls path-mtu rsvp mtu-signaling

[edit]
lab@R8# set protocols mpls admin-groups green 0

[edit]
lab@R8# set protocols mpls admin-groups purple 1

[edit]
lab@R8# set protocols mpls admin-groups blue 2

[edit]
lab@R8# set protocols mpls interface ge-0/0/4.118 admin-group purple

[edit]
lab@R8# set protocols mpls interface ge-0/0/4.138 admin-group blue

[edit]
lab@R8# set protocols mpls interface ge-0/0/4.168 admin-group green

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.118 subscription 120

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.118 authentication-key workbook

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.118 aggregate

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.138 subscription 120
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.138 authentication-key workbook

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.138 aggregate

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.168 authentication-key workbook

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.168 subscription 120

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.168 aggregate

[edit]
lab@R8# set protocols ospf traffic-engineering

[edit]
lab@R8# commit
commit complete
```

Verify the operation of the RSVP protocol on R8.

```
[edit]
lab@R8# run show mpls interface
Interface          State       Administrative groups (x: extended)
ge-0/0/4.118       Up          purple
ge-0/0/4.138       Up          blue
ge-0/0/4.168       Up          green
```

3) Configure full mesh of LSPs in the core between the routers except for R4, R5 and the route reflector.

a) Configure RSVP LSPs on R1.

Configure a full mesh of LSP from R1 using CSPF to constrain the paths (bandwidth/ link colors) as defined in the task requirements.

First let's check if the TED is populated so that CSPF can calculate the paths.

```
[edit]
lab@R1# run show ted database
TED database: 0 ISIS nodes 6 INET nodes
ID                              Type Age(s) LnkIn LnkOut Protocol
172.30.5.1                      Rtr      40     3      3 OSPF(0.0.0.0)
    To: 172.30.5.2, Local: 172.30.0.1, Remote: 172.30.0.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
      Local interface index: 69, Remote interface index: 0
    To: 172.30.5.7, Local: 172.30.0.5, Remote: 172.30.0.6
      Local interface index: 73, Remote interface index: 0
    To: 172.30.5.8, Local: 172.30.0.9, Remote: 172.30.0.10
      Local interface index: 74, Remote interface index: 0
ID                           Type Age(s) LnkIn LnkOut Protocol
172.30.5.2                   Rtr      33     3      3 OSPF(0.0.0.0)
    To: 172.30.5.1, Local: 172.30.0.2, Remote: 172.30.0.1
      Local interface index: 69, Remote interface index: 0
    To: 172.30.5.3, Local: 172.30.0.13, Remote: 172.30.0.14
      Local interface index: 75, Remote interface index: 0
    To: 172.30.5.6, Local: 172.30.0.17, Remote: 172.30.0.18
      Local interface index: 76, Remote interface index: 0
ID                           Type Age(s) LnkIn LnkOut Protocol
172.30.5.3                   Rtr      30     3      3 OSPF(0.0.0.0)
    To: 172.30.5.2, Local: 172.30.0.14, Remote: 172.30.0.13
      Local interface index: 74, Remote interface index: 0
    To: 172.30.5.7, Local: 172.30.0.29, Remote: 172.30.0.30
      Local interface index: 76, Remote interface index: 0
    To: 172.30.5.8, Local: 172.30.0.33, Remote: 172.30.0.34
      Local interface index: 77, Remote interface index: 0
ID                           Type Age(s) LnkIn LnkOut Protocol
172.30.5.6                   Rtr      17     3      3 OSPF(0.0.0.0)
    To: 172.30.5.2, Local: 172.30.0.18, Remote: 172.30.0.17
      Local interface index: 73, Remote interface index: 0
    To: 172.30.5.7, Local: 172.30.0.45, Remote: 172.30.0.46
      Local interface index: 75, Remote interface index: 0
    To: 172.30.5.8, Local: 172.30.0.21, Remote: 172.30.0.22
      Local interface index: 76, Remote interface index: 0
ID                           Type Age(s) LnkIn LnkOut Protocol
172.30.5.7                   Rtr      13     3      3 OSPF(0.0.0.0)
    To: 172.30.5.1, Local: 172.30.0.6, Remote: 172.30.0.5
      Local interface index: 70, Remote interface index: 0
    To: 172.30.5.3, Local: 172.30.0.30, Remote: 172.30.0.29
      Local interface index: 71, Remote interface index: 0
    To: 172.30.5.6, Local: 172.30.0.46, Remote: 172.30.0.45
      Local interface index: 72, Remote interface index: 0
ID                           Type Age(s) LnkIn LnkOut Protocol
172.30.5.8                   Rtr       9     3      3 OSPF(0.0.0.0)
    To: 172.30.5.1, Local: 172.30.0.10, Remote: 172.30.0.9
      Local interface index: 70, Remote interface index: 0
    To: 172.30.5.3, Local: 172.30.0.34, Remote: 172.30.0.33
      Local interface index: 71, Remote interface index: 0
    To: 172.30.5.6, Local: 172.30.0.22, Remote: 172.30.0.21
      Local interface index: 72, Remote interface index: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

R4 and R5 are missing, because OSPF traffic-engineering is not enabled on them, but even if they are configured the TE LSAs scope is limited only to a single Area. LDP tunneling will be used to connect Area 1 with the routers in Area 0.

Configure LSPs on R1.

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R2 to 172.30.5.2 bandwidth 100m
```

You have to increase the default hold priority of the LSP to ensure that it is higher than default value of 0. Set the hold priority of the LSP to 5. It is recommended that if you change the hold priority to also change the setup value to the same value.

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R2 priority 5 5

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R2 admin-group include-any [ green
blue ]

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R3 to 172.30.5.3

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R3 bandwidth 100m

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R3 priority 5 5

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R3 admin-group include-any [ green
blue ]

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 to 172.30.5.6

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 bandwidth 100m

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 priority 5 5

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 admin-group include-any [ green
blue ]

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set protocols mpls label-switched-path R1-to-R7 to 172.30.5.7

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R7 bandwidth 100m

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R7 priority 5 5

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R7 admin-group include-any [ green
blue ]

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 to 172.30.5.8

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 bandwidth 100m

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 priority 5 5

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 admin-group include-any [ green
blue ]

[edit]
lab@R1# commit
commit complete
```

List the MPLS LSP table.

```
[edit]
lab@R1# run show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.2      172.30.5.1      Up     0 *                      R1-to-R2
172.30.5.3      172.30.5.1      Up     0 *                      R1-to-R3
172.30.5.6      172.30.5.1      Up     0 *                      R1-to-R6
172.30.5.7      172.30.5.1      Up     0 *                      R1-to-R7
172.30.5.8      172.30.5.1      Up     0 *                      R1-to-R8
Total 5 displayed, Up 5, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

inet.3 is populated with the routes resolved via RSVP signaled LSPs.

```
[edit]
lab@R1# run show route table inet.3

inet.3: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.2/32       *[RSVP/7/1] 00:54:14, metric 11
                     > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R2
                     [LDP/9] 02:15:48, metric 11
                     > to 172.30.0.66 via ge-0/0/4.206, Push 100016
172.30.5.3/32       *[RSVP/7/1] 00:54:14, metric 21
                     > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R3
172.30.5.6/32       *[RSVP/7/1] 00:54:14, metric 30
                     > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R6
172.30.5.7/32       *[RSVP/7/1] 00:54:14, metric 12
                     > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R7
172.30.5.8/32       *[RSVP/7/1] 00:54:14, metric 10
                     > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R8
172.30.5.41/32      *[LDP/9] 02:18:48, metric 10
                     > to 172.30.0.66 via ge-0/0/4.206
```

Routes for R4 and R5 are not reachable via the LSP. R4 and R5 are in a different OSPF area and only LDP is enabled as the label signaling protocol.

b) Configure RSVP LSPs on R2.

You must configure two LSP from R2 to R1 and R8 using a different physical path. Those four LSPs should automatically adjust their bandwidth every 24 hours.

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first to 172.30.5.1

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first priority 5 5

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first admin-group include-any [
green blue ]

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first auto-bandwidth adjust-
interval 86400
```


© 2020 Juniper Networks, Inc. All rights reserved.

LSP bandwidth reservation should be in the range of minimum 50m and maximum of 100m.

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first auto-bandwidth minimum-
bandwidth 50m

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first auto-bandwidth maximum-
bandwidth 100m
```

The **auto-bandwidth** option should be enabled globally, and a file should be created to collect the bandwidth statistics.

```
[edit]
lab@R2# set protocols mpls statistics auto-bandwidth

[edit]
lab@R2# set protocols mpls statistics file mpls-bandwidth
```

The second LSP from R1 to R2 should use a different preconfigured path then the first LSP. Configure static ERO to manually select the path through the network.

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-second to 172.30.5.1

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-second priority 5 5

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-second primary path-1

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-second auto-bandwidth adjust-
interval 86400

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-second auto-bandwidth minimum-
bandwidth 50m

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-second auto-bandwidth maximum-
bandwidth 100m
```

Configure the static ERO used for the second LSPs.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set protocols mpls path path-1 172.30.5.3

[edit]
lab@R2# set protocols mpls path path-1 172.30.5.8

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R6 to 172.30.5.6

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R6 bandwidth 100m

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R6 priority 5 5

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R6 admin-group include-any [ green
blue ]

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 to 172.30.5.7

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 bandwidth 100m

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 priority 5 5

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 admin-group include-any [ green
blue ]

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-first to 172.30.5.8

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-first priority 5 5

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-first admin-group include-any [
green blue ]

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-first auto-bandwidth adjust-
interval 86400

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-first auto-bandwidth minimum-
bandwidth 50m
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-first auto-bandwidth maximum-
bandwidth 100m

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-second to 172.30.5.8

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-second priority 5 5

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-second primary path-1

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-second auto-bandwidth adjust-
interval 86400

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-second auto-bandwidth minimum-
bandwidth 50m

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-second auto-bandwidth maximum-
bandwidth 100m

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R3 to 172.30.5.3

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R3 bandwidth 100m

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R3 priority 5 5

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R3 admin-group include-any [ green
blue ]

[edit]
lab@R2# commit
commit complete
```

Verify the MPLS LSP table.

```
[edit]
lab@R2# run show mpls lsp
Ingress LSP: 7 sessions
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
To              From            State Rt P    ActivePath      LSPname
172.30.5.1      172.30.5.2      Up     0 *                    R2-to-R1-first
172.30.5.1      172.30.5.2      Up     0 *    path-1          R2-to-R1-second
172.30.5.3      172.30.5.2      Up     0 *                    R2-to-R3
172.30.5.6      172.30.5.2      Up     0 *                    R2-to-R6
172.30.5.7      172.30.5.2      Up     0 *                    R2-to-R7
172.30.5.8      172.30.5.2      Up     0 *                    R2-to-R8-first
172.30.5.8      172.30.5.2      Up     0 *    path-1          R2-to-R8-second
Total 7 displayed, Up 7, Down 0

Egress LSP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.2      172.30.5.1      Up       0  1 FF       3         - R1-to-R2
Total 1 displayed, Up 1, Down 0

Transit LSP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.1      Up       0  1 FF  301024         3 R1-to-R6
172.30.5.8      172.30.5.1      Up       0  1 FF  301040    299904 R1-to-R8
Total 2 displayed, Up 2, Down 0
```

Verify the content of the **`inet.3`** routing table.

```
[edit]
lab@R2# run show route table inet.3

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       *[RSVP/7/1] 00:04:44, metric 5
                     > to 172.30.0.1 via ae0.0, label-switched-path R2-to-R1-first
                       to 172.30.0.14 via ge-0/0/4.123, label-switched-path R2-to-
R1-second
172.30.5.3/32       *[RSVP/7/1] 00:04:44, metric 10
                     > to 172.30.0.1 via ae0.0, label-switched-path R2-to-R3
172.30.5.6/32       *[RSVP/7/1] 00:04:44, metric 10
                     > to 172.30.0.18 via ge-0/0/4.126, label-switched-path R2-to-
R6
172.30.5.7/32       *[RSVP/7/1] 00:04:44, metric 15
                     > to 172.30.0.1 via ae0.0, label-switched-path R2-to-R7
172.30.5.8/32       *[RSVP/7/1] 00:04:44, metric 15
                       to 172.30.0.18 via ge-0/0/4.126, label-switched-path R2-to-
R8-first
                     > to 172.30.0.14 via ge-0/0/4.123, label-switched-path R2-to-
R8-second
172.30.5.41/32      *[LDP/9] 15:16:16, metric 10
                     > to 172.30.0.70 via ge-0/0/4.207
```


© 2020 Juniper Networks, Inc. All rights reserved.

c) Configure RSVP LSPs on R3.

You must configure all of the LSP originated from R3 with a higher setup priority. Set the setup and hold priorities to 3.

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 to 172.30.5.1

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 bandwidth 10m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 priority 3 3

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1admin-group include-any [
green blue ]

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 to 172.30.5.6

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 priority 3 3

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6admin-group include-any [
green blue ]

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 to 172.30.5.7

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 priority 3 3

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 admin-group include-any [
green blue ]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 to 172.30.5.8

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 priority 3 3

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 admin-group include-any [
green blue ]
```

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 to 172.30.5.2

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 priority 3 3

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 admin-group include-any [
green blue ]

[edit]
lab@R3# commit
commit complete
```

Verify the MPLS LSP table.

```
[edit]
lab@R3# run show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.1      172.30.5.3      Up     0 *                      R3-to-R1
172.30.5.2      172.30.5.3      Up     0 *                      R3-to-R2
172.30.5.6      172.30.5.3      Up     0 *                      R3-to-R6
172.30.5.7      172.30.5.3      Up     0 *                      R3-to-R7
172.30.5.8      172.30.5.3      Up     0 *                      R3-to-R8
Total 5 displayed, Up 5, Down 0

Egress LSP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.3      172.30.5.1      Up       0  1 FF       3         - R1-to-R3
172.30.5.3      172.30.5.2      Up       0  1 FF       3         - R2-to-R3
Total 2 displayed, Up 2, Down 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Verify the content of the **inet.3** routing table.

```
[edit]
lab@R3# run show route table inet.3

inet.3: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       *[RSVP/7/1] 00:16:29, metric 20
                     > to 172.30.0.30 via ge-0/0/4.137, label-switched-path R3-to-
R1
172.30.5.2/32       *[RSVP/7/1] 00:16:29, metric 10
                     > to 172.30.0.30 via ge-0/0/4.137, label-switched-path R3-to-
R2
172.30.5.4/32       *[LDP/9] 02:06:13, metric 5
                     > to 172.30.0.82 via ae0.0
172.30.5.5/32       *[LDP/9] 02:01:27, metric 10
                     > to 172.30.0.86 via ge-0/0/4.135
172.30.5.6/32       *[RSVP/7/1] 00:16:29, metric 20
                     > to 172.30.0.34 via ge-0/0/4.138, label-switched-path R3-to-
R6
172.30.5.7/32       *[RSVP/7/1] 00:16:29, metric 10
                     > to 172.30.0.30 via ge-0/0/4.137, label-switched-path R3-to-
R7
172.30.5.8/32       *[RSVP/7/1] 00:16:29, metric 10
                     > to 172.30.0.34 via ge-0/0/4.138, label-switched-path R3-to-
R8
```

On R3 all loopback IP addresses are reachable through LSPs, except the one from the route reflector. R3 is the OSPF ABR router between the Area 0 (RSVP) and Area 1 (LDP).

d) Configure RSVP LSPs on R6.

You have to configure all of the LSP originated from R6 with a higher setup priority. Set the setup and hold priorities to 3.

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 to 172.30.5.1

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 bandwidth 100m
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 priority 3 3

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 admin-group include-any [
purple blue ]


[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 to 172.30.5.2

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 bandwidth 100m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 priority 3 3

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 admin-group include-any [
purple blue ]

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 to 172.30.5.3

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 bandwidth 100m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 priority 3 3

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 admin-group include-any [
purple blue ]

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 to 172.30.5.7

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 bandwidth 100m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 priority 3 3

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 admin-group include-any [
purple blue ]

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 to 172.30.5.8
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 bandwidth 100m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 priority 3 3


[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 admin-group include-any [
purple blue ]

[edit]
lab@R6# commit
commit complete
```

Verify the MPLS LSP table.

```
[edit]
lab@R6# run show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.1      172.30.5.6      Up     0 *                      R6-to-R1
172.30.5.2      172.30.5.6      Up     0 *                      R6-to-R2
172.30.5.3      172.30.5.6      Up     0 *                      R6-to-R3
172.30.5.7      172.30.5.6      Up     0 *                      R6-to-R7
172.30.5.8      172.30.5.6      Up     0 *                      R6-to-R8
Total 5 displayed, Up 5, Down 0

Egress LSP: 3 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.1      Up       0  1 FF       3         - R1-to-R6
172.30.5.6      172.30.5.3      Up       0  1 FF       3         - R3-to-R6
172.30.5.6      172.30.5.2      Up       0  1 FF       3         - R2-to-R6
Total 3 displayed, Up 3, Down 0

Transit LSP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.8      172.30.5.1      Up       0  1 FF  299808         3 R1-to-R8
172.30.5.8      172.30.5.2      Up       0  1 FF  299824         3 R2-to-R8-first
Total 2 displayed, Up 2, Down 0
```

Verify the content of the **inet.3** routing table.

```
[edit]
lab@R6# run show route table inet.3

inet.3: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

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

172.30.5.1/32       *[RSVP/7/1] 00:10:55, metric 20
                     > to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R1
172.30.5.2/32       *[RSVP/7/1] 00:10:55, metric 10
                     > to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R2
172.30.5.3/32       *[RSVP/7/1] 00:10:55, metric 20
                     > to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R3
172.30.5.4/32       *[LDP/9] 01:59:23, metric 10
                     > to 172.30.0.89 via ge-0/0/4.146
172.30.5.5/32       *[LDP/9] 01:59:23, metric 5
                     > to 172.30.0.93 via ae0.0
172.30.5.7/32       *[RSVP/7/1] 00:10:55, metric 10
                     > to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R7
172.30.5.8/32       *[RSVP/7/1] 00:10:55, metric 10
                     > to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R8
```

e) Configure RSVP LSPs on R7.

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R1 to 172.30.5.1

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R1 bandwidth 100m

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R1 priority 5 5

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R1 admin-group include-any [
purple blue ]

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 to 172.30.5.2

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 bandwidth 100m

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 priority 5 5

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# set protocols mpls label-switched-path R7-to-R2 admin-group include-any [
purple blue ]

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R3 to 172.30.5.3

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R3 bandwidth 100m

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R3 priority 5 5

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R3 admin-group include-any [
purple blue ]

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R6 to 172.30.5.6

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R6 bandwidth 100m

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R6 priority 5 5

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R6 admin-group include-any [
purple blue ]

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R8 to 172.30.5.8

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R8 bandwidth 100m

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R8 priority 5 5

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R8 admin-group include-any [
purple blue ]

[edit]
lab@R7# commit
commit complete
```

Verify the MPLS LSP table.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.1      172.30.5.7      Up     0 *                      R7-to-R1
172.30.5.2      172.30.5.7      Up     0 *                      R7-to-R2
172.30.5.3      172.30.5.7      Up     0 *                      R7-to-R3
172.30.5.6      172.30.5.7      Up     0 *                      R7-to-R6
172.30.5.8      172.30.5.7      Up     0 *                      R7-to-R8
Total 5 displayed, Up 5, Down 0

Egress LSP: 4 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.1      Up       0  1 FF       3         - R1-to-R7
172.30.5.7      172.30.5.3      Up       0  1 FF       3         - R3-to-R7
172.30.5.7      172.30.5.2      Up       0  1 FF       3         - R2-to-R7
172.30.5.7      172.30.5.6      Up       0  1 FF       3         - R6-to-R7
Total 4 displayed, Up 4, Down 0

Transit LSP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.3      Up       0  1 FF  299808         3 R3-to-R1
172.30.5.1      172.30.5.6      Up       0  1 FF  299840         3 R6-to-R1
172.30.5.2      172.30.5.3      Up       0  1 FF  299824    299840 R3-to-R2
172.30.5.3      172.30.5.1      Up       0  1 FF  299776         3 R1-to-R3
172.30.5.3      172.30.5.2      Up       0  1 FF  299792         3 R2-to-R3
172.30.5.8      172.30.5.6      Up       0  1 FF  299856    299856 R6-to-R8
Total 6 displayed, Up 6, Down 0
```

Verify the content of the **`inet.3`** routing table.

```
[edit]
lab@R7# run show route table inet.3

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       *[RSVP/7/1] 00:06:04, metric 10
                     > to 172.30.0.5 via ge-0/0/4.117, label-switched-path R7-to-
R1
172.30.5.2/32       *[RSVP/7/1] 00:06:03, metric 20
                     > to 172.30.0.45 via ge-0/0/4.167, label-switched-path R7-to-
R2
172.30.5.3/32       *[RSVP/7/1] 00:06:03, metric 10
                     > to 172.30.0.45 via ge-0/0/4.167, label-switched-path R7-to-
R3
172.30.5.6/32       *[RSVP/7/1] 00:06:04, metric 10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.45 via ge-0/0/4.167, label-switched-path R7-to-
R6
172.30.5.8/32       *[RSVP/7/1] 00:06:03, metric 20
                    > to 172.30.0.5 via ge-0/0/4.117, label-switched-path R7-to-
R8
```

f) Configure RSVP LSPs on R8.

```
[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 to 172.30.5.1

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 bandwidth 100m

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 priority 5 5

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 admin-group include-any [
purple blue ]

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R2 to 172.30.5.2

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R2 bandwidth 100m

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R2 priority 5 5

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R2 admin-group include-any [
purple blue ]

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3 to 172.30.5.3

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3 bandwidth 100m

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3 priority 5 5

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3 admin-group include-any [
purple blue ]

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# set protocols mpls label-switched-path R8-to-R6 to 172.30.5.6

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R6 bandwidth 100m

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R6 priority 5 5

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R6 admin-group include-any [
purple blue ]

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R7 to 172.30.5.7

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R7 bandwidth 100m

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R7 priority 5 5

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R7 admin-group include-any [
purple blue ]

[edit]
lab@R8# commit
commit complete
```

Verify the MPLS LSP table.

```
[edit]
lab@R8# run show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P     ActivePath       LSPname
172.30.5.1      172.30.5.8      Up     0 *                      R8-to-R1
172.30.5.2      172.30.5.8      Up     0 *                      R8-to-R2
172.30.5.3      172.30.5.8      Up     0 *                      R8-to-R3
172.30.5.6      172.30.5.8      Up     0 *                      R8-to-R6
172.30.5.7      172.30.5.8      Up     0 *                      R8-to-R7
Total 5 displayed, Up 5, Down 0

Egress LSP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.8      172.30.5.6      Up       0  1 FF       3         - R6-to-R8
172.30.5.8      172.30.5.7      Up       0  1 FF       3         - R7-to-R8
172.30.5.8      172.30.5.2      Up       0  1 FF       3         - R2-to-R8-first
172.30.5.8      172.30.5.2      Up       0  1 FF       3         - R2-to-R8-second
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.8      172.30.5.1      Up         0  1 FF       3         - R1-to-R8
172.30.5.8      172.30.5.3      Up         0  1 FF       3         - R3-to-R8 Total
6 displayed, Up 6, Down 0

Transit LSP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.3      Up       0  1 FF  299776        3 R3-to-R6
Total 1 displayed, Up 1, Down 0
```

Verify the content of the **`inet.3`** routing table.

```
[edit]
lab@R8# run show route table inet.3

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       *[RSVP/7/1] 00:01:32, metric 10
                     > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R1
172.30.5.2/32       *[RSVP/7/1] 00:01:32, metric 20
                     > to 172.30.0.33 via ge-0/0/4.138, label-switched-path R8-to-
R2
172.30.5.3/32       *[RSVP/7/1] 00:01:32, metric 10
                     > to 172.30.0.33 via ge-0/0/4.138, label-switched-path R8-to-
R3
172.30.5.6/32       *[RSVP/7/1] 00:01:31, metric 10
                     > to 172.30.0.33 via ge-0/0/4.138, label-switched-path R8-to-
R6
172.30.5.7/32       *[RSVP/7/1] 00:01:32, metric 20
                     > to 172.30.0.9 via ge-0/0/4.118, label-switched-path R8-to-
R7
```

4) Configure LSPs traffic protection on all routers.

a) Configure LSP traffic protection on R1.

You must configure backup protection paths for each LSP. You can create two paths, one primary and one secondary, both using CSPF.

```
[edit]
lab@R1# set protocols mpls path primary-path

[edit]
lab@R1# set protocols mpls path secondary-path

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set protocols mpls label-switched-path R1-to-R2 primary primary-path

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R2 secondary secondary-path

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R3 primary primary-path

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 primary primary-path


[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R7 primary primary-path

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 primary primary-path

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R3 secondary secondary-path

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 secondary secondary-path

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R7 secondary secondary-path

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 secondary secondary-path
```

RSVP link-protection should be configured on the LSPs level and on all the RSVP enabled interfaces.

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.117 link-protection

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.118 link-protection

[edit]
lab@R1# set protocols rsvp interface ae0.0 link-protection

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R2 link-protection

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R3 link-protection
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 link-protection

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R7 link-protection

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 link-protection

[edit]
lab@R1# commit
commit complete
```

Verify the MPLS LSP table to verify that the primary paths are used for all LSPs.

```
[edit]
lab@R1# run show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P    ActivePath       LSPname
172.30.5.2      172.30.5.1      Up     0 *    primary-path     R1-to-R2
172.30.5.3      172.30.5.1      Up     0 *    primary-path     R1-to-R3
172.30.5.6      172.30.5.1      Up     0 *    primary-path     R1-to-R6
172.30.5.7      172.30.5.1      Up     0 *    primary-path     R1-to-R7
172.30.5.8      172.30.5.1      Up     0 *    primary-path     R1-to-R8
Total 5 displayed, Up 5, Down 0

Egress LSP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.8      Up       0  1 FF       3         - R8-to-R1
172.30.5.1      172.30.5.3      Up       0  1 FF       3         - R3-to-R1
172.30.5.1      172.30.5.7      Up       0  1 FF       3         - R7-to-R1
172.30.5.1      172.30.5.2      Up       0  1 FF       3         - R2-to-R1-first
172.30.5.1      172.30.5.2      Up       0  1 FF       3         - R2-to-R1-second
172.30.5.1      172.30.5.6      Up       0  1 FF       3         - R6-to-R1
Total 6 displayed, Up 6, Down 0

Transit LSP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.8      Up       0  1 FF  301088        3 R8-to-R7
172.30.5.8      172.30.5.7      Up       0  1 FF  301072        3 R7-to-R8
Total 2 displayed, Up 2, Down 0
```

Bypass LSPs are established to protect the LSPs in case of failure.

```
[edit]
lab@R1# run show rsvp session | no-more
Ingress RSVP: 7 sessions
To              From            State   Rt Style Labelin Labelout LSPname
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.2      172.30.5.1      Up       0  1 SE       -         3 R1-to-R2
172.30.5.2      172.30.5.1      Up       0  1 SE       -    299904 Bypass-
>172.30.0.2
172.30.5.3      172.30.5.1      Up       0  1 SE       -    299920 R1-to-R3
172.30.5.6      172.30.5.1      Up       0  1 SE       -    301312 R1-to-R6
172.30.5.7      172.30.5.1      Up       0  1 SE       -         3 R1-to-R7
172.30.5.7      172.30.5.1      Up       0  1 SE       -    301344 Bypass-
>172.30.0.6
172.30.5.8      172.30.5.1      Up       0  1 SE       -    301328 R1-to-R8
Total 7 displayed, Up 7, Down 0

Egress RSVP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.8      Up       0  1 FF       3         - R8-to-R1
172.30.5.1      172.30.5.3      Up       0  1 FF       3         - R3-to-R1
172.30.5.1      172.30.5.7      Up       0  1 FF       3         - R7-to-R1
172.30.5.1      172.30.5.2      Up       0  1 FF       3         - R2-to-R1-first
172.30.5.1      172.30.5.2      Up       0  1 FF       3         - R2-to-R1-second
172.30.5.1      172.30.5.6      Up       0  1 FF       3         - R6-to-R1
Total 6 displayed, Up 6, Down 0

Transit RSVP: 2 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.8      Up       0  1 FF  301088        3 R8-to-R7
172.30.5.8      172.30.5.7      Up       0  1 FF  301072        3 R7-to-R8
Total 2 displayed, Up 2, Down 0
```

A bypass LSP is signaled on each interface.

```
[edit]
lab@R1# run show rsvp interface extensive
RSVP interface: 3 active
ae0.0 Index 328, State Ena/Up
  Authentication, Aggregate, Reliable, LinkProtection
  HelloInterval 9(second)
  Address 172.30.0.1
  ActiveResv 4, PreemptionCnt 0, Update threshold 10%, MaxResvTh 0bps, 0%
  Subscription 100%, Actual 100%
  bc0 = ct0, StaticBW 2Gbps
  ct0: StaticBW 2Gbps, AvailableBW 1.7Gbps
    MaxAvailableBW 2Gbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 300Mbps[6] 0bps[7]
0bps
  Protection: On, Bypass: 1, LSP: 3, Protected LSP: 3, Unprotected LSP: 0
      1 Apr  9 10:19:19 New bypass Bypass->172.30.0.2
    Bypass: Bypass->172.30.0.2, State: Up, Type: LP, LSP: 3, Backup: 0
      4 Apr  9 10:19:21 Up
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
      3 Apr  9 10:19:21 Record Route:  172.30.0.10(Label=299904)
172.30.0.33(Label=300224) 172.30.0.13(Label=3)
      2 Apr  9 10:19:20 CSPF: computation result accepted
      1 Apr  9 10:19:19 Originate Call
...
```

b) Configure LSP traffic protection on R2.

```
[edit]
lab@R2# set protocols mpls path primary-path

[edit]
lab@R2# set protocols mpls path secondary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R3 primary primary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R6 primary primary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 primary primary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first primary primary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-first primary primary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first secondary secondary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-second secondary secondary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-first secondary secondary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-second secondary secondary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R3 secondary secondary-path

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R6 secondary secondary-path

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# set protocols mpls label-switched-path R2-to-R7 secondary secondary-path
```

Note: LSP R2-to-R1-second and LSP R2-to-R8-second have already explicit primary paths configured in the previous step.

```
[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.123 link-protection

[edit]
lab@R2# set protocols rsvp interface ge-0/0/4.126 link-protection

[edit]
lab@R2# set protocols rsvp interface ae0.0 link-protection

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first link-protection

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-second link-protection

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R3 link-protection

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R6 link-protection

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R7 link-protection

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-first link-protection

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R8-second link-protection

[edit]
lab@R2# commit
commit complete
```

Verify that the primary paths are used for all LSPs.

```
[edit]
lab@R2# run show mpls lsp
Ingress LSP: 7 sessions
To              From            State Rt P    ActivePath      LSPname
172.30.5.1      172.30.5.2      Up     0 *    primary-path    R2-to-R1-first
172.30.5.1      172.30.5.2      Up     0 *    path-1          R2-to-R1-second
172.30.5.3      172.30.5.2      Up     0 *    primary-path    R2-to-R3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.6      172.30.5.2      Up     0 *     primary-path     R2-to-R6
172.30.5.7      172.30.5.2      Up     0 *     primary-path     R2-to-R7
172.30.5.8      172.30.5.2      Up     0 *     primary-path     R2-to-R8-first
172.30.5.8      172.30.5.2      Up     0 *     path-1           R2-to-R8-second
Total 7 displayed, Up 7, Down 0

Egress LSP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.2      172.30.5.7      Up       0  1 FF       3         - R7-to-R2
172.30.5.2      172.30.5.6      Up       0  1 FF       3         - R6-to-R2
172.30.5.2      172.30.5.8      Up       0  1 FF       3         - R8-to-R2
172.30.5.2      172.30.5.1      Up       0  1 SE       3         - R1-to-R2
172.30.5.2      172.30.5.1      Up       0  1 SE       3         - Bypass-
>172.30.0.2
172.30.5.2      172.30.5.3      Up       0  1 FF       3         - R3-to-R2
Total 6 displayed, Up 6, Down 0

Transit LSP: 7 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.3      172.30.5.6      Up       0  1 FF  301104         3 R6-to-R3
172.30.5.3      172.30.5.7      Up       0  1 FF  301136         3 R7-to-R3
172.30.5.6      172.30.5.8      Up       0  1 FF  301152         3 R8-to-R6
172.30.5.6      172.30.5.1      Up       0  1 SE  301312         3 R1-to-R6
172.30.5.7      172.30.5.1      Up       0  1 SE  301344    300240 Bypass-
>172.30.0.6
172.30.5.8      172.30.5.1      Up       0  1 SE  301328    300160 R1-to-R8
172.30.5.8      172.30.5.6      Up       0  1 FF  301120    300064 R6-to-R8
Total 7 displayed, Up 7, Down 0
```

Secondary paths are not established. They will come UP only in the case that the primary paths are not available.

```
[edit]
lab@R2# run show mpls lsp name R2-to-R1-first extensive
Ingress LSP: 7 sessions

172.30.5.1
  From: 172.30.5.2, State: Up, ActiveRoute: 0, LSPname: R2-to-R1-first
  ActivePath: primary-path (primary)
  Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Follow destination IGP metric
  Autobandwidth
  MinBW: 50Mbps, MaxBW: 100Mbps
  AdjustTimer: 86400 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 21530 second(s).
  Overflow limit: 0, Overflow sample count: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  LSP Self-ping Status : Enabled
 *Primary   primary-path     State: Up
    Priorities: 5 5
    Bandwidth: 50Mbps
    SmartOptimizeTimer: 180
          Include Any: blue green
    Flap Count: 0
    MBB Count: 0
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 5)
 172.30.0.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.5.1(flag=0x20) 172.30.0.1(Label=3)
    9 Apr  9 10:24:45.594 Link-protection Up
    8 Apr  9 10:24:42.673 Self-ping ended successfully
    7 Apr  9 10:24:41.713 Selected as active path
    6 Apr  9 10:24:41.712 Up
    5 Apr  9 10:24:41.712 Self-ping started
    4 Apr  9 10:24:41.712 Self-ping enqueued
    3 Apr  9 10:24:41.712 Record Route:  172.30.5.1(flag=0x20)
172.30.0.1(Label=3)
    2 Apr  9 10:24:41.532 Originate Call
    1 Apr  9 10:24:41.532 CSPF: computation result accepted  172.30.0.1
  Secondary secondary-path   State: Dn
    Priorities: 5 5
```

Verify the RSVP session table.

```
[edit]
lab@R2# run show rsvp session | no-more
Ingress RSVP: 10 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.2      Up       0  1 SE       -         3 R2-to-R1-first
172.30.5.1      172.30.5.2      Up       0  1 SE       -    300272 R2-to-R1-second
172.30.5.1      172.30.5.2      Up       0  1 SE       -    300288 Bypass-
>172.30.0.1
172.30.5.3      172.30.5.2      Up       0  1 SE       -    301264 R2-to-R3
172.30.5.3      172.30.5.2      Up       0  1 SE       -    301280 Bypass-
>172.30.0.14
172.30.5.6      172.30.5.2      Up       0  1 SE       -         3 R2-to-R6
172.30.5.6      172.30.5.2      Up       0  1 SE       -    301296 Bypass-
>172.30.0.18
172.30.5.7      172.30.5.2      Up       0  1 SE       -    301248 R2-to-R7
172.30.5.8      172.30.5.2      Up       0  1 SE       -    300176 R2-to-R8-first
172.30.5.8      172.30.5.2      Up       0  1 SE       -    300256 R2-to-R8-second
Total 10 displayed, Up 10, Down 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Egress RSVP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.2      172.30.5.7      Up       0  1 FF       3         - R7-to-R2
172.30.5.2      172.30.5.6      Up       0  1 FF       3         - R6-to-R2
172.30.5.2      172.30.5.8      Up       0  1 FF       3         - R8-to-R2
172.30.5.2      172.30.5.1      Up       0  1 SE       3         - R1-to-R2
172.30.5.2      172.30.5.1      Up       0  1 SE       3         - Bypass-
>172.30.0.2
172.30.5.2      172.30.5.3      Up       0  1 FF       3         - R3-to-R2
Total 6 displayed, Up 6, Down 0

Transit RSVP: 7 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.3      172.30.5.6      Up       0  1 FF  301104        3 R6-to-R3
172.30.5.3      172.30.5.7      Up       0  1 FF  301136        3 R7-to-R3
172.30.5.6      172.30.5.8      Up       0  1 FF  301152        3 R8-to-R6
172.30.5.6      172.30.5.1      Up       0  1 SE  301312        3 R1-to-R6
172.30.5.7      172.30.5.1      Up       0  1 SE  301344   300240 Bypass-
>172.30.0.6
172.30.5.8      172.30.5.1      Up       0  1 SE  301328   300160 R1-to-R8
172.30.5.8      172.30.5.6      Up       0  1 FF  301120   300064 R6-to-R8
Total 7 displayed, Up 7, Down 0
```

c) Configure LSP traffic protection on R3.

```
[edit]
lab@R3# set protocols mpls path primary-path

[edit]
lab@R3# set protocols mpls path secondary-path

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 primary primary-path

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 primary primary-path

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 primary primary-path

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 primary primary-path

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 primary primary-path
```


© 2020 Juniper Networks, Inc. All rights reserved.

On R3 the secondary path should be pre-established. This is accomplished with the `standby` knob configured under the secondary path.

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 secondary secondary-path
standby

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 secondary secondary-path
standby

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 secondary secondary-path
standby

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 secondary secondary-path
standby

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 secondary secondary-path
standby
```

The `adaptive` knob is configured to change the reservation style of the LSPs from fixed filter to shared explicit. This will reserve the bandwidth only one time if both paths of a given LSP transit the same link.

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 adaptive

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 adaptive

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 adaptive

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 adaptive

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 adaptive
```

The LSPs originated from R3 should use fast reroute protection. Link-protection still has to be configured on all RSVP enabled interfaces to protect the LSPs that transit these interfaces.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.123 link-protection

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.137 link-protection

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.138 link-protection
```

You have to make sure that the bandwidth is inherited by the fast reroute detours, but without the administrative groups settings. The easiest way is to remove any of the constraints and add the bandwidth reservation again for the detours.

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 fast-reroute bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 fast-reroute no-include-any

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 fast-reroute bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 fast-reroute no-include-any

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 fast-reroute bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R6 fast-reroute no-include-any

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 fast-reroute bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 fast-reroute no-include-any

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 fast-reroute bandwidth 100m

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 fast-reroute no-include-any

[edit]
lab@R3# commit
commit complete
```

Verify that the primary paths are used for all LSPs.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show mpls lsp | no-more
Ingress LSP: 5 sessions
To              From            State Rt P    ActivePath      LSPname
172.30.5.1      172.30.5.3      Up     0 *    primary-path    R3-to-R1
172.30.5.2      172.30.5.3      Up     0 *    primary-path    R3-to-R2
172.30.5.6      172.30.5.3      Up     0 *    primary-path    R3-to-R6
172.30.5.7      172.30.5.3      Up     0 *    primary-path    R3-to-R7
172.30.5.8      172.30.5.3      Up     0 *    primary-path    R3-to-R8
Total 5 displayed, Up 5, Down 0

Egress LSP: 8 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.3      172.30.5.7      Up       0  1 FF       3         - R7-to-R3
172.30.5.3      172.30.5.2      Up       0  1 SE       3         - R2-to-R3
172.30.5.3      172.30.5.2      Up       0  1 SE       3         - R2-to-R3
172.30.5.3      172.30.5.2      Up       0  1 SE       3         - R2-to-R3
172.30.5.3      172.30.5.2      Up       0  1 SE       3         - Bypass-
>172.30.0.14
172.30.5.3      172.30.5.1      Up       0  1 SE       3         - R1-to-R3
172.30.5.3      172.30.5.6      Up       0  1 FF       3         - R6-to-R3
172.30.5.3      172.30.5.8      Up       0  1 FF       3         - R8-to-R3
Total 8 displayed, Up 8, Down 0
Transit LSP: 8 sessions

To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.2      Up       0  1 SE  300272   299920 R2-to-R1-second
172.30.5.1      172.30.5.2      Up       0  1 SE  300288   299936 Bypass-
>172.30.0.1
172.30.5.2      172.30.5.8      Up       0  1 FF  300080        3 R8-to-R2
172.30.5.2      172.30.5.1      Up       0  1 SE  300224        3 Bypass-
>172.30.0.2
172.30.5.6      172.30.5.8      Up       0  1 FF  300096   301152 R8-to-R6
172.30.5.7      172.30.5.1      Up       0  1 SE  300240        3 Bypass-
>172.30.0.6
172.30.5.8      172.30.5.2      Up       0  1 SE  300256        3 R2-to-R8-second
172.30.5.8      172.30.5.6      Up       0  1 FF  300064        3 R6-to-R8
Total 8 displayed, Up 8, Down 0
```

Verify the RSVP session table.

```
lab@R3# run show rsvp session | no-more
Ingress RSVP: 11 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.3      Up       0  1 SE       -   299968 R3-to-R1
172.30.5.1      172.30.5.3      Up       0  1 SE       -   300000 R3-to-R1
172.30.5.2      172.30.5.3      Up       0  1 SE       -   299984 R3-to-R2
172.30.5.2      172.30.5.3      Up       0  1 SE       -   299984 R3-to-R2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.6      172.30.5.3      Up       0  1 SE       -   299968 R3-to-R6
172.30.5.6      172.30.5.3      Up       0  1 SE       -   300016 R3-to-R6
172.30.5.7      172.30.5.3      Up       0  1 SE       -        3 R3-to-R7
172.30.5.7      172.30.5.3      Up       0  1 SE       -   300016 R3-to-R7
172.30.5.8      172.30.5.3      Up       0  1 SE       -        3 R3-to-R8
172.30.5.8      172.30.5.3      Up       0  1 SE       -   300000 R3-to-R8
172.30.5.8      172.30.5.3      Up       0  1 SE       -   301424 Bypass-
>172.30.0.34
Total 11 displayed, Up 11, Down 0

Egress RSVP: 6 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.3      172.30.5.6      Up       0  1 FF       3         - R6-to-R3
172.30.5.3      172.30.5.2      Up       0  1 SE       3         - R2-to-R3
172.30.5.3      172.30.5.2      Up       0  1 SE       3         - Bypass-
>172.30.0.14
172.30.5.3      172.30.5.7      Up       0  1 FF       3         - R7-to-R3
172.30.5.3      172.30.5.1      Up       0  1 SE       3         - R1-to-R3
172.30.5.3      172.30.5.8      Up       0  1 FF       3         - R8-to-R3
Total 6 displayed, Up 6, Down 0

Transit RSVP: 9 sessions, 2 detours
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.2      Up       0  1 SE  300272   299920 R2-to-R1-second
172.30.5.1      172.30.5.2      Up       0  1 SE  300288   299936 Bypass-
>172.30.0.1
172.30.5.2      172.30.5.8      Up       0  1 FF  300080        3 R8-to-R2
172.30.5.2      172.30.5.1      Up       0  1 SE  300224        3 Bypass-
>172.30.0.2
172.30.5.6      172.30.5.8      Up       0  1 FF  300096   301152 R8-to-R6
172.30.5.7      172.30.5.3      Up       0  1 SE  300320        3 R3-to-R7
172.30.5.7      172.30.5.1      Up       0  1 SE  300240        3 Bypass-
>172.30.0.6
172.30.5.8      172.30.5.2      Up       0  1 SE  300256        3 R2-to-R8-second
172.30.5.8      172.30.5.6      Up       0  1 FF  300064        3 R6-to-R8
Total 9 displayed, Up 9, Down 0
```

Link-protection is enabled on the RSVP enabled interfaces.

```
[edit]
lab@R3# run show rsvp interface extensive
RSVP interface: 3 active
...

ge-0/0/4.123 Index 336, State Ena/Up
  Authentication, Aggregate, Reliable, LinkProtection
  HelloInterval 9(second)
  Address 172.30.0.14
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  ActiveResv 9, PreemptionCnt 0, Update threshold 10%, MaxResvTh 0bps, 0%
  Subscription 120%, Actual 120%
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 500Mbps
    MaxAvailableBW 1.2Gbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 500Mbps[4] 0bps[5] 200Mbps[6] 0bps[7]
0bps
  Protection: On, Bypass: 0, LSP: 0, Protected LSP: 0, Unprotected LSP: 0
...
```

Fast-reroute is desired on the LSPs and the detours are established over alternative paths to the egress LSR.

```
[edit]
lab@R3# run show mpls lsp name R3-to-R1 extensive
Ingress LSP: 5 sessions

172.30.5.1
  From: 172.30.5.3, State: Up, ActiveRoute: 0, LSPname: R3-to-R1
  ActivePath: primary-path (primary)
  FastReroute desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Follow destination IGP metric
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  LSP Self-ping Status : Enabled
 *Primary   primary-path     State: Up
    Priorities: 3 3
    Bandwidth: 100Mbps
    SmartOptimizeTimer: 180
          Include Any: blue green
    Flap Count: 0
    MBB Count: 0
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
 172.30.0.30 S 172.30.0.5 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.30.0.30(flag=5 Label=299968) 172.30.0.5(Label=3)
   13 Apr  9 10:29:38.128 Self-ping ended successfully
   12 Apr  9 10:29:37.337 Record Route:  172.30.0.30(flag=5 Label=299968)
172.30.0.5(Label=3)
   11 Apr  9 10:29:37.233 Fast-reroute Detour Up
   10 Apr  9 10:29:37.233 Self-ping started
    9 Apr  9 10:29:37.233 Self-ping enqueued
    8 Apr  9 10:28:50.790 Self-ping ended successfully
    7 Apr  9 10:28:50.148 Selected as active path
    6 Apr  9 10:28:50.147 Up
    5 Apr  9 10:28:50.146 Self-ping started
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    4 Apr  9 10:28:50.146 Self-ping enqueued
    3 Apr  9 10:28:50.146 Record Route:  172.30.0.30(Label=299968)
172.30.0.5(Label=3)
    2 Apr  9 10:28:50.110 Originate Call
    1 Apr  9 10:28:50.110 CSPF: computation result accepted  172.30.0.30
172.30.0.5
  Standby   secondary-path  State: Up
    Priorities: 3 3
    Bandwidth: 100Mbps
    SmartOptimizeTimer: 180
          Include Any: blue green
    Flap Count: 0
    MBB Count: 0
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 35)
 172.30.0.34 S 172.30.0.21 S 172.30.0.17 S 172.30.0.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
         172.30.0.34(flag=0xd Label=300000) 172.30.0.21(flag=0xd Label=300224)
172.30.0.17(flag=5 Label=301376) 172.30.0.1(Label=3)
   13 Apr  9 10:30:01.209 Self-ping ended successfully
   12 Apr  9 10:30:00.596 Fast-reroute Detour Up
   11 Apr  9 10:30:00.596 Self-ping started
   10 Apr  9 10:30:00.596 Self-ping enqueued
    9 Apr  9 10:29:37.830 Record Route:  172.30.0.34(flag=0xd Label=300000)
172.30.0.21(flag=0xd Label=300224) 172.30.0.17(flag=5 Label=301376)
172.30.0.1(Label=3)
    8 Apr  9 10:29:33.496 Record Route:  172.30.0.34(Label=300000)
172.30.0.21(Label=300224) 172.30.0.17(flag=5 Label=301376) 172.30.0.1(Label=3)
    7 Apr  9 10:28:52.793 Self-ping ended successfully
    6 Apr  9 10:28:52.595 Up
    5 Apr  9 10:28:52.595 Self-ping started
    4 Apr  9 10:28:52.595 Self-ping enqueued
    3 Apr  9 10:28:52.595 Record Route:  172.30.0.34(Label=300000)
172.30.0.21(Label=300224) 172.30.0.17(Label=301376) 172.30.0.1(Label=3)
    2 Apr  9 10:28:50.790 Originate Call
    1 Apr  9 10:28:50.790 CSPF: computation result accepted  172.30.0.34
172.30.0.21 172.30.0.17 172.30.0.1
  Created: Tue Apr  9 10:28:50 2019
Total 1 displayed, Up 1, Down 0

Egress LSP: 6 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 8 sessions, 3 detours
Total 0 displayed, Up 0, Down 0
```

d) Configure LSP traffic protection on R6.

© 2020 Juniper Networks, Inc. All rights reserved.

The requirements for traffic protection on R6 are the same as on R3. Secondary paths should be pre-established and fast reroute should be configured on all LSPs.

```
[edit]
lab@R6# set protocols mpls path primary-path

[edit]
lab@R6# set protocols mpls path secondary-path

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 primary primary-path

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 primary primary-path

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 primary primary-path

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 primary primary-path

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 primary primary-path

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 secondary secondary-path
standby

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 secondary secondary-path
standby

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 secondary secondary-path
standby

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 secondary secondary-path
standby

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 secondary secondary-path
standby

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 adaptive
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 adaptive

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 adaptive

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 adaptive

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 adaptive
```

Link-protection should be enabled on all RSVP interfaces.

```
[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.126 link-protection

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.167 link-protection

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.168 link-protection

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 fast-reroute bandwidth 100m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 fast-reroute bandwidth 100m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 fast-reroute bandwidth 100m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 fast-reroute bandwidth 100m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 fast-reroute bandwidth 100m

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 fast-reroute no-include-any

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 fast-reroute no-include-any

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R3 fast-reroute no-include-any

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 fast-reroute no-include-any
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 fast-reroute no-include-any

[edit]
lab@R6# commit
commit complete
```

Verify that the primary paths are used for all LSPs.

```
[edit]
lab@R6# run show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P    ActivePath      LSPname
172.30.5.1      172.30.5.6      Up     0 *    primary-path    R6-to-R1
172.30.5.2      172.30.5.6      Up     0 *    primary-path    R6-to-R2
172.30.5.3      172.30.5.6      Up     0 *    primary-path    R6-to-R3
172.30.5.7      172.30.5.6      Up     0 *    primary-path    R6-to-R7
172.30.5.8      172.30.5.6      Up     0 *    primary-path    R6-to-R8
Total 5 displayed, Up 5, Down 0

Egress LSP: 7 sessions, 3 detours
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.6      172.30.5.8      Up       0  1 FF       3         - R8-to-R6
172.30.5.6      172.30.5.3      Up       0  1 SE       3         - R3-to-R6
172.30.5.6      172.30.5.3      Up       0  1 SE       3         - R3-to-R6
172.30.5.6      172.30.5.1      Up       0  1 SE       3         - R1-to-R6
172.30.5.6      172.30.5.2      Up       0  1 SE       3         - R2-to-R6
172.30.5.6      172.30.5.2      Up       0  1 SE       3         - Bypass-
>172.30.0.18
172.30.5.6      172.30.5.7      Up       0  1 FF       3         - R7-to-R6
Total 7 displayed, Up 7, Down 0

Transit LSP: 9 sessions, 2 detours
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.3      Up       0  1 SE  300304   300048 R3-to-R1
172.30.5.1      172.30.5.3      Up       0  1 SE  300224   301376 R3-to-R1
172.30.5.2      172.30.5.7      Up       0  1 FF  300016        3 R7-to-R2
172.30.5.2      172.30.5.3      Up       0  1 SE  300208        3 R3-to-R2
172.30.5.3      172.30.5.7      Up       0  1 FF  300032   301136 R7-to-R3
172.30.5.7      172.30.5.3      Up       0  1 SE  300240   301408 R3-to-R7
172.30.5.8      172.30.5.2      Up       0  1 SE  300176        3 R2-to-R8-first
172.30.5.8      172.30.5.1      Up       0  1 SE  300160        3 R1-to-R8
172.30.5.8      172.30.5.3      Up       0  1 SE  300192        3 R3-to-R8
Total 9 displayed, Up 9, Down 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

The fast-reroute option is desired on the LSPs and detours are established over alternative path to the Egress LSR. The extensive view of the RSVP sessions shows the calculated path of the fast reroute detours.

```
[edit]
lab@R6# run show rsvp session name R3-to-R1 extensive
Ingress RSVP: 11 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 7 sessions, 3 detours
Total 0 displayed, Up 0, Down 0

Transit RSVP: 11 sessions, 2 detours

172.30.5.1
  From: 172.30.5.3, LSPstate: Up, ActiveRoute: 0
  LSPname: R3-to-R1, LSPpath: Secondary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300048
  Resv style: 1 SE, Label in: 300304, Label out: 300048
  Time left:  153, Since: Tue Apr  9 10:29:32 2019
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 9192
  Port number: sender 2 receiver 12920 protocol 0
  Enhanced FRR: Disabled, Reason: LSP Type, Refresh: 30 secs
  Detour branch from 172.30.0.17, to skip 172.30.5.1, Up
    Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 9192
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 172.30.0.17 (ge-0/0/4.126) 3 pkts
       incoming message handle: P-62/3, Message ID: 185, Epoch: 11876550
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 172.30.0.22 (ge-0/0/4.168) 2 pkts
       outgoing message state: refreshing, Message ID: 126, Epoch: 11876506
    RESV rcvfrom: 172.30.0.22 (ge-0/0/4.168) 3 pkts, Entropy label: Yes
       incoming message handle: R-63/3, Message ID: 122, Epoch: 11876451
    RESV
       outgoing message state: refreshing, Message ID: 129, Epoch: 11876506
    Explct route: 172.30.0.22 172.30.0.9
    Record route: 172.30.0.33 172.30.0.22 172.30.0.18 172.30.0.17 <self>
172.30.0.22 172.30.0.9
    Label in: 300304, Label out: 300048

172.30.5.1
  From: 172.30.5.3, LSPstate: Up, ActiveRoute: 0
  LSPname: R3-to-R1, LSPpath: Secondary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 301376
  Resv style: 1 SE, Label in: 300224, Label out: 301376
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Time left:  134, Since: Tue Apr  9 10:28:51 2019
Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 9192
Port number: sender 2 receiver 12920 protocol 0
FastReroute desired
Enhanced FRR: Disabled, Reason: LSP Type, Refresh: 30 secs
PATH rcvfrom: 172.30.0.22 (ge-0/0/4.168) 3 pkts
     incoming message handle: P-54/3, Message ID: 103, Epoch: 11876451
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 172.30.0.17 (ge-0/0/4.126) 5 pkts
     outgoing message state: refreshing, Message ID: 104, Epoch: 11876506
RESV rcvfrom: 172.30.0.17 (ge-0/0/4.126) 5 pkts, Entropy label: Yes
     incoming message handle: R-60/5, Message ID: 186, Epoch: 11876550
RESV
     outgoing message state: refreshing, Message ID: 120, Epoch: 11876506
  Explct route: 172.30.0.17 172.30.0.1
  Record route: 172.30.0.33 172.30.0.22 <self> 172.30.0.17 172.30.0.1
    Detour is Up
    Detour Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 9192
    Detour adspec: received MTU 1500 sent MTU 1500
    Path MTU: received 1500
    Detour PATH sentto: 172.30.0.46 (ge-0/0/4.167) 2 pkts
       outgoing message state: refreshing, Message ID: 105, Epoch: 11876506
    Detour RESV rcvfrom: 172.30.0.46 (ge-0/0/4.167) 2 pkts, Entropy label: Yes
       incoming message handle: R-69/2, Message ID: 120, Epoch: 11876463
    Detour RESV
       outgoing message state: refreshing, Message ID: 120, Epoch: 11876506
    Detour Explct route: 172.30.0.46 172.30.0.5
    Detour Record route: 172.30.0.33 172.30.0.22 <self> 172.30.0.46 172.30.0.5
    Detour Label out: 300048
Total 2 displayed, Up 2, Down 0
```

e) Configure LSP traffic protection on R7.

```
[edit]
lab@R7# set protocols mpls path primary-path

[edit]
lab@R7# set protocols mpls path secondary-path

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R1 primary primary-path

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 primary primary-path

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R3 primary primary-path
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R6 primary primary-path

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R8 primary primary-path

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R1 secondary secondary-path

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 secondary secondary-path

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R3 secondary secondary-path

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R6 secondary secondary-path


[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R8 secondary secondary-path

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.117 link-protection

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.137 link-protection

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.167 link-protection

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R1 link-protection

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R2 link-protection

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R3 link-protection

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R6 link-protection

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R8 link-protection

[edit]
lab@R7# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Verify that the primary paths are used for all LSPs.

```
[edit]
lab@R7# run show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P    ActivePath      LSPname
172.30.5.1      172.30.5.7      Up     0 *    primary-path    R7-to-R1
172.30.5.2      172.30.5.7      Up     0 *    primary-path    R7-to-R2
172.30.5.3      172.30.5.7      Up     0 *    primary-path    R7-to-R3
172.30.5.6      172.30.5.7      Up     0 *    primary-path    R7-to-R6
172.30.5.8      172.30.5.7      Up     0 *    primary-path    R7-to-R8
Total 5 displayed, Up 5, Down 0

Egress LSP: 9 sessions, 6 detours
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.7      172.30.5.8      Up       0  1 FF       3         - R8-to-R7
172.30.5.7      172.30.5.3      Up       0  1 SE       3         - R3-to-R7
172.30.5.7      172.30.5.3      Up       0  1 SE       3         - R3-to-R7
172.30.5.7      172.30.5.3      Up       0  1 SE       3         - Bypass-
>172.30.0.30
172.30.5.7      172.30.5.6      Up       0  1 SE       3         - R6-to-R7
172.30.5.7      172.30.5.6      Up       0  1 SE       3         - R6-to-R7
172.30.5.7      172.30.5.2      Up       0  1 SE       3         - R2-to-R7
172.30.5.7      172.30.5.1      Up       0  1 SE       3         - R1-to-R7
172.30.5.7      172.30.5.1      Up       0  1 SE       3         - Bypass-
>172.30.0.6
Total 9 displayed, Up 9, Down 0

Transit LSP: 20 sessions, 1 detours
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.3      Up       0  1 SE  299968         3 R3-to-R1
172.30.5.1      172.30.5.3      Up       0  1 SE  300400         3 R3-to-R1
172.30.5.1      172.30.5.2      Up       0  1 SE  300272         3 R2-to-R1-first
172.30.5.1      172.30.5.2      Up       0  1 SE  300192         3 R2-to-R1-first
172.30.5.1      172.30.5.6      Up       0  1 SE  300080         3 R6-to-R1
172.30.5.1      172.30.5.6      Up       0  1 SE  300080         3 R6-to-R1
172.30.5.2      172.30.5.6      Up       0  1 SE  300112    301568 R6-to-R2
172.30.5.2      172.30.5.6      Up       0  1 SE  300288    300720 Bypass-
>172.30.0.17
172.30.5.2      172.30.5.3      Up       0  1 SE  300208    300544 R3-to-R2
172.30.5.3      172.30.5.6      Up       0  1 SE  300128    301584 R6-to-R3
172.30.5.3      172.30.5.2      Up       0  1 SE  300224         3 Bypass-
>172.30.0.14
172.30.5.3      172.30.5.1      Up       0  1 SE  299920         3 R1-to-R3
172.30.5.6      172.30.5.3      Up       0  1 SE  300384    301920 R3-to-R6
172.30.5.6      172.30.5.1      Up       0  1 SE  300240    300624 R1-to-R6
172.30.5.8      172.30.5.1      Up       0  1 SE  300176    300528 R1-to-R8
172.30.5.8      172.30.5.3      Up       0  1 SE  300352    300592 R3-to-R8
172.30.5.8      172.30.5.3      Up       0  1 SE  300368    301904 R3-to-R8
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.8      172.30.5.3      Up        0  1 SE  300320   301824 Bypass-
>172.30.0.34
172.30.5.8      172.30.5.6      Up        0  1 SE  300336   301856 R6-to-R8
172.30.5.8      172.30.5.6      Up        0  1 SE  300096   301536 R6-to-R8
172.30.5.8      172.30.5.6      Up        0  1 SE  300304   301808 Bypass-
>172.30.0.22
Total 20 displayed, Up 20, Down 0
```

f) Configure LSP traffic protection on R8.

```
[edit]
lab@R8# set protocols mpls path primary-path

[edit]
lab@R8# set protocols mpls path secondary-path


[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 primary primary-path

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R2 primary primary-path

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3 primary primary-path

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R6 primary primary-path

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R7 primary primary-path

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 secondary secondary-path

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R2 secondary secondary-path

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3 secondary secondary-path

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R6 secondary secondary-path

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R7 secondary secondary-path

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# set protocols rsvp interface ge-0/0/4.118 link-protection

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.138 link-protection

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.168 link-protection

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 link-protection

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R2 link-protection


[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3 link-protection


[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R6 link-protection

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R7 link-protection

[edit]
lab@R8# commit
commit complete
```

Verify that the primary paths are used for all LSPs.

```
[edit]
lab@R8# run show mpls lsp
Ingress LSP: 5 sessions
To              From            State Rt P    ActivePath      LSPname
172.30.5.1      172.30.5.8      Up     0 *    primary-path    R8-to-R1
172.30.5.2      172.30.5.8      Up     0 *    primary-path    R8-to-R2
172.30.5.3      172.30.5.8      Up     0 *    primary-path    R8-to-R3
172.30.5.6      172.30.5.8      Up     0 *    primary-path    R8-to-R6
172.30.5.7      172.30.5.8      Up     0 *    primary-path    R8-to-R7
Total 5 displayed, Up 5, Down 0

Egress LSP: 12 sessions, 7 detours
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.8      172.30.5.2      Up       0  1 SE       3         - R2-to-R8-first
172.30.5.8      172.30.5.2      Up       0  1 SE       3         - R2-to-R8-second
172.30.5.8      172.30.5.1      Up       0  1 SE       3         - R1-to-R8
172.30.5.8      172.30.5.1      Up       0  1 SE       3         - R1-to-R8
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.5.8      172.30.5.1      Up       0  1 SE       3        - Bypass-
>172.30.0.10
172.30.5.8      172.30.5.3      Up       0  1 SE       3        - R3-to-R8
172.30.5.8      172.30.5.3      Up       0  1 SE       3        - R3-to-R8
172.30.5.8      172.30.5.3      Up       0  1 SE       3        - Bypass-
>172.30.0.34
172.30.5.8      172.30.5.7      Up       0  1 SE       3        - R7-to-R8
172.30.5.8      172.30.5.6      Up       0  1 SE       3        - R6-to-R8
172.30.5.8      172.30.5.6      Up       0  1 SE       3        - R6-to-R8
172.30.5.8      172.30.5.6      Up       0  1 SE       3        - Bypass-
>172.30.0.22
Total 12 displayed, Up 12, Down 0

Transit LSP: 23 sessions, 5 detours
To              From            State   Rt Style Labelin Labelout LSPname
172.30.5.1      172.30.5.3      Up       0  1 SE  300464   300640 R3-to-R1
172.30.5.1      172.30.5.2      Up       0  1 SE  300432   300704 R2-to-R1-first
172.30.5.1      172.30.5.2      Up       0  1 SE  299920        3 R2-to-R1-second
172.30.5.1      172.30.5.2      Up       0  1 SE  299936        3 Bypass-
>172.30.0.1
172.30.5.1      172.30.5.6      Up       0  1 SE  300256        3 R6-to-R1
172.30.5.1      172.30.5.6      Up       0  1 SE  300096        3 R6-to-R1
172.30.5.2      172.30.5.6      Up       0  1 SE  300080   300368 R6-to-R2
172.30.5.2      172.30.5.6      Up       0  1 SE  300368   300608 R6-to-R2
172.30.5.2      172.30.5.1      Up       0  1 SE  299904   300224 Bypass-
>172.30.0.2
172.30.5.2      172.30.5.3      Up       0  1 SE  299984   300208 R3-to-R2
172.30.5.2      172.30.5.3      Up       0  1 SE  300272   300416 R3-to-R2
172.30.5.3      172.30.5.6      Up       0  1 SE  300128        3 R6-to-R3
172.30.5.3      172.30.5.6      Up       0  1 SE  300144        3 R6-to-R3
172.30.5.3      172.30.5.2      Up       0  1 SE  300240        3 R2-to-R3
172.30.5.3      172.30.5.7      Up       0  1 SE  300544        3 R7-to-R3
172.30.5.6      172.30.5.3      Up       0  1 SE  300512        3 R3-to-R6
172.30.5.6      172.30.5.3      Up       0  1 SE  299968        3 R3-to-R6
172.30.5.6      172.30.5.2      Up       0  1 SE  300400        3 Bypass-
>172.30.0.18
172.30.5.7      172.30.5.3      Up       0  1 SE  300448   301872 R3-to-R7
172.30.5.7      172.30.5.3      Up       0  1 SE  300480   300688 R3-to-R7
172.30.5.7      172.30.5.6      Up       0  1 SE  300352   300592 R6-to-R7
172.30.5.7      172.30.5.6      Up       0  1 SE  300112   301552 R6-to-R7
172.30.5.7      172.30.5.1      Up       0  1 SE  300416   300496 Bypass-
>172.30.0.6
Total 23 displayed, Up 23, Down 0
```

5) Configure LDP tunneling from R4 and R5 to the route reflector.

© 2020 Juniper Networks, Inc. All rights reserved.

LDP is used on R4, R5 and the route reflector, whereas RSVP is used on the other routers. To provide full MPLS services you need to have a full mesh of LSP paths between the loopbacks of all backbone routers.

LDP tunneling on the RSVP LSPs in Area 0 allows LDP to establish extended sessions over RSVP.

Because you have to make sure that a failure of single link will not cause LSPs to break, the most efficient way is to enable `ldp-tunneling` on all direct RSVP LSPs between the routers. Doing so will provide for the required redundancy.

a) Enable `ldp-tunneling` on R1.

To establish extended LDP sessions you have to enable LDP on the loopback interface.

```
[edit]
lab@R1# set protocols ldp interface lo0.0

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R2 ldp-tunneling

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R8 ldp-tunneling

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R7 ldp-tunneling

[edit]
lab@R1# commit
commit complete
```

b) Enable `ldp-tunneling` on R2.

```
[edit]
lab@R2# set protocols ldp interface lo0.0

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-first ldp-tunneling

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R1-second ldp-tunneling

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R3 ldp-tunneling

[edit]
lab@R2# set protocols mpls label-switched-path R2-to-R6 ldp-tunneling
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# commit
commit complete
```

c) Enable ldp-tunneling on R3.

```
[edit]
lab@R3# set protocols ldp interface lo0.0

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R2 ldp-tunneling

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R7 ldp-tunneling

[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R8 ldp-tunneling

[edit]
lab@R3# commit
commit complete
```

d) Enable ldp-tunneling on R6.

```
[edit]
lab@R6# set protocols ldp interface lo0.0

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R2 ldp-tunneling

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R7 ldp-tunneling

[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R8 ldp-tunneling

[edit]
lab@R6# commit
commit complete
```

e) Enable ldp-tunneling on R7.

```
[edit]
lab@R7# set protocols ldp interface lo0.0

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R6 ldp-tunneling
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R3 ldp-tunneling

[edit]
lab@R7# set protocols mpls label-switched-path R7-to-R1 ldp-tunneling

[edit]
lab@R7# commit
commit complete
```

f) Enabling `ldp-tunneling` on R8.

```
[edit]
lab@R8# set protocols ldp interface lo0.0

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R6 ldp-tunneling

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R3 ldp-tunneling

[edit]
lab@R8# set protocols mpls label-switched-path R8-to-R1 ldp-tunneling

[edit]
lab@R8# commit
commit complete
```

Verify the operation of `ldp-tunneling` and the full mesh of LSP between the routers.

Router R1 has three extended LDP session over the RSVP LSPs to the directly connected routers.

```
[edit]
lab@R1# run show ldp neighbor
Address            Interface          Label space ID         Hold time
172.30.0.66        ge-0/0/4.206       172.30.5.41:0            12
172.30.5.2         lo0.0              172.30.5.2:0             44
172.30.5.7         lo0.0              172.30.5.7:0             32
172.30.5.8         lo0.0              172.30.5.8:0             34
```

Verify if R4 and R5 have all loopback IP addresses in the `inet.3` routing table.

```
[edit]
lab@R4# run show route table inet.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.3: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.2/32      *[LDP/9] 00:32:08, metric 15
                    > to 172.30.0.81 via ae0.0, Push 300576
172.30.5.3/32      *[LDP/9] 03:06:50, metric 5
                    > to 172.30.0.81 via ae0.0
172.30.5.5/32      *[LDP/9] 00:32:09, metric 15
                      to 172.30.0.81 via ae0.0, Push 299808
                    > to 172.30.0.90 via ge-0/0/4.146, Push 299872
172.30.5.6/32      *[LDP/9] 00:33:35, metric 10
                    > to 172.30.0.90 via ge-0/0/4.146
172.30.5.7/32      *[LDP/9] 00:32:08, metric 15
                    > to 172.30.0.81 via ae0.0, Push 300560
172.30.5.8/32      *[LDP/9] 00:32:09, metric 15
                    > to 172.30.0.81 via ae0.0, Push 300544
172.30.5.41/32     *[LDP/9] 00:32:08, metric 25
                    > to 172.30.0.81 via ae0.0, Push 300592
```

```
[edit]
lab@R5# run show route table inet.3

inet.3: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.2/32      *[LDP/9] 00:00:26, metric 15
                    > to 172.30.0.94 via ae0.0, Push 300480
172.30.5.3/32      *[LDP/9] 02:34:31, metric 10
                    > to 172.30.0.85 via ge-0/0/4.135
172.30.5.4/32      *[LDP/9] 00:00:22, metric 15
                      to 172.30.0.94 via ae0.0, Push 299776
                    > to 172.30.0.85 via ge-0/0/4.135, Push 299776
172.30.5.6/32      *[LDP/9] 02:34:00, metric 5
                    > to 172.30.0.94 via ae0.0
172.30.5.7/32      *[LDP/9] 00:01:07, metric 15
                    > to 172.30.0.94 via ae0.0, Push 300432
172.30.5.8/32      *[LDP/9] 00:00:32, metric 15
                    > to 172.30.0.94 via ae0.0, Push 300464
172.30.5.41/32     *[LDP/9] 00:00:26, metric 25
                    > to 172.30.0.94 via ae0.0, Push 300496
```

R1's loopback address is missing in the `inet.3` table on both routers.

R3 and R6 are the gateways for R4 and R5 to the other part of the network. Verify on R6 the reason why the loopback address of R1 is missing.

```
[edit]
lab@R6# run show route table inet.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.3: 16 destinations, 25 routes (8 active, 0 holddown, 14 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       *[RSVP/7/1] 00:46:17, metric 20
                     > to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R1
                       to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R1
                       to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R1
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R1
172.30.5.2/32       *[RSVP/7/1] 01:14:58, metric 10
                     > to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R2
                       to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R2
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R2
                       to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R2
                      [LDP/9] 00:05:58, metric 10
                     > to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R2
                       to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R2
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R2
                       to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R2
172.30.5.3/32       *[RSVP/7/1] 01:14:58, metric 20
                     > to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R3
                       to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R3
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R3
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R3
172.30.5.4/32       *[LDP/9] 00:07:21, metric 10
                     > to 172.30.0.89 via ge-0/0/4.146
172.30.5.5/32       *[LDP/9] 00:07:21, metric 5
                     > to 172.30.0.93 via ae0.0
172.30.5.7/32       *[RSVP/7/1] 01:14:58, metric 10
                     > to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R7
                       to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R7
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                      to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R7
                      to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R7
                    [LDP/9] 00:06:04, metric 10
                    > to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R7
                      to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R7
                      to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R7
                      to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R7
172.30.5.8/32       *[RSVP/7/1] 00:46:17, metric 10
                     > to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R8
                       to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R8
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R8
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R8
                     [LDP/9] 00:06:04, metric 10
                     > to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R8
                       to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R8
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R8
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R8
172.30.5.41/32      *[LDP/9] 00:05:58, metric 20
                     > to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R2
                       to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R2
                       to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R2
                       to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R2
```

The route of R1 loopback IP address is known only via the RSVP LSP and not via LDP like the route to R2's loopback IP address.

Verify the hidden routes installed in `inet.3` because of the `ldp-tunneling` feature.

```
[edit]
lab@R6# run show route table inet.3 172.30.5.1 hidden
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.3: 16 destinations, 25 routes (8 active, 0 holddown, 14 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       [OSPF] 00:28:04, metric 20
                    > to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R1
                      to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R1
                      to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R1
                      to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R1
```

```
[edit]
lab@R6# run show route table inet.3 172.30.5.2 hidden

inet.3: 16 destinations, 25 routes (8 active, 0 holddown, 14 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.2/32       [OSPF] 00:26:56, metric 10
                    > to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R2
                      to 172.30.0.46 via ge-0/0/4.167, label-switched-path R6-to-
R2
                      to 172.30.0.22 via ge-0/0/4.168, label-switched-path R6-to-
R2
                      to 172.30.0.17 via ge-0/0/4.126, label-switched-path R6-to-
R2
```

The route for R1 is known via LSP `R6-to-R1` without `ldp-tunneling` enabled, whereas R2 loopback IP address is known via LSP with enabled `ldp-tunneling`.

Fix this with enabling `ldp-tunneling` on R6 and R3 on the LSP to R1 and on R1 for the LSP to R3 and R6.

R3:

```
[edit]
lab@R3# set protocols mpls label-switched-path R3-to-R1 ldp-tunneling

[edit]
lab@R3# commit
commit complete
```

R6:

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 ldp-tunneling

[edit]
lab@R6# commit
commit complete
```

R1:

```
[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R3 ldp-tunneling

[edit]
lab@R1# set protocols mpls label-switched-path R1-to-R6 ldp-tunneling

[edit]
lab@R1# commit
commit complete
```

Verify the `inet.3` table for R1 loopback IP address after enabling `ldp-tunneling` on the LSP from R1 to R3 and R6. You may need to wait a minute or two for the protocols to synchronize before R1's loopback IP address shows up.

```
[edit]
lab@R4# run show route table inet.3

inet.3: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       *[LDP/9] 00:00:30, metric 25
                     > to 172.30.0.81 via ae0.0, Push 300720
172.30.5.2/32       *[LDP/9] 00:33:53, metric 15
                     > to 172.30.0.81 via ae0.0, Push 300576
172.30.5.3/32       *[LDP/9] 03:08:35, metric 5
                     > to 172.30.0.81 via ae0.0
172.30.5.5/32       *[LDP/9] 00:33:54, metric 15
                       to 172.30.0.81 via ae0.0, Push 299808
                     > to 172.30.0.90 via ge-0/0/4.146, Push 299872
172.30.5.6/32       *[LDP/9] 00:35:20, metric 10
                     > to 172.30.0.90 via ge-0/0/4.146
172.30.5.7/32       *[LDP/9] 00:33:53, metric 15
                     > to 172.30.0.81 via ae0.0, Push 300560
172.30.5.8/32       *[LDP/9] 00:33:54, metric 15
                     > to 172.30.0.81 via ae0.0, Push 300544
172.30.5.41/32      *[LDP/9] 00:33:53, metric 25
                     > to 172.30.0.81 via ae0.0, Push 300592
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# run show route table inet.3

inet.3: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32       *[LDP/9] 00:00:44, metric 25
                     > to 172.30.0.94 via ae0.0, Push 300592
172.30.5.2/32       *[LDP/9] 00:32:05, metric 15
                     > to 172.30.0.94 via ae0.0, Push 300480
172.30.5.3/32       *[LDP/9] 03:06:10, metric 10
                     > to 172.30.0.85 via ge-0/0/4.135
172.30.5.4/32       *[LDP/9] 00:32:01, metric 15
                       to 172.30.0.94 via ae0.0, Push 299776
                     > to 172.30.0.85 via ge-0/0/4.135, Push 299776
172.30.5.6/32       *[LDP/9] 03:05:39, metric 5
                     > to 172.30.0.94 via ae0.0
172.30.5.7/32       *[LDP/9] 00:32:46, metric 15
                     > to 172.30.0.94 via ae0.0, Push 300432
172.30.5.8/32       *[LDP/9] 00:32:11, metric 15
                     > to 172.30.0.94 via ae0.0, Push 300464
172.30.5.41/32      *[LDP/9] 00:32:05, metric 25
                     > to 172.30.0.94 via ae0.0, Push 300496
```

R1 loopback is resolved in the `inet.3` table using the LDP protocol.

6) You have to make sure that IPv4 and IPv6 traffic from CE to P1 are mapped to different LSPs.

In this Step you are asked to configure a forwarding table policy mapping IPv4 traffic from C3 to P1 to one LSP and IPv6 traffic to the other.

P1 has two connections to your AS, one on R1 and other on R8. You configured two LSPs from R2 to R1 and two others to R8. You can use a regular expression to use the LSPs to R1 for IPv4 traffic and both LSPs to R8 for IPv6 traffic.

Configure the mapping policy applied to the forwarding table.

```
[edit]
lab@R2# set policy-options policy-statement lsp-mapping term 1 from family inet

[edit]
lab@R2# set policy-options policy-statement lsp-mapping term 1 from protocol bgp

[edit]
lab@R2# set policy-options policy-statement lsp-mapping term 1 from as-path p1-neighbor
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set policy-options policy-statement lsp-mapping term 1 then install-nexthop
lsp-regex R2-to-R1-first

[edit]
lab@R2# set policy-options policy-statement lsp-mapping term 2 from family inet6

[edit]
lab@R2# set policy-options policy-statement lsp-mapping term 2 from protocol bgp

[edit]
lab@R2# set policy-options policy-statement lsp-mapping term 2 from as-path p1-neighbor

[edit]
lab@R2# set policy-options policy-statement lsp-mapping term 2 then install-nexthop
lsp-regex R2-to-R8-second

[edit]
lab@R2# set policy-options as-path p1-neighbor "110047427 .*"
```

Applying the policy to R2's forwarding table.

```
[edit]
lab@R2# set routing-options forwarding-table export lsp-mapping

[edit]
lab@R2# commit
commit complete
```

Verify the forwarding table for the IPv4 and IPv6 LSP mapping.

```
[edit]
lab@R2# run show route forwarding-table destination 1.64.0.0/10
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
1.64.0.0/10        user     0                    indr 262143   403
                              172.30.0.1         ucst   620     2 ae0.0
…
```

It looks that IPv4 traffic is not mapped to the LSP, however looking at the routing table, it is clear that the particular next hop is LSP R2-to-R1 first. In the forwarding table, no label is pushed to the traffic because the Egress LSR, R1 signaled an implicit null label for this LSP.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show route aspath-regex "110047427 .*"

inet.0: 524 destinations, 524 routes (521 active, 0 holddown, 3 hidden)
+ = Active Route, - = Last Active, * = Both

1.64.0.0/10         *[BGP/170] 00:35:36, localpref 150, from 172.30.5.41
                       AS path: 110047427 1620 61671 I
                       to 172.30.0.1 via ae0.0, label-switched-path R2-to-R1-first
                       to 172.30.0.18 via ge-0/0/4.126, label-switched-path
Bypass->172.30.0.1
...
```

The forwarding table shows that for IPv6 traffic two labels are pushed to the packet. The reason for this is that an Explicit Null label is used for 6PE. (These outputs were taken after enabling 6PE.)

```
[edit]
lab@R2# run show route forwarding-table destination fd01:aaaa:bbbb::/80
Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
fd01:aaaa:bbbb::/80 user     0                    indr 262145    17
                              172.30.0.14         Push 2, Push 300848(top)   719
2 ge-0/0/4.123
…
```

```
[edit]
lab@R2# run show route table inet6.0 aspath-regex "110047427 .*"

inet6.0: 49 destinations, 53 routes (49 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fd01:aaaa:bbbb::/80*[BGP/170] 00:38:33, localpref 100, from 172.30.5.41
                      AS path: 110047427 I
                      to 172.30.0.14 via ge-0/0/4.123, label-switched-path R2-to-
R1-second
                      to 172.30.0.1 via ae0.0, label-switched-path Bypass-
>172.30.0.14
...
```

7) Configure load balancing in the core.

You should configure the load-balancing policy applied to the forwarding table to utilize all the LSPs in the MPLS backbone network.

© 2020 Juniper Networks, Inc. All rights reserved.

The load balancing policy should be configured on all routes in the same was a done on R1.

Repeat this on all eight devices.

```
[edit]
lab@R1# set policy-options policy-statement load-balancing term 1 then load-balance
per-packet

[edit]
lab@R1# set routing-options forwarding-table export load-balancing

[edit]
lab@R1# commit
commit complete
```

8) Enable IPv6 tunneling.

IPv6 tunneling makes it possible to forward IPv6 traffic over MPLS. This is known as the 6PE model, more information on this topic is provided in Chapter 4.

`ipv6-tunneling` under `protocols mpls` should be configured on all routers. This will copy the routes from `inet.3` to `inet6.3` in v4-mapped IPv6 addresses.

Configure IPv6 tunneling on R1, R2, R3, R5, R6, R7, R8 and router reflector.

```
[edit]
lab@R1# set protocols mpls ipv6-tunneling

[edit]
lab@R1# commit
commit complete
```

After committing the configuration routing table `inet6.3` is populated.

```
[edit]
lab@R1# run show route table inet6.3

inet6.3: 8 destinations, 13 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::ffff:172.30.5.2/128
                   *[RSVP/7/1] 00:00:42, metric 11
                    > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R2
                      to 172.30.0.10 via ge-0/0/4.118, label-switched-path
Bypass->172.30.0.2
                    [LDP/9] 00:00:34, metric 11
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R2
                      to 172.30.0.10 via ge-0/0/4.118, label-switched-path R1-to-
R2
::ffff:172.30.5.3/128
                   *[RSVP/7/1] 00:00:41, metric 65545
                    > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R3
                      to 172.30.0.2 via ae0.0, label-switched-path Bypass-
>172.30.0.6
                    [LDP/9] 00:00:34, metric 65545
                    > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R3
                      to 172.30.0.2 via ae0.0, label-switched-path R1-to-R3
::ffff:172.30.5.4/128
                   *[LDP/9] 00:00:34, metric 65550
                    > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R3
                      to 172.30.0.2 via ae0.0, label-switched-path R1-to-R3
::ffff:172.30.5.5/128
                   *[LDP/9] 00:00:34, metric 65550
                    > to 172.30.0.2 via ae0.0, label-switched-path R1-to-R6
                      to 172.30.0.10 via ge-0/0/4.118, label-switched-path R1-to-
R6
...
```

Enable the IPv6 `labeled-unicast family` in BGP with the `explicit null` option. This BGP family will advertise routes with an explicit null label of 2 for all IPv6 routes in `inet6.0`.

Repeat this on all devices including the route reflector.

```
[edit]
lab@R1# set protocols bgp group ibgp family inet6 labeled-unicast explicit-null

[edit]
lab@R1# set protocols bgp group ibgp family inet unicast

[edit]
lab@R1# commit
commit complete
```

Remember to add the IPv4 unicast family explicitly, otherwise only the BGP `family inet6 labeled-unicast explicit-null` will be enabled.

```
[edit]
lab@R1# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       650        488          0          0          0          0
inet6.0
                        32         32          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591         55        370       0       0        8:06
Establ
  inet.0: 86/86/86/0
  inet6.0: 16/16/16/0
192.168.0.38      110047427       2498       1565       0       1     3:50:32
402/564/402/0        0/0/0/0
fc09:c0:ffee::6   110047427       1513       1530       0       0     3:49:52
Establ
  inet6.0: 16/16/16/0
```

16 IPv6 routes are received from the route reflector.

An extensive view of the IPv6 route advertised from R2 shows that next-hop is ::ffff:172.30.5.2, resolved via a RSVP LSP. The transport label value is 2.

```
[edit]
lab@R1# run show route fd18:cccc:dddd:2:2::/80 extensive

inet6.0: 49 destinations, 49 routes (49 active, 0 holddown, 0 hidden)
fd18:cccc:dddd:2:2::/80 (1 entry, 1 announced)
TSI:
KRT in-kernel fd18:cccc:dddd:2:2::/80 -> {indirect(1048602)}
Page 0 idx 0, (group P1-1-ipv6 type External) Type 1 val 0xd9db29c (adv_entry)
   Advertised metrics:
     Nexthop: Self
     AS path: [54591] I
     Communities: no-export
Path fd18:cccc:dddd:2:2:: from 172.30.5.41 Vector len 4.  Val: 0
        *BGP    Preference: 170/-201
                Next hop type: Indirect, Next hop index: 0
                Address: 0xd989b50
                Next-hop reference count: 32
                Source: 172.30.5.41
                Next hop type: Router, Next hop index: 0
                Next hop: 172.30.0.2 via ae0.0 weight 0x1, selected
                Label-switched-path R1-to-R2
                Label operation: Push 2
                Label TTL action: prop-ttl
                Load balance label: Label 2: None;
                Label element ptr: 0xd989660
                Label parent element ptr: 0xd982c40
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 172.30.0.2 via ae0.0 weight 0x2001
Label-switched-path R1-to-R2
Label operation: Push 2
Label TTL action: prop-ttl
Load balance label: Label 2: None;
Label element ptr: 0xd9890c0
Label parent element ptr: 0xd988b80
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 172.30.0.10 via ge-0/0/4.118 weight 0x8001 uflags Mark
connection protection tlv
Label-switched-path Bypass->172.30.0.2
Label operation: Push 2, Push 299904(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 2: None; Label 299904: None;
Label element ptr: 0xd98dec0
```

The route is successfully advertised from R1 to P1-1.

```
[edit]
lab@R1# run show route advertising-protocol bgp fc09:c0:ffee::6
fd18:cccc:dddd:2:2::/80

inet6.0: 49 destinations, 53 routes (49 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  fd18:cccc:dddd:2:2::/80
*                         Self                                    I
```

Task 6: VPN Configuration

In this task, the following VPN services are configured.

- Implement L3VPNs with customer sites running either OSPF or BGP as PE-CE protocols.
- Hub-and-spoke topology
- Internet access for the VPN customers,
- Multicast
- Inter-provider VPNs.
- L2 VPLS infrastructure including dual-homed customer sites and VLAN normalization.

© 2020 Juniper Networks, Inc. All rights reserved.

- EVPN

Customer	Site	Router	PE-CE Protocol	Protocol details
CE1	S1	CE1-1	OSPF	Area 0.0.0.0
	S2	CE1-2	OSPF	Area 0.0.0.0
CE2	S1	CE2-1	BGP	AS 64600
	S2	CE2-2	BGP	AS 64600
	S3	CE2-3	BGP	AS 64600

Table 31

1) Make sure that all PE routers receive only the routes for route-targets they specifically request.

To accomplish this requirement, you have to enable family route-target on all routers that have a BGP session with route reflector. Remember that to resolve the routes in table `bgp.rtarget.0` you should have the route reflector loopback IP address in the `inet.3` routing table. When the route reflector is enabled for LDP, its loopback IP address should be present in `inet.3` and you should not have to perform any fancy tricks (as done in the previous chapters).

Enable `family route-target` on all routers including the route reflector.

Note: The route reflector uses cluster-1 and cluster-2 as the BGP group.

```
[edit]
lab@R1# set protocols bgp group ibgp family route-target

[edit]
lab@R1# commit
commit complete
```

2) Configure **`family inet-vpn unicast`** to exchange VPNv4 routes between PE routers.

```
[edit]
lab@R1# set protocols bgp group ibgp family inet-vpn unicast

[edit]
lab@R1# commit
commit complete
```

3) Configure customer CE1 L3VPN.

© 2020 Juniper Networks, Inc. All rights reserved.

Customer CE1 has two sites attached on R1 and R7. The PE-CE protocol used by the customer is OSPF.

a) Configure R1 for customer CE1 L3VPN.

First you have to configure the PE-CE interface towards the customer router.

```
[edit]
lab@R1# set interfaces ge-0/0/5 unit 318 vlan-id 318

[edit]
lab@R1# set interfaces ge-0/0/5 unit 318 description "CE1-2 connection"

[edit]
lab@R1# set interfaces ge-0/0/5 unit 318 family inet address 192.168.0.69/30
```

The automatic route-distinguisher is based on Router ID.

```
[edit]
lab@R1# set routing-options route-distinguisher-id 172.30.5.1
```

Configure the routing-instance dedicated for customer CE1.

```
[edit]
lab@R1# set routing-instances CE1 instance-type vrf

[edit]
lab@R1# set routing-instances CE1 interface ge-0/0/5.318

[edit]
lab@R1# set routing-instances CE1 vrf-export CE1-export

[edit]
lab@R1# set routing-instances CE1 vrf-import CE1-import
```

Note: VRF policies will be defined later.

OSPF (Area 0) is enabled for the PE-CE interface.

```
[edit]
lab@R1# set routing-instances CE1 protocols ospf area 0 interface all

[edit]
lab@R1# set policy-options policy-statement CE1-export term 1 from protocol ospf

[edit]
lab@R1# set policy-options policy-statement CE1-export term 1 then community add CE1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set policy-options policy-statement CE1-export term 1 then accept

[edit]
lab@R1# set policy-options policy-statement CE1-import term 1 from protocol bgp


[edit]
lab@R1# set policy-options policy-statement CE1-import term 1 from community CE1

[edit]
lab@R1# set policy-options policy-statement CE1-import term 1 then accept

[edit]
lab@R1# set policy-options community CE1 members target:54591:100
```

Unlike using BGP as the PE-CE protocol, where routes received from the MPLS core are automatically advertised to the CE routers, using OSPF you need explicit policy to match BGP routes and export them to OSPF.

```
[edit]
lab@R1# set policy-options policy-statement CE1-bgp-to-ospf term 1 from protocol bgp

[edit]
lab@R1# set policy-options policy-statement CE1-bgp-to-ospf term 1 then accept

[edit]
lab@R1# set routing-instances CE1 protocols ospf export CE1-bgp-to-ospf

[edit]
lab@R1# commit
commit complete
```

b) Configure R7 for customer CE1 L3VPN.

```
[edit]
lab@R7# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R7# set interfaces ge-0/0/5 unit 324 vlan-id 324

[edit]
lab@R7# set interfaces ge-0/0/5 unit 324 description "CE1-1 connection"

[edit]
lab@R7# set interfaces ge-0/0/5 unit 324 family inet address 192.168.0.93/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set routing-options route-distinguisher-id 172.30.5.7

[edit]
lab@R7# set routing-instances CE1 instance-type vrf

[edit]
lab@R7# set routing-instances CE1 interface ge-0/0/5.324

[edit]
lab@R7# set routing-instances CE1 vrf-import CE1-import

[edit]
lab@R7# set routing-instances CE1 vrf-export CE1-export

[edit]
lab@R7# set routing-instances CE1 protocols ospf area 0 interface all

[edit]
lab@R7# set policy-options policy-statement CE1-import term 1 from protocol bgp

[edit]
lab@R7# set policy-options policy-statement CE1-import term 1 from community CE1

[edit]
lab@R7# set policy-options policy-statement CE1-import term 1 then accept

[edit]
lab@R7# set policy-options policy-statement CE1-export term 1 from protocol ospf

[edit]
lab@R7# set policy-options policy-statement CE1-export term 1 then community add CE1

[edit]
lab@R7# set policy-options policy-statement CE1-export term 1 then accept

[edit]
lab@R7# set policy-options community CE1 members target:54591:100

[edit]
lab@R7# set policy-options policy-statement CE1-bgp-to-ospf term 1 from protocol bgp

[edit]
lab@R7# set policy-options policy-statement CE1-bgp-to-ospf term 1 then accept

[edit]
lab@R7# set routing-instances CE1 protocols ospf export CE1-bgp-to-ospf

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# commit
commit complete
```

c) Verify the operation of customer CE1 L3VPN.

The OSPF neighbors are in FULL state with both CE routers.

```
[edit]
lab@R1# run show ospf neighbor instance CE1
Address          Interface              State     ID               Pri  Dead
192.168.0.70     ge-0/0/5.318           Full      172.31.63.2      128    32
```

```
[edit]
lab@R7# run show ospf neighbor instance CE1
Address          Interface              State     ID               Pri  Dead
192.168.0.94     ge-0/0/5.324           Full      172.31.63.1      128    38
```

Verify the routing instance table on both PE routers.

```
[edit]
lab@R1# run show route table CE1.inet.0

CE1.inet.0: 16 destinations, 29 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.48.0/30     *[OSPF/10] 00:08:10, metric 6
                    > to 192.168.0.70 via ge-0/0/5.318
                    [BGP/170] 00:01:09, MED 6, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R7
                      to 172.30.0.2 via ae0.0, label-switched-path Bypass-
>172.30.0.6
172.31.48.4/30     *[OSPF/10] 00:08:10, metric 2
                    > to 192.168.0.70 via ge-0/0/5.318
                    [BGP/170] 00:01:09, MED 7, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R7
                      to 172.30.0.2 via ae0.0, label-switched-path Bypass-
>172.30.0.6
...
```

```
[edit]
lab@R7# run show route table CE1.inet.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
CE1.inet.0: 16 destinations, 29 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.48.0/30     *[OSPF/10] 00:01:02, metric 6
                    > to 192.168.0.94 via ge-0/0/5.324
                    [BGP/170] 00:00:32, MED 6, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.5 via ge-0/0/4.117, label-switched-path R7-to-
R1
                      to 172.30.0.29 via ge-0/0/4.137, label-switched-path
Bypass->172.30.0.5
172.31.48.4/30     *[OSPF/10] 00:01:02, metric 7
                    > to 192.168.0.94 via ge-0/0/5.324
                    [BGP/170] 00:00:32, MED 2, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.5 via ge-0/0/4.117, label-switched-path R7-to-
R1
                      to 172.30.0.29 via ge-0/0/4.137, label-switched-path
Bypass->172.30.0.5
...
```

OSPF routes are received from the MPLS backbone and from a local OSPF neighbor. Customer CE1 has backdoor link, this is the reason why the same routes are received in both routing instances via the local OSPF neighbors. You have applied an OSPF export policy to advertise the routes received from BGP to the local CE routers, however those routes are not active because of higher route preference of BGP.

```
[edit]
lab@R1# run show ospf database instance CE1

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.31.63.1      172.31.63.1      0x8000000f   647  0x22 0xfc39  72
Router   172.31.63.2      172.31.63.2      0x80000011  1035  0x22 0xaf4  120
Router   172.31.63.3      172.31.63.3      0x8000000f  1039  0x22 0xe59d 120
Router   172.31.63.4      172.31.63.4      0x8000000d  1040  0x22 0xc719  72
Router   172.31.63.5      172.31.63.5      0x8000000c  1061  0x22 0x2f57  84
Router  *192.168.0.69     192.168.0.69     0x80000005    37  0x22 0x1e52  36
Router   192.168.0.93     192.168.0.93     0x80000005    21  0x22 0xd839  36
Network  192.168.0.70     172.31.63.2      0x80000001  1035  0x22 0x4480  32
Network  192.168.0.94     172.31.63.1      0x80000001   647  0x22 0xa0f5  32
```

Only an OSPF Router LSA Type 1 is advertised by R1 to the local OSPF neighbor.

To make the BGP routes active in the routing table you can make the route preference of the OSPF routes higher than BGP routes.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set routing-instances CE1 protocols ospf preference 180

[edit]
lab@R1# commit
commit complete

[edit]
lab@R7# set routing-instances CE1 protocols ospf preference 180

[edit]
lab@R7# commit
commit complete
```

The OSPF routes are active in the routing table because of better preference.

```
[edit]
lab@R1# run show route table CE1.inet.0

CE1.inet.0: 16 destinations, 29 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.48.0/30     *[BGP/170] 00:01:50, MED 6, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R7
                      to 172.30.0.2 via ae0.0, label-switched-path Bypass-
>172.30.0.6
                    [OSPF/180] 00:03:24, metric 6
                    > to 192.168.0.70 via ge-0/0/5.318
172.31.48.4/30     *[BGP/170] 00:01:50, MED 7, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R7
                      to 172.30.0.2 via ae0.0, label-switched-path Bypass-
>172.30.0.6
                    [OSPF/180] 00:03:24, metric 2
                    > to 192.168.0.70 via ge-0/0/5.318
172.31.48.8/30     *[BGP/170] 00:01:50, MED 7, localpref 100, from 172.30.5.41
                      AS path: I

                    > to 172.30.0.6 via ge-0/0/4.117, label-switched-path R1-to-
R7
                      to 172.30.0.2 via ae0.0, label-switched-path Bypass-
>172.30.0.6
                    [OSPF/180] 00:03:24, metric 2
                    > to 192.168.0.70 via ge-0/0/5.318
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

After changing the preference of the routes both PEs export the BGP routes to the local OSPF neighbor.

```
[edit]
lab@R1# run show ospf database instance CE1

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.31.63.1      172.31.63.1      0x8000000f   754  0x22 0xfc39  72
Router   172.31.63.2      172.31.63.2      0x80000011  1142  0x22 0xaf4  120
Router   172.31.63.3      172.31.63.3      0x8000000f  1146  0x22 0xe59d 120
Router   172.31.63.4      172.31.63.4      0x8000000d  1147  0x22 0xc719  72
Router   172.31.63.5      172.31.63.5      0x8000000c  1168  0x22 0x2f57  84
Router  *192.168.0.69     192.168.0.69     0x80000007    14  0x22 0x1a54  36
Router   192.168.0.93     192.168.0.93     0x80000006    58  0x22 0xd63a  36
Network  192.168.0.70     172.31.63.2      0x80000001  1142  0x22 0x4480  32
Network  192.168.0.94     172.31.63.1      0x80000001   754  0x22 0xa0f5  32
Summary *172.31.48.0      192.168.0.69     0x80000001    14  0xa2 0x34da  28
Summary *172.31.48.4      192.168.0.69     0x80000001    14  0xa2 0x16f3  28
Summary *172.31.48.8      192.168.0.69     0x80000001    14  0xa2 0xed18  28
Summary *172.31.48.12     192.168.0.69     0x80000001    14  0xa2 0xcf31  28
Summary *172.31.48.16     192.168.0.69     0x80000001    14  0xa2 0xcf29  28
Summary *172.31.63.1      192.168.0.69     0x80000001    14  0xa2 0x649c  28
Summary *172.31.63.2      192.168.0.69     0x80000001    14  0xa2 0x8c6e  28
Summary *172.31.63.3      192.168.0.69     0x80000001    14  0xa2 0x8c6c  28
Summary *172.31.63.4      192.168.0.69     0x80000001    14  0xa2 0xb43e  28
Summary *172.31.63.5      192.168.0.69     0x80000001    14  0xa2 0x787e  28
Summary *192.168.0.72     192.168.0.69     0x80000001    14  0xa2 0x1047  28
Summary *192.168.0.84     192.168.0.69     0x80000001    14  0xa2 0xc97c  28
```

```
[edit]
lab@R7# run show ospf database instance CE1

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.31.63.1      172.31.63.1      0x8000000f   765  0x22 0xfc39  72
Router   172.31.63.2      172.31.63.2      0x80000011  1156  0x22 0xaf4  120
Router   172.31.63.3      172.31.63.3      0x8000000f  1160  0x22 0xe59d 120
Router   172.31.63.4      172.31.63.4      0x8000000d  1161  0x22 0xc719  72
Router   172.31.63.5      172.31.63.5      0x8000000c  1182  0x22 0x2f57  84
Router   192.168.0.69     192.168.0.69     0x80000007    29  0x22 0x1a54  36
Router  *192.168.0.93     192.168.0.93     0x80000007    16  0x22 0xd43b  36
Network  192.168.0.70     172.31.63.2      0x80000001  1156  0x22 0x4480  32
Network  192.168.0.94     172.31.63.1      0x80000001   765  0x22 0xa0f5  32
Summary  *172.31.48.0      192.168.0.93     0x80000001    29  0xa2 0x34da  28
Summary  *172.31.48.4      192.168.0.93     0x80000001    29  0xa2 0x16f3  28
Summary  *172.31.48.8      192.168.0.93     0x80000001    29  0xa2 0xed18  28
Summary  *172.31.48.12     192.168.0.93     0x80000001    29  0xa2 0xcf31  28
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Summary   *172.31.48.16      192.168.0.93      0x80000001    29   0xa2 0xcf29   28
Summary   *172.31.63.1       192.168.0.93      0x80000001    29   0xa2 0x649c   28
Summary   *172.31.63.2       192.168.0.93      0x80000001    29   0xa2 0x8c6e   28
Summary   *172.31.63.3       192.168.0.93      0x80000001    29   0xa2 0x8c6c   28
Summary   *172.31.63.4       192.168.0.93      0x80000001    29   0xa2 0xb43e   28
Summary   *172.31.63.5       192.168.0.93      0x80000001    29   0xa2 0x787e   28
Summary   *192.168.0.72      192.168.0.93      0x80000001    29   0xa2 0x1047   28
Summary   *192.168.0.84      192.168.0.93      0x80000001    29   0xa2 0xc97c   28
```

The routes are advertised as OSPF Summary Type 3 LSA, Inter Area routes.

Why are the routes advertised as Summary OSPF routes? This depends on two communities. These two communities are the `domain-id` and `rte-type`. Domain-id is used to signal if routes received from other sites are part of the same OSPF domain or not. The `rte-type` community signals the type of the route received from the local OSPF process. This community is always attached to the OSPF routes advertised to the other PEs via Multiprotocol BGP.

```
[edit]
lab@R1# run show route advertising-protocol bgp 172.30.5.41 table CE1.inet.0
extensive
CE1.inet.0: 16 destinations, 29 routes (16 active, 0 holddown, 0 hidden)
* 172.31.48.0/30 (2 entries, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.30.5.1:4
     VPN Label: 301792
     Nexthop: Self
     Flags: Nexthop Change
     MED: 6
     Localpref: 100
     AS path: [54591] I
     Communities: target:54591:100 rte-type:0.0.0.0:1:0
```

For example, R1 is advertising the routes with `rte-type` community `rte-type:0.0.0.0:1:0`. Intra-area routes from LSA 1.

If the `domain-id` match on each side of the L3VPN, then the OSPF routes appear as Summary Type 3 LSAs. If the domain-id does not match, then the OSPF routes appear as External Type 5 LSAs. If the domain-id is not attached, then Junos will see this as "matching" as the value will default to 0.

According to the requirements you have to make sure that customer routes appear as OSPF External in the L3VPN. To accomplish this, you must ensure that the routes from both VPN sites are part of a different OSPF domain.

© 2020 Juniper Networks, Inc. All rights reserved.

Configure a different local domain-id on R1 and R7.

R1 is configured with domain-id 2.

```
[edit]
lab@R1# set routing-instances CE1 protocols ospf domain-id 2

[edit]
lab@R1# commit
commit complete
```

R7 is configured with domain-id 1.

```
[edit]
lab@R7# set routing-instances CE1 protocols ospf domain-id 1

[edit]
lab@R7# commit
commit complete
```

Next you have to attach the local domain-id community to the routes by modifying the vrf-export policy.

```
[edit]
lab@R1# set policy-options policy-statement CE1-export term 1 then community add CE1-
domain

[edit]
lab@R1# set policy-options community CE1-domain members domain:2:0

[edit]
lab@R1# commit
commit complete
```

Modifying the policy on router R7.

```
[edit]
lab@R7# set policy-options policy-statement CE1-export term 1 then community add CE1-
domain

[edit]
lab@R7# set policy-options community CE1-domain members domain:1:0

[edit]
lab@R7# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

Verify if the routes are advertised with the domain-id attached.

```
[edit]
lab@R1# run show route advertising-protocol bgp 172.30.5.41 table CE1.inet.0
extensive

CE1.inet.0: 16 destinations, 29 routes (16 active, 0 holddown, 0 hidden)
* 172.31.48.0/30 (2 entries, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.30.5.1:4
     VPN Label: 301792
     Nexthop: Self
     Flags: Nexthop Change
     MED: 6
     Localpref: 100
     AS path: [54591] I
     Communities: target:54591:100 domain-id:2:0 rte-type:0.0.0.0:1:0
```

After making the domain-id different on both sites, the PE routers are advertising External LSA to their OSPF neighbors.

```
[edit]
lab@R1# run show ospf database instance CE1 | no-more
    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.31.63.1      172.31.63.1      0x8000000f   846  0x22 0xfc39  72
Router   172.31.63.2      172.31.63.2      0x80000011  1234  0x22 0xaf4  120
Router   172.31.63.3      172.31.63.3      0x8000000f  1238  0x22 0xe59d 120
Router   172.31.63.4      172.31.63.4      0x8000000d  1239  0x22 0xc719  72
Router   172.31.63.5      172.31.63.5      0x8000000c  1260  0x22 0x2f57  84
Router  *192.168.0.69     192.168.0.69     0x80000008    31  0x22 0x1855  36
Router   192.168.0.93     192.168.0.93     0x80000008    15  0x22 0xd23c  36
Network  192.168.0.70     172.31.63.2      0x80000001  1234  0x22 0x4480  32
Network  192.168.0.94     172.31.63.1      0x80000001   846  0x22 0xa0f5  32
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *172.31.48.0      192.168.0.69     0x80000001    12  0xa2 0xa15   36
Extern  *172.31.48.4      192.168.0.69     0x80000001    11  0xa2 0xeb2e  36
Extern  *172.31.48.8      192.168.0.69     0x80000001    11  0xa2 0xc352  36
Extern  *172.31.48.12     192.168.0.69     0x80000001    11  0xa2 0xa56b  36
Extern  *172.31.48.16     192.168.0.69     0x80000001    11  0xa2 0xa563  36
Extern  *172.31.63.1      192.168.0.69     0x80000001    11  0xa2 0x3ad6  36
Extern  *172.31.63.2      192.168.0.69     0x80000001    12  0xa2 0x62a8  36
Extern  *172.31.63.3      192.168.0.69     0x80000001    11  0xa2 0x62a6  36
Extern  *172.31.63.4      192.168.0.69     0x80000001    11  0xa2 0x8a78  36
Extern  *172.31.63.5      192.168.0.69     0x80000001    11  0xa2 0x4eb8  36
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Extern   *192.168.0.72     192.168.0.69     0x80000001    11  0xa2 0xe581  36
Extern   *192.168.0.84     192.168.0.69     0x80000001    10  0xa2 0x9fb6  36
```

```
[edit]
lab@R7# run show ospf database instance CE1 | no-more

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.31.63.1      172.31.63.1      0x8000000f   857  0x22 0xfc39  72
Router   172.31.63.2      172.31.63.2      0x80000011  1248  0x22 0xaf4  120
Router   172.31.63.3      172.31.63.3      0x8000000f  1252  0x22 0xe59d 120
Router   172.31.63.4      172.31.63.4      0x8000000d  1253  0x22 0xc719  72
Router   172.31.63.5      172.31.63.5      0x8000000c  1274  0x22 0x2f57  84
Router   192.168.0.69     192.168.0.69     0x80000008    46  0x22 0x1855  36
Router  *192.168.0.93     192.168.0.93     0x80000008    24  0x22 0xd23c  36
Network  192.168.0.70     172.31.63.2      0x80000001  1248  0x22 0x4480  32
Network  192.168.0.94     172.31.63.1      0x80000001   857  0x22 0xa0f5  32
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   *172.31.48.0      192.168.0.93     0x80000001    27  0xa2 0xa15   36
Extern   *172.31.48.4      192.168.0.93     0x80000001    27  0xa2 0xeb2e  36
Extern   *172.31.48.8      192.168.0.93     0x80000001    27  0xa2 0xc352  36
Extern   *172.31.48.12     192.168.0.93     0x80000001    27  0xa2 0xa56b  36
Extern   *172.31.48.16     192.168.0.93     0x80000001    27  0xa2 0xa563  36
Extern   *172.31.63.1      192.168.0.93     0x80000001    27  0xa2 0x3ad6  36
Extern   *172.31.63.2      192.168.0.93     0x80000001    27  0xa2 0x62a8  36
Extern   *172.31.63.3      192.168.0.93     0x80000001    27  0xa2 0x62a6  36
Extern   *172.31.63.4      192.168.0.93     0x80000001    27  0xa2 0x8a78  36
Extern   *172.31.63.5      192.168.0.93     0x80000001    27  0xa2 0x4eb8  36
Extern   *192.168.0.72     192.168.0.93     0x80000001    27  0xa2 0xe581  36
Extern   *192.168.0.84     192.168.0.93     0x80000001    26  0xa2 0x9fb6  36
```

4) Configure customer CE2 L3VPN.

Customer CE2 has two sites connected to the local AS and one connected via partner AS. S1 connected to R7 should be used as the central site in a hub-and-spoke fashion to exchange traffic between the spoke sites.

a) Configure router R4 as spoke.

Because all traffic will be exchanged through the hub site it is very likely that the AS path of the routes will hold the local AS number up to two or three times. You must allow routes which contain the local AS number on the RR and on all PEs that connect to customer CE2 sites.

```
[edit]
lab@R4# set routing-options autonomous-system 54591 loops 3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R4# set interfaces ge-0/0/5.323 vlan-id 323

[edit]
lab@R4# set interfaces ge-0/0/5.323 description "CE2-2 connection"

[edit]
lab@R4# set interfaces ge-0/0/5.323 family inet address 192.168.0.89/30
```

Configure an additional loopback IP address for the CE2 routing instance.

```
[edit]
lab@R4# set interfaces lo0.1 family inet address 172.30.5.21/32

[edit]
lab@R4# set routing-options route-distinguisher-id 172.30.5.4

[edit]
lab@R4# set routing-instances CE2-spoke instance-type vrf

[edit]
lab@R4# set routing-instances CE2-spoke interface ge-0/0/5.323

[edit]
lab@R4# set routing-instances CE2-spoke interface lo0.1
```

Configure vrf-table-label to advertise a single label per vrf. This will enable connectivity to the directly connected networks in the VRF.

```
[edit]
lab@R4# set routing-instances CE2-spoke vrf-table-label

[edit]
lab@R4# set routing-instances CE2-spoke vrf-import CE2-spoke-import

[edit]
lab@R4# set routing-instances CE2-spoke vrf-export CE2-spoke-export

[edit]
lab@R4# set routing-instances CE2-spoke protocols bgp group ce peer-as 64600

[edit]
lab@R4# set routing-instances CE2-spoke protocols bgp group ce neighbor 192.168.0.90
```


© 2020 Juniper Networks, Inc. All rights reserved.

You must configure specific vrf-import and vrf-export policies to match different route target communities. The spoke community should be attached in the export direction whereas the hub community should be matched on the import direction. This will force all the traffic to be exchanged only through the hub site.

```
[edit]
lab@R4# set policy-options policy-statement CE2-spoke-import term 1 from protocol bgp

[edit]
lab@R4# set policy-options policy-statement CE2-spoke-import term 1 from community CE2-
hub

[edit]
lab@R4# set policy-options policy-statement CE2-spoke-import term 1 then accept

[edit]
lab@R4# set policy-options policy-statement CE2-spoke-export term 1 from protocol [
direct bgp ]

[edit]
lab@R4# set policy-options policy-statement CE2-spoke-export term 1 then community add
CE2-spoke

[edit]
lab@R4# set policy-options policy-statement CE2-spoke-export term 1 then accept

[edit]
lab@R4# set policy-options community CE2-hub members target:54591:200

[edit]
lab@R4# set policy-options community CE2-spoke members target:54591:201

[edit]
lab@R4# commit
commit complete
```

b) Configure router R7 as the hub site.

The hub site should have two routing instances, one hub and one spoke. The hub routing instances will advertise routes to the spoke sites, whereas the spoke instance will receive the route advertisements from the spoke sites.

```
[edit]
lab@R7# set interfaces ge-0/0/5.311 vlan-id 311

[edit]
lab@R7# set interfaces ge-0/0/5.311 description "CE2-1 connection hub"
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set interfaces ge-0/0/5.311 family inet address 192.168.0.41/30

[edit]
lab@R7# set interfaces ge-0/0/5.312 vlan-id 312

[edit]
lab@R7# set interfaces ge-0/0/5.312 description "CE2-1 connection spoke"

[edit]
lab@R7# set interfaces ge-0/0/5.312 family inet address 192.168.0.45/30

[edit]
lab@R7# set interfaces lo0.1 family inet address 172.30.5.33/32

[edit]
lab@R7# set interfaces lo0.2 family inet address 172.30.5.34/32

[edit]
lab@R7# set routing-options autonomous-system 54591 loops 3
```

Configure the hub routing instance and route target policies.

```
[edit]
lab@R7# set routing-instances CE2-hub instance-type vrf

[edit]
lab@R7# set routing-instances CE2-hub interface ge-0/0/5.311

[edit]
lab@R7# set routing-instances CE2-hub interface lo0.1

[edit]
lab@R7# set routing-instances CE2-hub vrf-table-label

[edit]
lab@R7# set routing-instances CE2-hub vrf-import CE2-hub-import

[edit]
lab@R7# set routing-instances CE2-hub vrf-export CE2-hub-export

[edit]
lab@R7# set routing-instances CE2-hub protocols bgp group ce peer-as 64600 neighbor
192.168.0.42
```

`vrf-import` policy has a single term which rejects everything.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set policy-options policy-statement CE2-hub-import term 1 then reject

[edit]
lab@R7# set policy-options policy-statement CE2-hub-export term 1 from protocol [ bgp
direct ]

[edit]
lab@R7# set policy-options policy-statement CE2-hub-export term 1 then community add
CE2-hub

[edit]
lab@R7# set policy-options policy-statement CE2-hub-export term 1 then accept
```

Configure the spoke routing instance and route target policies.

```
[edit]
lab@R7# set routing-instances CE2-spoke instance-type vrf

[edit]
lab@R7# set routing-instances CE2-spoke interface lo0.2

[edit]
lab@R7# set routing-instances CE2-spoke interface ge-0/0/5.312

[edit]
lab@R7# set routing-instances CE2-spoke vrf-table-label

[edit]
lab@R7# set routing-instances CE2-spoke vrf-import CE2-spoke-import

[edit]
lab@R7# set routing-instances CE2-spoke vrf-export CE2-spoke-export

[edit]
lab@R7# set routing-instances CE2-spoke protocols bgp group ce peer-as 64600

[edit]
lab@R7# set routing-instances CE2-spoke protocols bgp group ce neighbor 192.168.0.46
```

The spoke routing instance does not export any routes. its role is to accept routes from other spoke sites and send it to the HUB CE.

```
[edit]
lab@R7# set policy-options policy-statement CE2-spoke-export term 1 then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set policy-options policy-statement CE2-spoke-import term 1 from protocol bgp

[edit]
lab@R7# set policy-options policy-statement CE2-spoke-import term 1 from community CE2-
spoke

[edit]
lab@R7# set policy-options policy-statement CE2-spoke-import term 1 then accept

[edit]
lab@R7# set policy-options community CE2-hub members target:54591:200

[edit]
lab@R7# set policy-options community CE2-spoke members target:54591:201

[edit]
lab@R7# commit
commit complete
```

Verify the operation of the hub and spoke L3VPN of customer CE2.

The Hub routing instance routing table on R7 is populated with routes received from the CE router.

```
[edit]
lab@R7# run show route table CE2-hub.inet.0

CE2-hub.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.21/32     *[BGP/170] 00:02:48, localpref 100
                      AS path: 64600 54591 I, validation-state: unverified
                    > to 192.168.0.42 via ge-0/0/5.311
172.30.5.33/32     *[Direct/0] 00:02:57
                    > via lo0.1
172.31.64.0/24     *[BGP/170] 00:02:50, localpref 100
                      AS path: 64600 I, validation-state: unverified
                    > to 192.168.0.42 via ge-0/0/5.311
172.31.65.0/24     *[BGP/170] 00:02:50, localpref 100
                      AS path: 64600 I, validation-state: unverified
                    > to 192.168.0.42 via ge-0/0/5.311
172.31.66.0/24     *[BGP/170] 00:02:50, localpref 100
                      AS path: 64600 I, validation-state: unverified
                    > to 192.168.0.42 via ge-0/0/5.311
172.31.67.0/24     *[BGP/170] 00:02:50, localpref 100
                      AS path: 64600 I, validation-state: unverified
                    > to 192.168.0.42 via ge-0/0/5.311
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
...
```

The spoke routing instance receives routes from R4 through Multiprotocol BGP.

```
[edit]
lab@R7# run show route table CE2-spoke.inet.0

CE2-spoke.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.21/32     *[BGP/170] 00:02:41, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.45 via ge-0/0/4.167, label-switched-path R7-to-
R3
172.30.5.34/32     *[Direct/0] 00:25:13
                    > via lo0.2
172.31.76.0/24     *[BGP/170] 00:02:41, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.45 via ge-0/0/4.167, label-switched-path R7-to-
R3
172.31.77.0/24     *[BGP/170] 00:02:41, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.45 via ge-0/0/4.167, label-switched-path R7-to-
R3
...
```

Verify the EBGP advertisements to the CE router from the spoke routing instance.

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.46

CE2-spoke.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.5.21/32          Self                                    I
* 192.168.0.88/30         Self                                    I
```

No routes are advertised except the local routes on the R4 spoke routing instance. The reason is the default behavior of Junos to not advertise route learned from EBGP back to EBGP peer from the same AS. To change this behavior, you can use **as-override** knob. This will override the AS path received from R4 with the local AS number and will allow Junos to advertise the routes to the CE router.

```
[edit]
lab@R7# set routing-instances CE2-spoke protocols bgp group ce as-override

[edit]
lab@R7# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

```
[edit]
lab@R7# run show route advertising-protocol bgp 192.168.0.46

CE2-spoke.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.5.21/32          Self                                    I
* 172.31.76.0/24          Self                                    54591 I
* 172.31.77.0/24          Self                                    54591 I
* 192.168.0.88/30         Self                                    I
```

After modifying the default Junos behavior the routes are advertised with changed AS path.

On R4 routes received from the R7 hub routing instance are inserted into the CE2-spoke routing table because of the hub route target.

```
[edit]
lab@R4# run show route table CE2-spoke.inet.0

CE2-spoke.inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.21/32     *[Direct/0] 00:04:57
                    > via lo0.1
172.30.5.33/32     *[BGP/170] 00:00:47, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.81 via ae0.0, Push 16, Push 299808(top)
172.31.64.0/24     *[BGP/170] 00:00:47, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.81 via ae0.0, Push 16, Push 299808(top)
172.31.65.0/24     *[BGP/170] 00:00:47, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.81 via ae0.0, Push 16, Push 299808(top)
172.31.66.0/24     *[BGP/170] 00:00:47, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.81 via ae0.0, Push 16, Push 299808(top)
172.31.67.0/24     *[BGP/170] 00:00:47, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.81 via ae0.0, Push 16, Push 299808(top)
172.31.68.0/24     *[BGP/170] 00:00:47, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.81 via ae0.0, Push 16, Push 299808(top)
172.31.69.0/24     *[BGP/170] 00:00:47, localpref 100, from 172.30.5.41
                      AS path: 64600 I
                    > to 172.30.0.81 via ae0.0, Push 16, Push 299808(top)
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

You need to overwrite the AS path also on R4.

```
[edit]
lab@R4# set routing-instances CE2-spoke protocols bgp group ce as-override

[edit]
lab@R4# commit
commit complete
```

```
[edit]
lab@R4# run show route advertising-protocol bgp 192.168.0.90

CE2-spoke.inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.30.5.33/32          Self                                    I
* 172.31.64.0/24          Self                                    54591 I
* 172.31.65.0/24          Self                                    54591 I
* 172.31.66.0/24          Self                                    54591 I
* 172.31.67.0/24          Self                                    54591 I
* 172.31.68.0/24          Self                                    54591 I
* 172.31.69.0/24          Self                                    54591 I
* 172.31.70.0/24          Self                                    54591 I
* 172.31.71.0/24          Self                                    54591 I
* 192.168.0.40/30         Self                                    I
```

5) Allow route exchange between customer CE1 site S1 and customer CE2 site S1 at R7.

You have to enable route exchange between both customers on R7. You cannot use `rib-groups` to accomplish this requirement, so you have to use another solution. The `auto-export` feature can be used to copy routes between the customer routing instances. Unlike `rib-groups`, `auto-export` uses route target communities to import routes from one routing table to the other.

Enable auto-export on customer CE1 and CE2. You have to enable auto-export on both the hub and spoke routing instances for customer CE2. Routes from customer CE1 should be copied to the spoke routing instance and routes from hub routing instance should be copied to customer CE1. This way you will preserve the hub-and-spoke design of customer CE2.

```
[edit]
lab@R7# set routing-instances CE1 routing-options auto-export

[edit]
lab@R7# set routing-instances CE2-hub routing-options auto-export
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set routing-instances CE2-spoke routing-options auto-export
```

Configure new route target community and modifying the vrf-target policies.

```
[edit]
lab@R7# set policy-options community exchange members target:54591:111

[edit]
lab@R7# set policy-options policy-statement CE1-import term 2 from community exchange

[edit]
lab@R7# set policy-options policy-statement CE1-import term 2 then accept

[edit]
lab@R7# set policy-options policy-statement CE1-export term 1 then community add
exchange

[edit]
lab@R7# set policy-options policy-statement CE2-hub-export term 1 then community add
exchange

[edit]
lab@R7# set policy-options policy-statement CE2-spoke-import term 2 from community
exchange

[edit]
lab@R7# set policy-options policy-statement CE2-spoke-import term 2 then accept

[edit]
lab@R7# commit
commit complete
```

Verify if the routes are copied between the routing instances.

```
[edit]
lab@R7# run show route table CE1.inet.0 172.31.64.0/24 extensive

CE1.inet.0: 26 destinations, 39 routes (26 active, 0 holddown, 0 hidden)
172.31.64.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.31.64.0/24 -> {192.168.0.42}
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 759
                Address: 0x155cc88
                Next-hop reference count: 48
                Source: 192.168.0.42
                Next hop: 192.168.0.42 via ge-0/0/5.311, selected
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
State: <Secondary Active Ext>
Peer AS: 64600
Age: 4:09
Task: BGP_64600.192.168.0.42+179
Announcement bits (1): 1-KRT
AS path: 64600 I
Communities: target:54591:111 target:54591:200
Accepted
Localpref: 100
Router ID: 172.31.71.1
Primary Routing Table CE2-hub.inet.0
```

BGP routes from the hub routing instance are copied to the CE1 routing instance.

```
[edit]
lab@R7# run show route table CE2-spoke.inet.0 192.168.0.68/30 extensive

CE2-spoke.inet.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)

192.168.0.68/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.68/30 -> {192.168.0.94}
        *OSPF   Preference: 180
                Next hop type: Router, Next hop index: 618
                Address: 0x155dcdc
                Next-hop reference count: 17
                Next hop: 192.168.0.94 via ge-0/0/5.324, selected
                State: <Secondary Active Int>
                Age: 1:51       Metric: 7
                Area: 0.0.0.0
                Task: CE1-OSPF
                Announcement bits (1): 0-KRT
                AS path: I
                Communities: target:54591:100 target:54591:111 domain-id:1:0 rte-
type:0.0.0.0:1:0
                Primary Routing Table CE1.inet.0
```

The OSPF routes are copied from the CE1 routing instance to the spoke routing instance of customer CE2. Primary routing table shown is `CE1.inet.0`.

You have to make sure that the copied routes are not advertised to the remote VRFs on other PEs. Unlike using `rib-groups` when `auto-export` is used, the copied routes are not eligible for advertisement.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show route advertising-protocol bgp 172.30.5.41 172.31.64.0/24

CE2-hub.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.31.64.0/24          Self                         100        64600 I
```

The route is advertised only from the primary routing table to the route reflector.

6) Customer CE2 VPN must be provided with Internet access at the customer site S1 using a single customer facing VRF interface.

You must configure a static route with a next hop to `inet.0`, which is allowed in for this step. This route should be created in the spoke routing instance and advertised to the CE router from where it will be then advertised to the hub routing instance. In the hub-routing instance it will be advertised to the other spoke sites. One more static route should be created in `inet.0`, which summarizes all the routes used by CE2 customer. This summary route should be advertised to the Internet as well.

You must copy the direct interface routes from CE2-hub routing instance to `inet.0`. Having the interface route in `inet.0` will resolve the next hop of the summary route which will be configured later.

Traffic directed to the Internet will land in hub routing instance where a default route from the CE is received. Packets will reach the spoke routing instance to the CE router and from there they are able to reach the Internet. Traffic in reverse direction will hit the summary route and will be sent to the CE and from there to the right spoke.

```
[edit]
lab@R7# set routing-options rib-groups CE2-inet import-rib [ CE2-hub.inet.0 inet.0 ]

[edit]
lab@R7# set routing-instances CE2-hub routing-options interface-routes rib-group inet
CE2-inet

[edit]
lab@R7# set routing-instances CE2-spoke routing-options static route 0/0 next-table
inet.0

[edit]
lab@R7# set policy-options policy-statement default-to-ce term 1 from protocol static
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set policy-options policy-statement default-to-ce term 1 from route-filter 0/0
exact

[edit]
lab@R7# set policy-options policy-statement default-to-ce term 1 then accept

[edit]
lab@R7# set routing-instances CE2-spoke protocols bgp group ce export default-to-ce
```

Configure static summary route in `inet.0` pointing to next-hop in CE2-hub routing instance.

```
[edit]
lab@R7# set routing-options static route 172.31.64.0/20 next-hop 192.168.0.42
```

Configure policy to advertise the summary route to the route reflector.

```
[edit]
lab@R7# set policy-options policy-statement CE2-routes term 1 from protocol static

[edit]
lab@R7# set policy-options policy-statement CE2-routes term 1 from route-filter
172.31.64.0/20 exact

[edit]
lab@R7# set policy-options policy-statement CE2-routes term 1 then accept

[edit]
lab@R7# set protocols bgp group ibgp export CE2-routes

[edit]
lab@R7# commit
commit complete
```

7) Configure NG MVPN in customer CE2 L3VPN.

Customer CE2 VPN has two physical interfaces to the CE router. One of the interfaces terminates in the hub routing instance, and another one terminates in the spoke routing instance. R4 is a spoke site with single spoke routing instance.

The customer hub route target is `target:54591:200` and spoke route target is `target:54591:201`. VPNv4 routes are exchanged as expected and the traffic is always passing hub site.

© 2020 Juniper Networks, Inc. All rights reserved.

This customer is requesting a multicast service over the MPLS core network. NG-MVPN using Multiprotocol BGP and Multicast LDP are chosen to provide this service to the customer. The multicast flow is from a multicast source towards a multicast receiver. Keep in mind that a hub-and-spoke topology does not provide efficient way to transport multicast traffic.

To providing multicast in customer C2 VPN only the spoke routing instance is configured for NG-MVPN. Different symmetric route targets are used on both R7 and R4 spoke routing instances. The route target used for the mvpn routes is `target:54591:202`.

a) Configure R7 for NG VPN.

The provider PE will act as the RP for the customer VPN.

```
[edit]
lab@R7# set interfaces lo0.2 family inet address 172.30.5.253/32

[edit]
lab@R7# set interfaces lo0.2 family inet address 172.30.5.34/32 primary

[edit]
lab@R7# set routing-instances CE2-spoke protocols pim interface all

[edit]
lab@R7# set routing-instances CE2-spoke protocols pim rp local address 172.30.5.253
```

Multicast LDP configuration is very simple compared to the RSVP P2MP LSPs. You just have to enable the `p2mp` option under the `protocols ldp` hierarchy. This should be added to all routers enabled for LDP. Because LDP was already enabled on some of the routers you have to ensure that there is an LDP signaled LSP from R4 and R7 using the physical interfaces. LDP tunneling over RSVP signaled LSPs cannot be used.

```
[edit]
lab@R7# set protocols ldp interface ge-0/0/4.137

[edit]
lab@R7# set protocols ldp p2mp
```

Configure the provider tunnel using Multicast LDP in the CE2-spoke routing-instance.

```
[edit]
lab@R7# set routing-instances CE2-spoke provider-tunnel ldp-p2mp
```

Configure `protocol mvpn` under the routing instance.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set routing-instances CE2-spoke protocols mvpn mvpn-mode spt-only

[edit]
lab@R7# set routing-instances CE2-spoke protocols mvpn route-target import-target
target target:54591:202

[edit]
lab@R7# set routing-instances CE2-spoke protocols mvpn route-target export-target
target target:54591:202
```

You must configure `family inet-mvpn signaling` on all MVPN PEs routers, including the route reflector.

```
[edit]
lab@R7# set protocols bgp group ibgp family inet-mvpn signaling

[edit]
lab@R7# commit
commit complete
```

b) Configure R4 for NG VPN.

```
[edit]
lab@R4# set interfaces lo0.1 family inet address 172.30.5.253/32

[edit]
lab@R4# set interfaces lo0.1 family inet address 172.30.5.21/32 primary

[edit]
lab@R4# set routing-instances CE2-spoke protocols pim interface all

[edit]
lab@R4# set routing-instances CE2-spoke protocols pim rp local address 172.30.5.253

[edit]
lab@R4# set routing-instances CE2-spoke protocols mvpn mvpn-mode spt-only

[edit]
lab@R4# set routing-instances CE2-spoke protocols mvpn route-target import-target
target target:54591:202

[edit]
lab@R4# set routing-instances CE2-spoke protocols mvpn route-target export-target
target target:54591:202

[edit]
lab@R4# set routing-instances CE2-spoke provider-tunnel ldp-p2mp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set protocols ldp p2mp

[edit]
lab@R4# set protocols bgp group ibgp family inet-mvpn signaling

[edit]
lab@R4# commit
commit complete
```

a) Configure route reflector for NLRI.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family inet-mvpn signaling

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family inet-mvpn signaling

[edit]
lab@route-reflector# commit
commit complete
```

Verify the operation of the NG MVPN.

Junos applies a hidden export and import route target policy to the routing instance. MVPN auto discovery routes are exchanged between R7 and R4 and are inserted in the right `mvpn.0` routing tables.

```
[edit]
lab@R4# run show route table CE2-spoke.mvpn.0

CE2-spoke.mvpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:172.30.5.4:4:172.30.5.4/240
                   *[MVPN/70] 00:02:25, metric2 1
                      Indirect
1:172.30.5.7:6:172.30.5.7/240
                   *[BGP/170] 00:00:56, localpref 100, from 172.30.5.41
                      AS path: I
                    > to 172.30.0.81 via ae0.0, Push 300240
```

```
[edit]
lab@R7# run show route table CE2-spoke.mvpn.0 extensive

CE2-spoke.mvpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
1:172.30.5.4:4:172.30.5.4/240 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                PMSI: Flags 0x0: Label 0: LDP-P2MP: Root 172.30.5.4 , lsp-id
16777217
                Next hop type: Indirect
                Address: 0x188122c
                Next-hop reference count: 2
                Source: 172.30.5.41
                Protocol next hop: 172.30.5.4
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 1:57       Metric2: 1
                Task: BGP_54591.172.30.5.41+179
                Announcement bits (2): 0-PIM.CE2-spoke 1-mvpn global task
                AS path: I (Originator)
                Cluster list:  0.0.0.2
                Originator ID: 172.30.5.4
                Communities: target:54591:202
                Import Accepted
                Localpref: 100
                Router ID: 172.30.5.41
                Primary Routing Table bgp.mvpn.0
                Indirect next hops: 1
                        Protocol next hop: 172.30.5.4 Metric: 1
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 172.30.0.45 via ge-0/0/4.167 weight 0x1
                                Next hop: 172.30.0.5 via ge-0/0/4.117 weight
0x8001
                        172.30.5.4/32 Originating RIB: inet.3
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 172.30.0.45 via ge-0/0/4.167
                                Nexthop: 172.30.0.5 via ge-0/0/4.117
...
```

These routes are used for auto discovery of the PEs participating in the particular MVPN, which is based on the route target configured under `protocols mvpn`.

Routes type 6: (Shared tree join route) and Type 7: (Source tree join route) are used for transporting customer PIM messages in MBGP are use two new extended communities, `rt-import` and `src-as`.

The `src-as` is used for Intra-AS operation and is always in the format `src-as:54591:0`. The local AS number is 54591 and 0 is always set to zero.

© 2020 Juniper Networks, Inc. All rights reserved.

The `rt-import` community is used to select the active sender. That community is randomly generated from the PE lo0.0 IP address and a unique number is assigned to the VRF. The format of the `rt-import` community is `rt-import:172.30.5.7:6`. These two communities are attached to all VPNv4 routes advertised to other PE routers.

PE routers with attached receivers which want to join a multicast group send type 7: prefix to the sender PE. This value is based on the rt-import communities attached to the VPNv4 routes of the multicast source.

The sender PE uses an internal hidden policy which imports routes with the route target community based on the random generated `rt-import` community.

```
[edit]
lab@R4# run show route receive-protocol bgp 172.30.5.41 table bgp.l3vpn.0
extensive
* 172.30.5.7:5:172.31.64.0/24 (1 entry, 0 announced)
     Import Accepted
     Route Distinguisher: 172.30.5.7:5
     VPN Label: 16
     Nexthop: 172.30.5.7
     Localpref: 100
     AS path: 64600 I (Originator)
     Cluster list:  0.0.0.1
     Originator ID: 172.30.5.7
     Communities: target:54591:111 target:54591:200
...
```

The problem with hub-and-spoke topology is that hub routing instance is not MVPN enabled. In the output above the `rt-import` community is not attached. So, it does not attach an `rt-import` community to the IPv4 routes. Even if you enable the `protocols mvpn` in hub routing instance which will attach `rt-import` community it will not help because the `rt-import` community is randomly generated for each VRF. The type 7: routes will be imported in hub routing instance, but Type 1: are imported to the spoke routing instance and the p2mp LDP label is attached to the spoke routing instance. So, enabling `protocol mvpn` in hub routing instance will not solve the problem.

To solve this problem, you can change the route target export policy on the hub routing instance to attach the `rt-import` community generated for the spoke routing instance.

To find the randomly generated rt-community you can check the internal hidden policies on R7.

Note: The randomly generated rt-community value will be different on your device so you will need to make adjustments to your policy to account for the value you get from the

© 2020 Juniper Networks, Inc. All rights reserved.

command below. The remaining examples are based on the *172.30.5.7:7* rt-community value.

```
[edit]
lab@R7# run show policy __vrf-mvpn-import-cmcast-CE2-spoke-internal__
Policy __vrf-mvpn-import-cmcast-CE2-spoke-internal__:
    Term unnamed:
        from community __vrf-mvpn-community-rt_import-target-CE2-spoke-internal__
[target:172.30.5.7:7 ]
        then accept
    Term unnamed:
        then reject
```

R7 is expecting customer multicast routes with the community `target:172.30.5.7:7` attached, which will be imported in the `spoke.mvpn.0` routing table.

Attach the `rt-import` community on R7's hub routing instance. For consistency you can attach the `src-as` community also.

```
[edit]
lab@R7# set policy-options community src-as members src-as:54591:0

[edit]
lab@R7# set policy-options community rt-import-spoke members rt-import:172.30.5.7:7

[edit]
lab@R7# set policy-options policy-statement CE2-hub-export term 1 then community add
rt-import-spoke

[edit]
lab@R7# set policy-options policy-statement CE2-hub-export term 1 then community add
src-as

[edit]
lab@R7# commit
commit complete
```

After modifying the policy, the IPv4 routes are received with both MVPN communities.

```
[edit]
lab@R4# run show route table CE2-spoke.inet.0 172.31.67.0/24 extensive

CE2-spoke.inet.0: 21 destinations, 23 routes (21 active, 0 holddown, 0 hidden)
172.31.67.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.31.67.0/24 -> {indirect(262153)}
        *BGP    Preference: 170/-101
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Route Distinguisher: 172.30.5.7:5
Next hop type: Indirect
Address: 0x155e78c
Next-hop reference count: 46
Source: 172.30.5.41
Next hop type: Router, Next hop index: 671
Next hop: 172.30.0.81 via ae0.0, selected
Label operation: Push 16, Push 299888(top)
Label TTL action: prop-ttl, prop-ttl(top)
Protocol next hop: 172.30.5.7
Push 16
Indirect next hop: 16ed790 262153
State: <Secondary Active Int Ext>
Local AS: 54591 Peer AS: 54591
Age: 9  Metric2: 15
Task: BGP_54591.172.30.5.41+179
Announcement bits (1): 0-KRT
AS path: 64600 I (Originator)
Cluster list:  0.0.0.1
Originator ID: 172.30.5.7
Communities: target:54591:111 target:54591:200 src-as:54591:0 rt-
import:172.30.5.7:7
Import Accepted
VPN Label: 16
Localpref: 100
Router ID: 172.30.5.41
Primary Routing Table bgp.l3vpn.0
Indirect next hops: 1
        Protocol next hop: 172.30.5.7 Metric: 15
        Push 16
        Indirect next hop: 16ed790 262153
        Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 172.30.0.81 via ae0.0
        172.30.5.7/32 Originating RIB: inet.3
          Metric: 15                        Node path count: 1
          Forwarding nexthops: 1
                Nexthop: 172.30.0.81 via ae0.0
```

Verify the operation of the MVPN generating a PIM join on R4.

```
[edit]
lab@R4# set protocols igmp interface ge-0/0/5.323 static group 232.1.1.1 source
172.31.65.5
```

Configure static IGMP join for group 232.1.1.1 with a source attached to the hub site.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set protocols igmp interface ge-0/0/5.323 version 3
```

Source specific multicast is supported only with IGMP version 3

```
[edit]
lab@R4# commit
commit complete
```

Verify if R4 is generating customer multicast route Type 7: with the correct community to R7.

```
[edit]
lab@R4# run show route table CE2-spoke.mvpn.0 extensive | find 232.1.1.1/240

7:172.30.5.7:5:54591:32:172.31.65.5:32:232.1.1.1/240 (1 entry, 1 announced)
TSI:
RTF { 1287:7/64 }
Page 0 idx 0 Type 1 val 17e1c38
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [54591] I
    Communities: target:172.30.5.7:7
Path 7:172.30.5.7:5:54591:32:172.31.65.5:32:232.1.1.1 Vector len 4.  Val: 0
        *PIM    Preference: 105
                Next hop type: Multicast (IPv4) Composite, Next hop index: 262155
                Address: 0x1524424
                Next-hop reference count: 3
                State: <Active Int Ext>
                Age: 1:18
                Task: PIM.CE2-spoke
                Announcement bits (4): 0-PIM.CE2-spoke 1-mvpn global task 2-BGP
Route Target 3-BGP_RT_Background
                AS path: I
                Communities: target:172.30.5.7:7
```

Correct community is attached for the spoke routing instance. Based on the Type 7: the prefix from R4 is imported into the CE2-spoke.mvpn.0.

```
[edit]
lab@R7# run show route table CE2-spoke.mvpn.0

CE2-spoke.mvpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:172.30.5.4:32767:172.30.5.4/240
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    *[BGP/170] 02:37:35, localpref 100, from 172.30.5.41
                      AS path: I
                     > to 172.30.0.45 via ge-0/0/4.167, label-switched-path R7-to-
R3
                       to 172.30.0.5 via ge-0/0/4.117, label-switched-path R7-to-
R3
1:172.30.5.7:32767:172.30.5.7/240
                    *[MVPN/70] 02:39:49, metric2 1
                      Indirect
7:172.30.5.7:5:54591:32:172.31.65.5:32:232.1.1.1/240
                    *[BGP/170] 00:00:27, localpref 100, from 172.30.5.41
                      AS path: I
                     > to 172.30.0.5 via ge-0/0/4.117, label-switched-path R7-to-
R1
                       to 172.30.0.45 via ge-0/0/4.167, label-switched-path R7-to-
R1
```

PIM join state is seen with the Pseudo-MVPN as the downstream interface.

```
[edit]
lab@R7# run show pim join instance CE2-spoke extensive
Instance: PIM.CE2-spoke Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
    Source: 172.31.65.5
    Flags: sparse,spt
    Upstream interface: ge-0/0/5.312
    Upstream neighbor: Direct
    Upstream state: Local Source, Local RP
    Keepalive timeout:
    Uptime: 00:01:24
    Downstream neighbors:
        Interface: Pseudo-MVPN
            Uptime: 00:01:24 Time since last Join: 00:01:24

Instance: PIM.CE2-spoke Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

Verify the data plain operation for multicast LDP.

```
[edit]
lab@R4# run show mvpn instance

MVPN instance:
Legend for provider tunnel
S-    Selective provider tunnel
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Family : INET

Instance : CE2-spoke
  MVPN Mode : SPT-ONLY
  Provider tunnel: I-P-tnl:LDP-P2MP:172.30.5.4, lsp-id 16777217
  Neighbor                              Inclusive Provider Tunnel
  172.30.5.7                            LDP-P2MP:172.30.5.7, lsp-id 16777217
  C-mcast IPv4 (S:G)              Provider Tunnel                          St
  172.31.65.5/32:232.1.1.1/32     LDP-P2MP:172.30.5.7, lsp-id 16777217

MVPN instance:
Legend for provider tunnel
S-    Selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Family : INET6

Instance : CE2-spoke
  MVPN Mode : SPT-ONLY
  Provider tunnel: I-P-tnl:LDP-P2MP:172.30.5.4, lsp-id 16777217
```

```
[edit]
lab@R4# run show multicast route instance CE2-spoke extensive
Instance: CE2-spoke Family: INET

Group: 232.1.1.1
    Source: 172.31.65.5/32
    Upstream interface: lsi.0
    Downstream interface list:
        ge-0/0/5.323
    Session description: Source specific multicast
    Statistics: 0 kBps, 0 pps, 0 packets
    Next-hop ID: 262155
    Upstream protocol: MVPN
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: forever
    Wrong incoming interface notifications: 0
    Uptime: 00:10:56

Instance: CE2-spoke Family: INET6
```

R4 is a leaf of the inclusive path sourced from R7.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show ldp database p2mp
Input label database, 172.30.5.4:0--172.30.5.3:0
  Label     Prefix
 300288     P2MP root-addr 172.30.5.4, lsp-id 16777217

Output label database, 172.30.5.4:0--172.30.5.3:0
  Label     Prefix
     16     P2MP root-addr 172.30.5.7, lsp-id 16777217

Input label database, 172.30.5.4:0--172.30.5.6:0

Output label database, 172.30.5.4:0--172.30.5.6:0
```

```
[edit]
lab@R3# run show ldp database p2mp
Input label database, 172.30.5.3:0--172.30.5.1:0

Output label database, 172.30.5.3:0--172.30.5.1:0

Input label database, 172.30.5.3:0--172.30.5.2:0

Output label database, 172.30.5.3:0--172.30.5.2:0

Input label database, 172.30.5.3:0--172.30.5.4:0
  Label     Prefix
     16     P2MP root-addr 172.30.5.7, lsp-id 16777217

Output label database, 172.30.5.3:0--172.30.5.4:0
  Label     Prefix
 300288     P2MP root-addr 172.30.5.4, lsp-id 16777217

Input label database, 172.30.5.3:0--172.30.5.5:0

Output label database, 172.30.5.3:0--172.30.5.5:0

Input label database, 172.30.5.3:0--172.30.5.7:0
  Label     Prefix
     17     P2MP root-addr 172.30.5.4, lsp-id 16777217

Output label database, 172.30.5.3:0--172.30.5.7:0
  Label     Prefix
 300272     P2MP root-addr 172.30.5.7, lsp-id 16777217

Input label database, 172.30.5.3:0--172.30.5.8:0

Output label database, 172.30.5.3:0--172.30.5.8:0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show ldp database p2mp
Input label database, 172.30.5.7:0--172.30.5.1:0

Output label database, 172.30.5.7:0--172.30.5.1:0

Input label database, 172.30.5.7:0--172.30.5.3:0
  Label      Prefix
 300272      P2MP root-addr 172.30.5.7, lsp-id 16777217

Output label database, 172.30.5.7:0--172.30.5.3:0
  Label      Prefix
     17      P2MP root-addr 172.30.5.4, lsp-id 16777217

Input label database, 172.30.5.7:0--172.30.5.6:0

Output label database, 172.30.5.7:0--172.30.5.6:0
```

8) Configure VPN Option C to the remote site S3 of customer CE2.

Customer CE2 has one more site attached to remote AS 43208.365 reachable via R8. Labeled BGP is required when Option C is used for interconnecting MPLS VPN. The Router Reflector is the best choice to establish the L3VPN EBGP session with remote AS.

a) Configure the route reflector for Inter-AS connectivity.

```
[edit]
lab@route-reflector# set protocols bgp group P2-remote-pe local-address 172.30.5.41

[edit]
lab@route-reflector# set protocols bgp group P2-remote-pe neighbor 172.17.47.3
```

Because the local route reflector router is using and old Junos version and the remote AS is using a 4-byte AS, you must configure the **peer-as** with transition AS 23456.

```
[edit]
lab@route-reflector# set protocols bgp group P2-remote-pe peer-as 23456
```

You must configure the session as multihop without changing the next hop.

```
[edit]
lab@route-reflector# set protocols bgp group P2-remote-pe multihop no-nexthop-change
```

Only the `inet-vpn` unicast family is enabled on the EBGP session with the remote PE.

```
[edit]
lab@route-reflector# set protocols bgp group P2-remote-pe family inet-vpn unicast
```


© 2020 Juniper Networks, Inc. All rights reserved.

The Labeled BGP family should be enabled on the IBGP session.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family inet labeled-unicast rib
inet.3

[edit]
lab@route-reflector# commit
commit complete
```

b) Configure R8 for Inter-AS support.

R8 is an ASBR between both ASs. Labeled BGP should be configure to the session with P2-1 and towards the route reflector.

Family mpls should be enabled on the interface to P2-1.

```
[edit]
lab@R8# set interfaces ge-0/0/5.302 family mpls

[edit]
lab@R8# set protocols bgp group ibgp family inet labeled-unicast rib inet.3

[edit]
lab@R8# set protocols bgp group P2-1 family inet labeled-unicast rib inet.3

[edit]
lab@R8# set protocols bgp group P2-1 family inet unicast

[edit]
lab@R8# commit
commit complete
```

c) Configure R7 for Inter-AS.

On R7 the IBGP group should also be enabled for labeled BGP.

```
[edit]
lab@R7# set protocols bgp group ibgp family inet labeled-unicast rib inet.3

[edit]
lab@R7# commit
commit complete
```

Verify BGP operation.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# run show bgp summary | no-more
Groups: 3 Peers: 9 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 505        505          0          0          0          0
inet.3                   1          1          0          0          0          0
bgp.l3vpn.0             36         36          0          0          0          0
inet6.0                 48         32          0          0          0          0
bgp.l2vpn.0              0          0          0          0          0          0
bgp.mvpn.0               2          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
172.17.47.3           23456         13         24       0       1        4:40 Establ
  bgp.l3vpn.0: 3/3/0
...
```

The EBGP session with the remote PE is established and three routes are received.

```
[edit]
lab@route-reflector# run show route receive-protocol bgp 172.17.47.3 table
bgp.l3vpn.0 extensive

bgp.l3vpn.0: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)
* 172.17.47.2:200:172.31.78.0/24 (1 entry, 1 announced)
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 299792
     Nexthop: 172.17.47.3
     AS path: 23456 64600 I Unrecognized Attributes: 13 bytes
     AS path:  Attr flags e0 code 11: 02 02 a8 c8 01 6d 00 00 fc 58
     Communities: target:43208:200

* 172.17.47.2:200:172.31.79.0/24 (1 entry, 1 announced)
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 299792
     Nexthop: 172.17.47.3
     AS path: 23456 64600 I Unrecognized Attributes: 13 bytes
     AS path:  Attr flags e0 code 11: 02 02 a8 c8 01 6d 00 00 fc 58
     Communities: target:43208:200

* 172.17.47.2:200:192.168.0.100/30 (1 entry, 1 announced)
     Route Distinguisher: 172.17.47.2:200
     VPN Label: 299792
     Nexthop: 172.17.47.3
     AS path: 23456 I Unrecognized Attributes: 9 bytes
     AS path:  Attr flags e0 code 11: 02 01 a8 c8 01 6d
     Communities: target:43208:200
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Routes received from the remote PE are attached with a route target that is not recognized by R7. You will have to normalize the route targets.

```
[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-export term 1
from protocol bgp

[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-export term 1
from community CE2-hub

[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-export term 1
then community delete CE2-hub

[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-export term 1
then community add CE2-remote

[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-export term 1
then accept

[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-import term 1
from protocol bgp

[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-import term 1
from community CE2-remote

[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-import term 1
then community delete CE2-remote

[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-import term 1
then community add CE2-spoke

[edit]
lab@route-reflector# set policy-options policy-statement CE2-vpn-target-import term 1
then accept

[edit]
lab@route-reflector# set policy-options community CE2-hub members target:54591:200

[edit]
lab@route-reflector# set policy-options community CE2-remote members target:43208:200
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# set policy-options community CE2-spoke members target:54591:201

[edit]
lab@route-reflector# set protocols bgp group P2-remote-pe import CE2-vpn-target-import

[edit]
lab@route-reflector# set protocols bgp group P2-remote-pe export CE2-vpn-target-export

[edit]
lab@route-reflector# commit
commit complete
```

After route target normalization is applied, the routes received from the remote AS are inserted in the CE2-spoke routing table.

```
[edit]
lab@R7# run show route table CE2-spoke.inet.0 | find 172.31.7[8-9]
172.31.78.0/24     *[BGP/170] 00:01:38, localpref 100, from 172.30.5.41
                      AS path: 2831679853 64600 I, validation-state: unverified
                    > to 172.30.0.5 via ge-0/0/4.117, label-switched-path R7-to-
R8
                    [BGP/170] 00:01:37, localpref 100
                      AS path: 64600 54591 2831679853 54591 I, validation-state:
unverified
                    > to 192.168.0.42 via ge-0/0/5.311
172.31.79.0/24     *[BGP/170] 00:01:38, localpref 100, from 172.30.5.41
                      AS path: 2831679853 64600 I, validation-state: unverified
                    > to 172.30.0.5 via ge-0/0/4.117, label-switched-path R7-to-
R8
                    [BGP/170] 00:01:37, localpref 100
                      AS path: 64600 54591 2831679853 54591 I, validation-state:
unverified
                    > to 192.168.0.42 via ge-0/0/5.311
...
```

The labeled BGP routes are resolved `inet.3`.

```
[edit]
lab@R7# run show route table inet.3 172.17.47.3/32 extensive

inet.3: 20 destinations, 31 routes (9 active, 0 holddown, 19 hidden)
172.17.47.3/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xd708970
                Next-hop reference count: 2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Source: 172.30.5.41
Next hop type: Router, Next hop index: 0
Next hop: 172.30.0.5 via ge-0/0/4.117 weight 0x1, selected
Label-switched-path R7-to-R8
Label operation: Push 305568, Push 309568(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 305568: None; Label 309568: None;
Label element ptr: 0xd708780
Label parent element ptr: 0xd9f6a00
Label element references: 3
Label element child references: 1
Label element lsp id: 0
Session Id: 0x0
```

9) Configure customer CE3 VPLS service as shown in Table 32.

Customer	Site	Router	VPLS signaling	CE facing interface
CE3	S1	CE3-1	BGP	ge-0/0/3.601
	S2	CE3-2	BGP	ge-0/0/3.600

*Table 32*41

Customer CE3 requires a VPLS service between R2 and R3. the site connected to R3 is multihomed to R4 also. VLAN normalization should be used because the customer interfaces are configured with different VLAN numbers on both ends.

a) Configure R2 for CE3 VPLS.

```
[edit]
lab@R2# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R2# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R2# set interfaces ge-0/0/3 unit 601 description "CE3-1 connection"

[edit]
lab@R2# set interfaces ge-0/0/3 unit 601 encapsulation vlan-vpls
```

Configure the route distinguisher.

```
[edit]
lab@R2# set routing-options route-distinguisher-id 172.30.5.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

The VLAN configured on the interface is 601.

```
[edit]
lab@R2# set interfaces ge-0/0/3 unit 601 vlan-id 601
```

Create a VPLS routing instance.

```
[edit]
lab@R2# set routing-instances CE3-vpls instance-type vpls

[edit]
lab@R2# set routing-instances CE3-vpls interface ge-0/0/3.601
```

VLAN normalization is configured to 600. This command will overwrite the VLAN tag from 601 to 600 in inbound direction and vice versa in the outbound direction.

```
[edit]
lab@R2# set routing-instances CE3-vpls vlan-id 600

[edit]
lab@R2# set routing-instances CE3-vpls vrf-target target:54591:300

[edit]
lab@R2# set routing-instances CE3-vpls protocols vpls site-range 8

[edit]
lab@R2# set routing-instances CE3-vpls protocols vpls no-tunnel-services

[edit]
lab@R2# set routing-instances CE3-vpls protocols vpls site site-1 site-identifier 1
```

Family l2vpn signaling is configured on the IBGP session for each PE that participates in the VPLS.

```
[edit]
lab@R2# set protocols bgp group ibgp family l2vpn signaling

[edit]
lab@R2# commit
commit complete
```

b) Configure R3 for CE3 VPLS.

```
[edit]
lab@R3# set interfaces ge-0/0/3 vlan-tagging
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R3# set interfaces ge-0/0/3.600 encapsulation vlan-vpls

[edit]
lab@R3# set interfaces ge-0/0/3.600 description "CE3-2 connection"

[edit]
lab@R3# set interfaces ge-0/0/3.600 vlan-id 600

[edit]
lab@R3# set routing-options route-distinguisher-id 172.30.5.3

[edit]
lab@R3# set routing-instances CE3-vpls instance-type vpls

[edit]
lab@R3# set routing-instances CE3-vpls vlan-id 600

[edit]
lab@R3# set routing-instances CE3-vpls interface ge-0/0/3.600

[edit]
lab@R3# set routing-instances CE3-vpls vrf-target target:54591:300

[edit]
lab@R3# set routing-instances CE3-vpls protocols vpls no-tunnel-services

[edit]
lab@R3# set routing-instances CE3-vpls protocols vpls site site-2 site-identifier 2
```

Site is configured for multihoming.

```
[edit]
lab@R3# set routing-instances CE3-vpls protocols vpls site site-2 multi-homing
```

This site is configured as primary. This will automatically increase the local preference of the BGP announcements for this site to the maximum value of 65535.

```
[edit]
lab@R3# set routing-instances CE3-vpls protocols vpls site site-2 site-preference
primary

[edit]
lab@R3# set protocols bgp group ibgp family l2vpn signaling
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# commit
commit complete
```

c) Configure R3 for CE3 VPLS.

```
[edit]
lab@R4# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R4# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R4# set interfaces ge-0/0/3.600 description "CE3-2 connection"

[edit]
lab@R4# set interfaces ge-0/0/3.600 encapsulation vlan-vpls

[edit]
lab@R4# set interfaces ge-0/0/3.600 vlan-id 600

[edit]
lab@R4# set routing-instances CE3-vpls instance-type vpls

[edit]
lab@R4# set routing-instances CE3-vpls vlan-id 600

[edit]
lab@R4# set routing-instances CE3-vpls interface ge-0/0/3.600

[edit]
lab@R4# set routing-instances CE3-vpls vrf-target target:54591:300

[edit]
lab@R4# set routing-instances CE3-vpls protocols vpls

[edit]
lab@R4# set routing-instances CE3-vpls protocols vpls site-range 8

[edit]
lab@R4# set routing-instances CE3-vpls protocols vpls no-tunnel-services
```

The same site-identifier is configured as on R3 due to multi-homing.

```
[edit]
lab@R4# set routing-instances CE3-vpls protocols vpls site site-2 site-identifier 2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set routing-instances CE3-vpls protocols vpls site site-2 multi-homing
```

Router R4 is configured as the VPLS backup. BGP routes for this site are announced with a local preference value of 1.

```
[edit]
lab@R4# set routing-instances CE3-vpls protocols vpls site site-2 site-preference
backup

[edit]
lab@R4# set protocols bgp group ibgp family l2vpn signaling

[edit]
lab@R4# commit
commit complete
```

d) Configure route reflector to support VPLS.

Enable the Layer 2 VPN NLRI on the route reflector.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family l2vpn signaling

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family l2vpn signaling

[edit]
lab@route-reflector# commit
commit complete
```

Verify the operation of the VPLS configuration of CE3 customer.

```
[edit]
lab@R2# run show vpls connections
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
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
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           HS -- Hot-standby Connection

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

Instance: CE3-vpls
Edge protection: Not-Primary
  Local site: site-1 (1)
    connection-site           Type  St     Time last up          # Up trans
    2                         rmt   Up     Apr 10 07:15:46 2019           1
      Remote PE: 172.30.5.3, Negotiated control-word: No
      Incoming label: 262146, Outgoing label: 262145
      Local interface: lsi.1048576, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CE3-vpls local site 1 remote site 2
      Flow Label Transmit: No, Flow Label Receive: No
```

Connection to site 2 on primary PE, R3, is in an Up state.

```
[edit]
lab@R2# run show route forwarding-table family vpls
VPLS:
Enabled protocols: Single VLAN, ACKed by all peers,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      681     1
lsi.1048576        intf     0                    indr  1048602     5
                                                 ulst  1048601     2
                              172.30.0.18       Push 262145, Push 302768(top)
706     2 ge-0/0/4.126
                              172.30.0.14       Push 262145, Push 302768, Push
300688(top)      707     2 ge-0/0/4.123
00:50:56:a9:17:ac/48 user     0                  ucst      689     4 ge-0/0/3.601
...
```

R3 is the primary site for multi-homed customer site.

```
[edit]
lab@R3# run show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
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           HS -- Hot-standby Connection

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

Instance: CE3-vpls
Edge protection: Not-Primary
  Local site: site-2 (2)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Apr 10 07:15:53 2019           1
      Remote PE: 172.30.5.2, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262146
      Local interface: lsi.1048576, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls CE3-vpls local site 2 remote site 1
      Flow Label Transmit: No, Flow Label Receive: No
    2                         rmt   RN
```

The L2VPN BGP route is advertised with the maximum possible local preference value.

```
[edit]
lab@R3# run show route advertising-protocol  bgp 172.30.5.41 table CE3-
vpls.l2vpn.0 extensive

CE3-vpls.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
*  172.30.5.3:8:2:1/96 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.30.5.3:8
     Label-base: 262145, range: 8, offset: 1
     Nexthop: Self
     Flags: Nexthop Change
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
     Localpref: 65535
     AS path: [54591] I
     Communities: target:54591:300 Layer2-info: encaps: VPLS, control flags:[0x0]
, mtu: 0, site preference: 65535
```

R4 is a backup site and advertises the L2VPN BGP routes with the lowest possible local preference value.

```
[edit]
lab@R4# run show route advertising-protocol bgp 172.30.5.41 table CE3-
vpls.l2vpn.0 extensive

CE3-vpls.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
*  172.30.5.4:9:2:1/96 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 172.30.5.4:9
     Label-base: 262145, range: 8, offset: 1
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 1
     AS path: [54591] I
     Communities: target:54591:300 Layer2-info: encaps: VPLS, control flags:[0x0]
, mtu: 0, site preference: 1
```

Configure the VPLS that if 100 MAC addresses are learned that packets are dropped.

R2

```
[edit]
lab@R2# set routing-instances CE3-vpls protocols vpls interface-mac-limit 100

[edit]
lab@R2# set routing-instances CE3-vpls protocols vpls mac-table-size 100

[edit]
lab@R2# set routing-instances CE3-vpls protocols vpls mac-table-size packet-action drop

[edit]
lab@R2# commit
commit complete
```

R3

```
[edit]
lab@R3# set routing-instances CE3-vpls protocols vpls interface-mac-limit 100

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set routing-instances CE3-vpls protocols vpls mac-table-size 100

[edit]
lab@R3# set routing-instances CE3-vpls protocols vpls mac-table-size packet-action drop

[edit]
lab@R3# commit
commit complete
```

R4

```
[edit]
lab@R4# set routing-instances CE3-vpls protocols vpls interface-mac-limit 100

[edit]
lab@R4# set routing-instances CE3-vpls protocols vpls mac-table-size 100

[edit]
lab@R4# set routing-instances CE3-vpls protocols vpls mac-table-size packet-action drop

[edit]
lab@R4# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

1) Configure an EVPN service to enable communication between CE12-1 and CE34-1. Enable active-active multihoming for unicast traffic. Use the information as provided in table below.

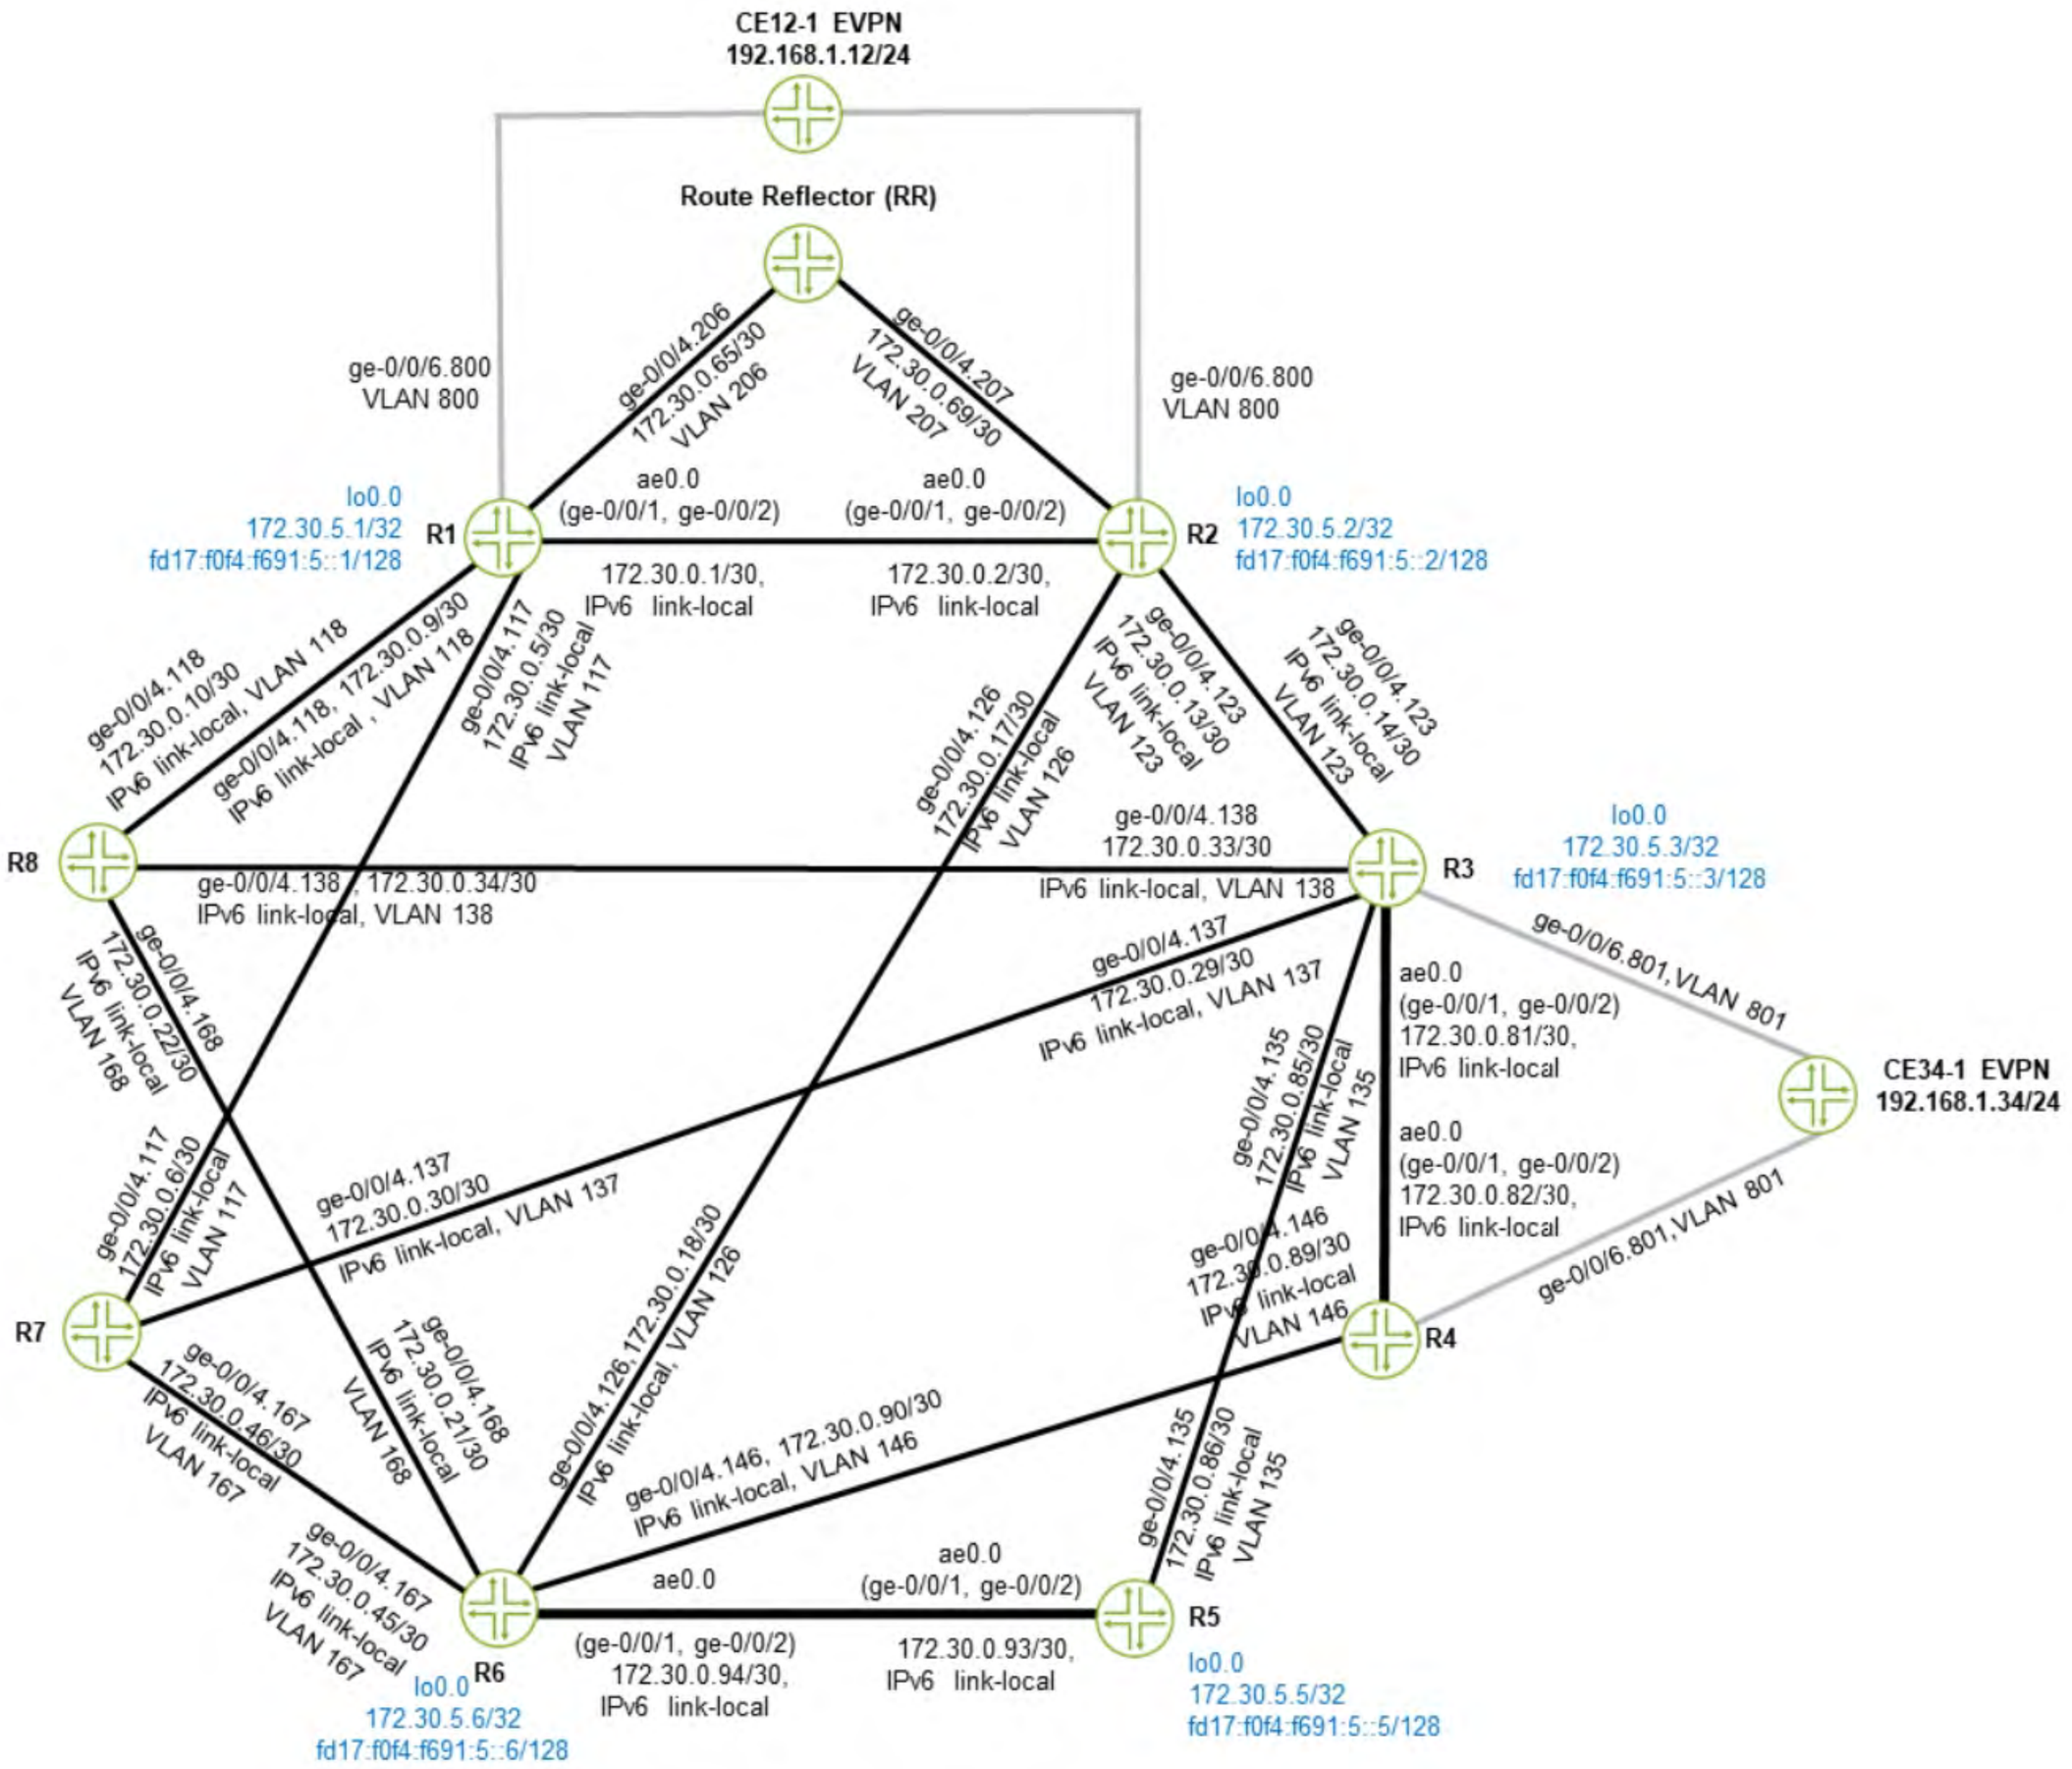


Figure 1 EVPN topology

© 2020 Juniper Networks, Inc. All rights reserved.

CE Router	CE L3 address	PE router	CE facing interface	LACP	Vlan ID	Default gateway
CE12-1	192.168.1.12/24	R1	ae1.800	Enabled	800	192.168.1.254/24
		R2				
CE34-1	192.168.1.34/24	R3	ae1.801		801	
		R4				

Note: the CE devices are preconfigured on VR-DEVICE2.

First, we need to enable EVPN signaling by exchanging EVPN NLRI in BGP for the required routers

RR

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family evpn signaling

[edit]
lab@route-reflector# set protocols bgp group cluster-2 family evpn signaling

[edit]
lab@route-reflector# commit
commit complete
```

R1

```
[edit]
lab@R1# set protocols bgp group ibgp family evpn signaling

[edit]
lab@R1# commit
commit complete
```

R2

```
[edit]
lab@R2# set protocols bgp group ibgp family evpn signaling

[edit]
lab@R2# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

R3

```
[edit]
lab@R3# set protocols bgp group ibgp family evpn signaling

[edit]
lab@R3# commit
commit complete
```

R4

```
[edit]
lab@R4# set protocols bgp group ibgp family evpn signaling

[edit]
lab@R4# commit
commit complete
```

Verify the exchange of EVPN NLRI in BGP.

```
lab@R1# run show bgp summary
Groups: 5 Peers: 6 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
…
…
…
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.41           54591        598        578       0       1     1:06:13 Establ
  bgp.rtarget.0: 11/17/17/0
  inet.0: 216/216/216/0
  inet6.0: 32/32/32/0
  bgp.l3vpn.0: 16/16/16/0
  bgp.mvpn.0: 4/4/4/0
  C2-spoke.inet.0: 12/16/16/0
  C2-spoke.mvpn.0: 4/4/4/0
  bgp.evpn.0: 0/0/0/0
```

At this stage it is expected that no EVPN routes are exchanged as no EVPN service is configured yet.

Configure the aggregate (member) interfaces and the IRB interface on R1, R2, R3 and R4. To support active-active multihoming in EVPN we need to make sure that the “all-active” knob is configured. Also make sure the LACP system id match and the ESI values match on each ethernet segment. The ESI value is signaled in BGP to ensure BUM traffic is not forwarded to ethernet segments that have the same ESI value configured. For this reason, the ESI values should match on each multihomed ethernet segment.

© 2020 Juniper Networks, Inc. All rights reserved.

R1, R2:

```
set chassis aggregated-devices ethernet device-count 2

set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 esi 00:00:00:00:00:00:00:00:00:12
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:12
set interfaces ae1 unit 800 encapsulation vlan-bridge
set interfaces ae1 unit 800 vlan-id 800
set interfaces irb unit 1 family inet address 192.168.2.254/24

set interfaces ge-0/0/6 gigether-options 802.3ad ae1

set routing-instances evpn1 instance-type evpn
set routing-instances evpn1 vlan-id 800
set routing-instances evpn1 interface ae1.800
set routing-instances evpn1 routing-interface irb.1
set routing-instances evpn1 vrf-target target:54591:1000
set routing-instances evpn1 protocols evpn interface ae1.800
```

R1:

```
set routing-instances evpn1 route-distinguisher 172.30.5.1:1000
commit
```

R2:

```
set routing-instances evpn1 route-distinguisher 172.30.5.2:1000
commit
```

R3, R4:

```
set chassis aggregated-devices ethernet device-count 2
set interfaces ae1 flexible-vlan-tagging
set interfaces ae1 encapsulation flexible-ethernet-services
set interfaces ae1 esi 00:00:00:00:00:00:00:00:00:34
set interfaces ae1 esi all-active
set interfaces ae1 aggregated-ether-options lacp active
set interfaces ae1 aggregated-ether-options lacp system-id 00:00:00:00:00:34
set interfaces ae1 unit 801 encapsulation vlan-bridge
set interfaces ae1 unit 801 vlan-id 801
set interfaces irb unit 1 family inet address 192.168.2.254/24

set interfaces ge-0/0/6 gigether-options 802.3ad ae1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
set routing-instances evpn1 instance-type evpn
set routing-instances evpn1 vlan-id 800
set routing-instances evpn1 interface ae1.801
set routing-instances evpn1 routing-interface irb.1
set routing-instances evpn1 vrf-target target:54591:1000
set routing-instances evpn1 protocols evpn interface ae1.801
```

R3:

```
set routing-instances evpn1 route-distinguisher 172.30.5.3:1000
commit
```

R4:

```
set routing-instances evpn1 route-distinguisher 172.30.5.4:1000
commit
```

Verification on R1 (the other routers should have similar outputs:

Interface ae1 should be up and LACP PDUs should be send and received.

```
lab@R1# run show interfaces terse | match ae1
ge-0/0/6.800               up    up   aenet    --> ae1.800
ge-0/0/6.32767          up    up   aenet    --> ae1.32767
ae1                                 up    up
ae1.800                        up    up   bridge
ae1.32767                   up    up   multiservice

lab@R1# run show lacp statistics interfaces ae1
Aggregated interface: ae1
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-0/0/6                1292        1275            0            0
```

Verify if EVPN NLRI is exchanged

```
lab@R1# run show bgp summary group ibgp | match evpn
bgp.evpn.0
  bgp.evpn.0: 21/21/21/0
  evpn1.evpn.0: 18/18/18/0
  __default_evpn__.evpn.0: 3/3/3/0
```

R2, R3 and R4 should have similar outputs.

© 2020 Juniper Networks, Inc. All rights reserved.

Verify connectivity between CE12-1 and CE34-4. Also verify reachability towards the default gateway.

VR-DEVICE2:

```
root@vr-device-2> show route table CE12-1.inet.0

CE12-1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:03:52
                    > to 192.168.2.254 via ae1.0
192.168.2.0/24     *[Direct/0] 00:03:52
                    > via ae1.0
192.168.2.12/32    *[Local/0] 00:03:52
                      Local via ae1.0

root@vr-device-2> show route table CE34-1.inet.0

CE34-1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:04:22
                    > to 192.168.2.254 via ae2.0
192.168.2.0/24     *[Direct/0] 00:04:22
                    > via ae2.0
192.168.2.34/32    *[Local/0] 00:04:22
                      Local via ae2.0
```

Verify the EVPN database:

```
lab@R1# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                  Timestamp        IP
address
800             2c:6b:f5:10:01:f0  172.30.5.3                     Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:1a:b8:f0  172.30.5.2                     Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:38:46:c0  00:00:00:00:00:00:00:00:00:12  Mar 05 14:44:28
192.168.2.12
800             2c:6b:f5:38:46:c1  00:00:00:00:00:00:00:00:00:34  Mar 05 14:44:27
192.168.2.34
800             2c:6b:f5:54:61:f0  172.30.5.4                     Mar 05 14:44:32
192.168.2.254
800             2c:6b:f5:86:c6:f0  irb.1                          Mar 05 14:44:24
192.168.2.254
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                      Timestamp        IP
address
800             2c:6b:f5:10:01:f0  172.30.5.3                         Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:1a:b8:f0  irb.1                              Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:38:46:c0  00:00:00:00:00:00:00:00:00:12      Mar 05 14:44:28
192.168.2.12
800             2c:6b:f5:38:46:c1  00:00:00:00:00:00:00:00:00:34      Mar 05 14:44:27
192.168.2.34
800             2c:6b:f5:54:61:f0  172.30.5.4                         Mar 05 14:44:32
192.168.2.254
800             2c:6b:f5:86:c6:f0  172.30.5.1                         Mar 05 14:44:25
192.168.2.254

lab@R3# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                      Timestamp        IP
address
800             2c:6b:f5:10:01:f0  irb.1                              Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:1a:b8:f0  172.30.5.2                         Mar 05 14:44:24
192.168.2.254
800             2c:6b:f5:38:46:c0  00:00:00:00:00:00:00:00:00:12      Mar 05 14:44:28
192.168.2.12
800             2c:6b:f5:38:46:c1  00:00:00:00:00:00:00:00:00:34      Mar 05 14:44:26
192.168.2.34
800             2c:6b:f5:54:61:f0  172.30.5.4                         Mar 05 14:44:31
192.168.2.254
800             2c:6b:f5:86:c6:f0  172.30.5.1                         Mar 05 14:44:24
192.168.2.254

lab@R4# run show evpn database
Instance: evpn1
VLAN  DomainId  MAC address        Active source                      Timestamp        IP
address
800             2c:6b:f5:10:01:f0  172.30.5.3                         Mar 05 14:44:33
192.168.2.254
800             2c:6b:f5:1a:b8:f0  172.30.5.2                         Mar 05 14:44:32
192.168.2.254
800             2c:6b:f5:38:46:c1  00:00:00:00:00:00:00:00:00:34      Mar 05 14:44:33
192.168.2.34
800             2c:6b:f5:54:61:f0  irb.1                              Mar 05 14:44:23
192.168.2.254
```

Ping CE34-1 to verify end to end connectivity

© 2020 Juniper Networks, Inc. All rights reserved.

```
root@vr-device-2> ping 192.168.2.34 routing-instance CE12-1
PING 192.168.1.34 (192.168.1.34): 56 data bytes
64 bytes from 192.168.1.34: icmp_seq=0 ttl=64 time=5.830 ms
64 bytes from 192.168.1.34: icmp_seq=1 ttl=64 time=5.424 ms
```

Verify reachability towards the default gateway

```
root@vr-device-2> ping 192.168.2.254 routing-instance CE12-1
PING 192.168.1.254 (192.168.1.254): 56 data bytes
64 bytes from 192.168.1.254: icmp_seq=0 ttl=64 time=4.837 ms
64 bytes from 192.168.1.254: icmp_seq=1 ttl=64 time=4.380 ms
```

Task 7: Class of Service Configuration

In this task you have to prepare the network for Class of Service support. You will configure classifiers, policers, forwarding classes, schedulers and rewrite rules.

To provide end-to-end Class of Service in the network you have to ensure a consistent configuration on all devices. You can configure the common configuration on one router and then copy it to the other routers.

Repeat this on all eight devices.

1) Configure code-point aliases. Based on information in Table 33.

Forwarding Class	Loss Priority	DSCP Value	EXP Value
be	low	0b000000	0b000
l3vpn	low	0b001000	0b001
l2vpn	low	0b001010	0b010
	high	0b001011	0b011
l3vpn-priority	low	0b101110	0b101
nc	low	0b110000	

Table 33

```
[edit]
lab@R4# set class-of-service code-point-aliases dscp be 000000

[edit]
lab@R4# set class-of-service code-point-aliases dscp l3vpn 001000

[edit]
lab@R4# set class-of-service code-point-aliases dscp l2vpn-low 001010
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set class-of-service code-point-aliases dscp l2vpn-high 001011

[edit]
lab@R4# set class-of-service code-point-aliases dscp l3vpn-priority 101110

[edit]
lab@R4# set class-of-service code-point-aliases dscp nc 110000

[edit]
lab@R4# set class-of-service code-point-aliases exp be 000

[edit]
lab@R4# set class-of-service code-point-aliases exp l3vpn 001

[edit]
lab@R4# set class-of-service code-point-aliases exp l2vpn-low 010

[edit]
lab@R4# set class-of-service code-point-aliases exp l2vpn-high 011

[edit]
lab@R4# set class-of-service code-point-aliases exp l3vpn-priority 101
```

2) Configure traffic classifiers according to Table 34.

Forwarding Class	Loss Priority	DSCP Value	EXP Value
be	low	0b000000	0b000
l3vpn	low	0b001000	0b001
l2vpn	low	0b001010	0b010
	high	0b001011	0b011
l3vpn-priority	low	0b101110	0b101
nc	low	0b110000	

Table 34

```
[edit]
lab@R4# set class-of-service classifiers dscp dscp-classifier forwarding-class be loss-
priority low code-points be

[edit]
lab@R4# set class-of-service classifiers dscp dscp-classifier forwarding-class l3vpn
loss-priority low code-points l3vpn

[edit]
lab@R4# set class-of-service classifiers dscp dscp-classifier forwarding-class l2vpn
loss-priority low code-points l2vpn-low
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set class-of-service classifiers dscp dscp-classifier forwarding-class l2vpn
loss-priority high code-points l2vpn-high

[edit]
lab@R4# set class-of-service classifiers dscp dscp-classifier forwarding-class l3vpn-
priority loss-priority low code-points l3vpn-priority

[edit]
lab@R4# set class-of-service classifiers dscp dscp-classifier forwarding-class nc loss-
priority low code-points nc

[edit]
lab@R4# set class-of-service classifiers exp mpls-classifier forwarding-class be loss-
priority low code-points be

[edit]
lab@R4# set class-of-service classifiers exp mpls-classifier forwarding-class l3vpn
loss-priority low code-points l3vpn

[edit]
lab@R4# set class-of-service classifiers exp mpls-classifier forwarding-class l2vpn
loss-priority low code-points l2vpn-low

[edit]
lab@R4# set class-of-service classifiers exp mpls-classifier forwarding-class l2vpn
loss-priority high code-points l2vpn-high

[edit]
lab@R4# set class-of-service classifiers exp mpls-classifier forwarding-class l3vpn-
priority loss-priority low code-points l3vpn-priority
```

3) Define the forwarding-classes to the queues according to Table 35.

Forwarding Class	Queue	Scheduler
be	0	be-sc
l3vpn	1	l3vpn-sc
l2vpn	2	l3vpn-pri-sc
l3vpn-priority	3	l2vpn-sc
nc	4	nc-sc

Table 35

```
[edit]
lab@R4# set class-of-service forwarding-classes queue 0 be
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set class-of-service forwarding-classes queue 1 l3vpn

[edit]
lab@R4# set class-of-service forwarding-classes queue 2 l2vpn

[edit]
lab@R4# set class-of-service forwarding-classes queue 3 l3vpn-priority

[edit]
lab@R4# set class-of-service forwarding-classes queue 4 nc
```

4) Configure drop-profiles according to Table 36.

Drop Profile	Fill Level	Drop Probability
low-drop	25	5
	50	15
	75	40
high-drop	25	10
	50	30
	75	65

Table 36

```
[edit]
lab@R4# set class-of-service drop-profiles low-drop interpolate fill-level 25

[edit]
lab@R4# set class-of-service drop-profiles low-drop interpolate fill-level 50

[edit]
lab@R4# set class-of-service drop-profiles low-drop interpolate fill-level 75

[edit]
lab@R4# set class-of-service drop-profiles low-drop interpolate drop-probability 5

[edit]
lab@R4# set class-of-service drop-profiles low-drop interpolate drop-probability 15

[edit]
lab@R4# set class-of-service drop-profiles low-drop interpolate drop-probability 40

[edit]
lab@R4# set class-of-service drop-profiles high-drop interpolate fill-level 25

[edit]
lab@R4# set class-of-service drop-profiles high-drop interpolate fill-level 50
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set class-of-service drop-profiles high-drop interpolate fill-level 75

[edit]
lab@R4# set class-of-service drop-profiles high-drop interpolate drop-probability 10

[edit]
lab@R4# set class-of-service drop-profiles high-drop interpolate drop-probability 30

[edit]
lab@R4# set class-of-service drop-profiles high-drop interpolate drop-probability 65
```

5) Configure the rewrite rules according to Table 37.

Forwarding Class	Loss Priority	DSCP Value	EXP Value
be	low	0b000000	0b000
l3vpn	low	0b001000	0b001
l2vpn	low	0b001010	0b010
	high	0b001011	0b011
l3vpn-priority	low	0b101110	0b101
nc	low	0b110000	

Table 37

```
[edit]
lab@R4# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class be loss-
priority low code-point be

[edit]
lab@R4# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class l3vpn
loss-priority low code-point l3vpn

[edit]
lab@R4# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class l2vpn
loss-priority low code-point l2vpn-low

[edit]
lab@R4# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class l2vpn
loss-priority high code-point l2vpn-high

[edit]
lab@R4# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class l3vpn-
priority loss-priority low code-point l3vpn-priority

[edit]
lab@R4# set class-of-service rewrite-rules dscp dscp-rewriter forwarding-class nc loss-
priority low code-point nc
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set class-of-service rewrite-rules exp mpls-rewriter forwarding-class be loss-
priority low code-point be

[edit]
lab@R4# set class-of-service rewrite-rules exp mpls-rewriter forwarding-class l3vpn
loss-priority low code-point l3vpn

[edit]
lab@R4# set class-of-service rewrite-rules exp mpls-rewriter forwarding-class l2vpn
loss-priority low code-point l2vpn-low

[edit]
lab@R4# set class-of-service rewrite-rules exp mpls-rewriter forwarding-class l2vpn
loss-priority high code-point l2vpn-high

[edit]
lab@R4# set class-of-service rewrite-rules exp mpls-rewriter forwarding-class l3vpn-
priority loss-priority low code-point l3vpn-priority
```

6) Configure the schedulers according to Table 38.

Scheduler	Parameter	Value
be-sc	Priority	low
	Transmit rate	remainder
	Buffer size	remainder
	Drop profile LP any	high-drop
l3vpn-sc	Priority	medium-low
	Transmit rate	20%
	Buffer size	20%
l2vpn-sc	Priority	medium-high
	Transmit rate	20%
	Buffer size	20%
	Drop profile LP low	low-drop
	Drop profile LP high	high-drop
l3vpn-pri-sc	Priority	high
	Transmit rate	10%
	Buffer size	5 msec
nc-sc	Priority	high
	Transmit rate	5%
	Buffer size	5%

Table 38

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set class-of-service schedulers be-sc transmit-rate remainder

[edit]
lab@R4# set class-of-service schedulers be-sc buffer-size remainder

[edit]
lab@R4# set class-of-service schedulers be-sc priority low

[edit]
lab@R4# set class-of-service schedulers be-sc drop-profile-map loss-priority any
protocol any drop-profile high-drop

[edit]
lab@R4# set class-of-service schedulers l3vpn-sc transmit-rate percent 20

[edit]
lab@R4# set class-of-service schedulers l3vpn-sc buffer-size percent 20

[edit]
lab@R4# set class-of-service schedulers l3vpn-sc priority medium-low

[edit]
lab@R4# set class-of-service schedulers l2vpn-sc transmit-rate percent 20

[edit]
lab@R4# set class-of-service schedulers l2vpn-sc buffer-size percent 20

[edit]
lab@R4# set class-of-service schedulers l2vpn-sc priority medium-high

[edit]
lab@R4# set class-of-service schedulers l2vpn-sc drop-profile-map loss-priority low
protocol any drop-profile low-drop

[edit]
lab@R4# set class-of-service schedulers l2vpn-sc drop-profile-map loss-priority high
protocol any drop-profile high-drop

[edit]
lab@R4# set class-of-service schedulers l3vpn-pri-sc transmit-rate percent 10

[edit]
lab@R4# set class-of-service schedulers l3vpn-pri-sc buffer-size temporal 5k

[edit]
lab@R4# set class-of-service schedulers l3vpn-pri-sc priority high
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set class-of-service schedulers nc-sc transmit-rate percent 5

[edit]
lab@R4# set class-of-service schedulers nc-sc buffer-size percent 5

[edit]
lab@R4# set class-of-service schedulers nc-sc priority high
```

7) Configure scheduler-map to attach the schedulers to the forwarding-classes according to Table 39.

Forwarding Class	Queue	Scheduler
be	0	be-sc
l3vpn	1	l3vpn-sc
l2vpn	2	l3vpn-pri-sc
l3vpn-priority	3	l2vpn-sc
nc	4	nc-sc

Table 39

```
[edit]
lab@R4# set class-of-service scheduler-maps core-interfaces forwarding-class be
scheduler be-sc

[edit]
lab@R4# set class-of-service scheduler-maps core-interfaces forwarding-class l3vpn
scheduler l3vpn-sc

[edit]
lab@R4# set class-of-service scheduler-maps core-interfaces forwarding-class l2vpn
scheduler l2vpn-sc

[edit]
lab@R4# set class-of-service scheduler-maps core-interfaces forwarding-class l3vpn-
priority scheduler l3vpn-pri-
sc

[edit]
lab@R4# set class-of-service scheduler-maps core-interfaces forwarding-class nc
scheduler nc-sc
```

8) Configure the interfaces with traffic classifiers and rewrite rules.

This step is specific for each router based on the interface's names. All core facing interfaces should be configured in the way shown as previously shown on R4.

© 2020 Juniper Networks, Inc. All rights reserved.

Note: Don't forget to configure the interface to P2-1 on R8 as core interface. R8 is ASBR between both ASs providing Option C Inter-AS and labeled packets are sent and received from remote AS.

Repeat this on R1, R2, R3, R4, R5, and R6.

```
[edit]
lab@R4# set class-of-service interfaces ge-0/0/4 scheduler-map core-interfaces

[edit]
lab@R4# set class-of-service interfaces ge-0/0/4 unit * classifiers dscp dscp-
classifier

[edit]
lab@R4# set class-of-service interfaces ge-0/0/4 unit * classifiers exp mpls-classifier

[edit]
lab@R4# set class-of-service interfaces ge-0/0/4 unit * rewrite-rules dscp dscp-
rewriter

[edit]
lab@R4# set class-of-service interfaces ge-0/0/4 unit * rewrite-rules exp mpls-rewriter
protocol mpls-inet-both
```

MPLS-rewriter will overwrite both EXP and IP precedence values.

```
[edit]
lab@R4# set class-of-service interfaces ae0 scheduler-map core-interfaces

[edit]
lab@R4# set class-of-service interfaces ae0 unit * classifiers dscp dscp-classifier

[edit]
lab@R4# set class-of-service interfaces ae0 unit * classifiers exp mpls-classifier

[edit]
lab@R4# set class-of-service interfaces ae0 unit * rewrite-rules dscp dscp-rewriter

[edit]
lab@R4# set class-of-service interfaces ae0 unit * rewrite-rules exp mpls-rewriter
protocol mpls-inet-both

[edit]
lab@R4# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Repeat this on R7 and R8.

```
[edit]
lab@R7# set class-of-service interfaces ge-0/0/4 scheduler-map core-interfaces

[edit]
lab@R7# set class-of-service interfaces ge-0/0/4 unit * classifiers dscp dscp-
classifier

[edit]
lab@R7# set class-of-service interfaces ge-0/0/4 unit * classifiers exp mpls-classifier

[edit]
lab@R7# set class-of-service interfaces ge-0/0/4 unit * rewrite-rules dscp dscp-
rewriter

[edit]
lab@R7# set class-of-service interfaces ge-0/0/4 unit * rewrite-rules exp mpls-rewriter
protocol mpls-inet-both

[edit]
lab@R7# commit
commit complete
```

9) Configure a Multi-field classifier for the L3VPN customer interfaces according to Table 40.

Traffic Type	Criteria	Forwarding Class
VPN regular	DSCP 0b000000	l3vpn
VPN priority	DSCP 0b101110	l3vpn-priority

Table 40

Repeat this on the R1, R4, and R7 devices.

```
[edit]
lab@R4# set firewall family inet filter l3vpn-classifier term 1 from dscp be

[edit]
lab@R4# set firewall family inet filter l3vpn-classifier term 1 then forwarding-class
l3vpn

[edit]
lab@R4# set firewall family inet filter l3vpn-classifier term 1 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set firewall family inet filter l3vpn-classifier term 2 from dscp ef

[edit]
lab@R4# set firewall family inet filter l3vpn-classifier term 2 then policer l3vpn-
priority-policer

[edit]
lab@R4# set firewall family inet filter l3vpn-classifier term 2 then forwarding-class
l3vpn-priority

[edit]
lab@R4# set firewall family inet filter l3vpn-classifier term 2 then accept

[edit]
lab@R4# set firewall family inet filter l3vpn-classifier term 3 then accept
```

Configure the policer for priority L3VPN traffic.

```
[edit]
lab@R4# set firewall policer l3vpn-priority-policer if-exceeding bandwidth-limit 25m

[edit]
lab@R4# set firewall policer l3vpn-priority-policer if-exceeding burst-size-limit 15k

[edit]
lab@R4# set firewall policer l3vpn-priority-policer then discard
```

Apply the Multifield classifier on the proper interfaces.

R1

```
[edit]
lab@R1# set interfaces ge-0/0/5.318 family inet filter input l3vpn-classifier

[edit]
lab@R1# commit
commit complete
```

R4

```
[edit]
lab@R4# set interfaces ge-0/0/5.323 family inet filter input l3vpn-classifier


[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# commit
commit complete
```

R7

```
[edit]
lab@R7# set interfaces ge-0/0/5.311 family inet filter input l3vpn-classifier

[edit]
lab@R7# set interfaces ge-0/0/5.312 family inet filter input l3vpn-classifier

[edit]
lab@R7# set interfaces ge-0/0/5.324 family inet filter input l3vpn-classifier

[edit]
lab@R7# commit
commit complete
```

10) Configure the multi-field classifier for the VPLS customer interfaces.

The multi-field classifier for VPLS traffic should be configured under **firewall family vpls** as it will be applied to **family vpls** under the **interface** hierarchy.

The VPLS filter will police all the traffic that enters the interface.

Repeat this on the R2, R3, and R4 devices.

```
[edit]
lab@R4# set firewall family vpls filter l2vpn-classifier term 1 then policer l2vpn-
policer

[edit]
lab@R4# set firewall family vpls filter l2vpn-classifier term 1 then forwarding-class
l2vpn

[edit]
lab@R4# set firewall policer l2vpn-policer if-exceeding bandwidth-limit 50m

[edit]
lab@R4# set firewall policer l2vpn-policer if-exceeding burst-size-limit 62k

[edit]
lab@R4# set firewall policer l2vpn-policer then loss-priority high
```

Apply the Multifield classifier on the interfaces.

R2

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set interfaces ge-0/0/3.601 family vpls filter input l2vpn-classifier

[edit]
lab@R2# commit
commit complete
```

R3

```
[edit]
lab@R3# set interfaces ge-0/0/3.600 family vpls filter input l2vpn-classifier

[edit]
lab@R3# commit
 commit complete
```

R4

```
[edit]
lab@R4# set interfaces ge-0/0/3.600 family vpls filter input l2vpn-classifier

[edit]
lab@R4# commit
commit complete
```

Verify the CoS configuration on R4.

```
[edit]
lab@R4# run show class-of-service classifier name dscp-classifier
Classifier: dscp-classifier, Code point type: dscp, Index: 51090
  Code point         Forwarding class                    Loss priority
  000000             be                                  low
  001000             l3vpn                               low
  001010             l2vpn                               low
  001011             l2vpn                               high
  101110             l3vpn-priority                      low
  110000             nc                                  low
```

```
[edit]
lab@R4# run show class-of-service classifier name mpls-classifier
Classifier: mpls-classifier, Code point type: exp, Index: 48975
  Code point         Forwarding class                    Loss priority
  000                be                                  low
  001                l3vpn                               low
  010                l2vpn                               low
  011                l2vpn                               high
  101                l3vpn-priority                      low
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show class-of-service rewrite-rule | find dscp-rewriter
Rewrite rule: dscp-rewriter, Code point type: dscp, Index: 20901
  Forwarding class                Loss priority       Code point
  be                              low                 000000
  l3vpn                           low                 001000
  l2vpn                           low                 001010
  l2vpn                           high                001011
  l3vpn-priority                  low                 101110
  nc                              low                 110000

Rewrite rule: mpls-rewriter, Code point type: exp, Index: 10617
  Forwarding class                Loss priority       Code point
  be                              low                 000
  l3vpn                           low                 001
  l2vpn                           low                 010
  l2vpn                           high                011
  l3vpn-priority                  low                 101
```

```
[edit]
lab@R4# run show class-of-service forwarding-class
Forwarding class                        ID      Queue  Restricted queue  Fabric
priority  Policing priority   SPU priority
  be                                    0       0          0              low
normal            low
  l3vpn                                 1       1          1              high
normal            low
  l2vpn                                 2       2          2              high
normal            low
  l3vpn-priority                        3       3          3              high
normal            low
  nc                                    4       4          0              high
normal            low
```

```
[edit]
lab@R4# run show class-of-service scheduler-map core-interfaces
Scheduler map: core-interfaces, Index: 58651

  Scheduler: be-sc, Forwarding class: be, Index: 15377
    Transmit rate: remainder, Rate Limit: none, Buffer size: remainder,
    Buffer Limit: none, Priority: low
    Excess Priority: unspecified
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             any         48162    high-drop
      Medium low      any         48162    high-drop
      Medium high     any         48162    high-drop
      High            any         48162    high-drop
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Scheduler: l3vpn-sc, Forwarding class: l3vpn, Index: 37464
  Transmit rate: 20 percent, Rate Limit: none, Buffer size: 20 percent,
  Buffer Limit: none, Priority: medium-low
  Excess Priority: unspecified
  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             any             1    <default-drop-profile>
    Medium low      any             1    <default-drop-profile>
    Medium high     any             1    <default-drop-profile>
    High            any             1    <default-drop-profile>

Scheduler: l2vpn-sc, Forwarding class: l2vpn, Index: 53848
  Transmit rate: 20 percent, Rate Limit: none, Buffer size: 20 percent,
  Buffer Limit: none, Priority: medium-high
  Excess Priority: unspecified
  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             any         59912    low-drop
    Medium low      any             1    <default-drop-profile>
    Medium high     any             1    <default-drop-profile>
    High            any         48162    high-drop

Scheduler: l3vpn-pri-sc, Forwarding class: l3vpn-priority, Index: 6318
  Transmit rate: 10 percent, Rate Limit: none, Buffer size: 5000 us,
  Buffer Limit: none, Priority: high
  Excess Priority: unspecified
  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             any             1    <default-drop-profile>
    Medium low      any             1    <default-drop-profile>
    Medium high     any             1    <default-drop-profile>
    High            any             1    <default-drop-profile>

Scheduler: nc-sc, Forwarding class: nc, Index: 15570
  Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
  Buffer Limit: none, Priority: high
  Excess Priority: unspecified
  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             any             1    <default-drop-profile>
    Medium low      any             1    <default-drop-profile>
    Medium high     any             1    <default-drop-profile>
    High            any             1    <default-drop-profile>
```

```
[edit]
lab@R4# run show class-of-service interface ge-0/0/4
Physical interface: ge-0/0/4, Index: 153
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Maximum usable queues: 8, Queues in use: 5
Exclude aggregate overhead bytes: disabled
Logical interface aggregate statistics: disabled
  Scheduler map: core-interfaces, Index: 58651
  Congestion-notification: Disabled

  Logical interface: ge-0/0/4.146, Index: 336
Object                  Name                    Type                    Index
Rewrite-Output          dscp-rewriter           dscp                    20901
Rewrite-Output          mpls-rewriter           exp (mpls-inet-both)     10618
Classifier              dscp-classifier         dscp                    51090
Classifier              dscp-ipv6-compatibility dscp-ipv6                   9
Classifier              mpls-classifier         exp                     16205

  Logical interface: ge-0/0/4.32767, Index: 337
```

```
[edit]
lab@R4# run show interfaces queue ge-0/0/4 egress
Physical interface: ge-0/0/4, Enabled, Physical link is Up
  Interface index: 139, SNMP ifIndex: 514
Forwarding classes: 8 supported, 5 in use
Egress queues: 8 supported, 5 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets              :                3479                     0 pps
    Bytes                :              350306                   424 bps
  Transmitted:
    Packets              :                3479                     0 pps
    Bytes                :              350306                   424 bps
    Tail-dropped packets :                   0                     0 pps
    RED-dropped packets  :                   0                     0 pps
     Low                 :                   0                     0 pps
     Medium-low          :                   0                     0 pps
     Medium-high         :                   0                     0 pps
     High                :                   0                     0 pps
    RED-dropped bytes    :                   0                     0 bps
     Low                 :                   0                     0 bps
     Medium-low          :                   0                     0 bps
     Medium-high         :                   0                     0 bps
     High                :                   0                     0 bps
Queue: 1, Forwarding classes: l3vpn
  Queued:
    Packets              :                   0                     0 pps
    Bytes                :                   0                     0 bps
  Transmitted:
    Packets              :                   0                     0 pps
    Bytes                :                   0                     0 bps
    Tail-dropped packets :                   0                     0 pps
    RED-dropped packets  :                   0                     0 pps
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
      Low                    :                       0                       0 pps
      Medium-low             :                       0                       0 pps
      Medium-high            :                       0                       0 pps
      High                   :                       0                       0 pps
    RED-dropped bytes        :                       0                       0 bps
      Low                    :                       0                       0 bps
      Medium-low             :                       0                       0 bps
      Medium-high            :                       0                       0 bps
      High                   :                       0                       0 bps
Queue: 2, Forwarding classes: l2vpn
  Queued:
    Packets                  :                       0                       0 pps
    Bytes                    :                       0                       0 bps
  Transmitted:
    Packets                  :                       0                       0 pps
    Bytes                    :                       0                       0 bps
    Tail-dropped packets     :                       0                       0 pps
    RED-dropped packets      :                       0                       0 pps
      Low                    :                       0                       0 pps
      Medium-low             :                       0                       0 pps
      Medium-high            :                       0                       0 pps
      High                   :                       0                       0 pps
    RED-dropped bytes        :                       0                       0 bps
      Low                    :                       0                       0 bps
      Medium-low             :                       0                       0 bps
      Medium-high            :                       0                       0 bps
      High                   :                       0                       0 bps
Queue: 3, Forwarding classes: l3vpn-priority
  Queued:
    Packets                  :                   10226                       0 pps
    Bytes                    :                  964620                     384 bps
  Transmitted:
    Packets                  :                   10226                       0 pps
    Bytes                    :                  964620                     384 bps
    Tail-dropped packets     :                       0                       0 pps
    RED-dropped packets      :                       0                       0 pps
      Low                    :                       0                       0 pps
      Medium-low             :                       0                       0 pps
      Medium-high            :                       0                       0 pps
      High                   :                       0                       0 pps
    RED-dropped bytes        :                       0                       0 bps
      Low                    :                       0                       0 bps
      Medium-low             :                       0                       0 bps
      Medium-high            :                       0                       0 bps
      High                   :                       0                       0 bps
Queue: 4, Forwarding classes: nc
  Queued:
    Packets                  :                       0                       0 pps
    Bytes                    :                       0                       0 bps
```


© 2020 Juniper Networks, Inc. All rights reserved.

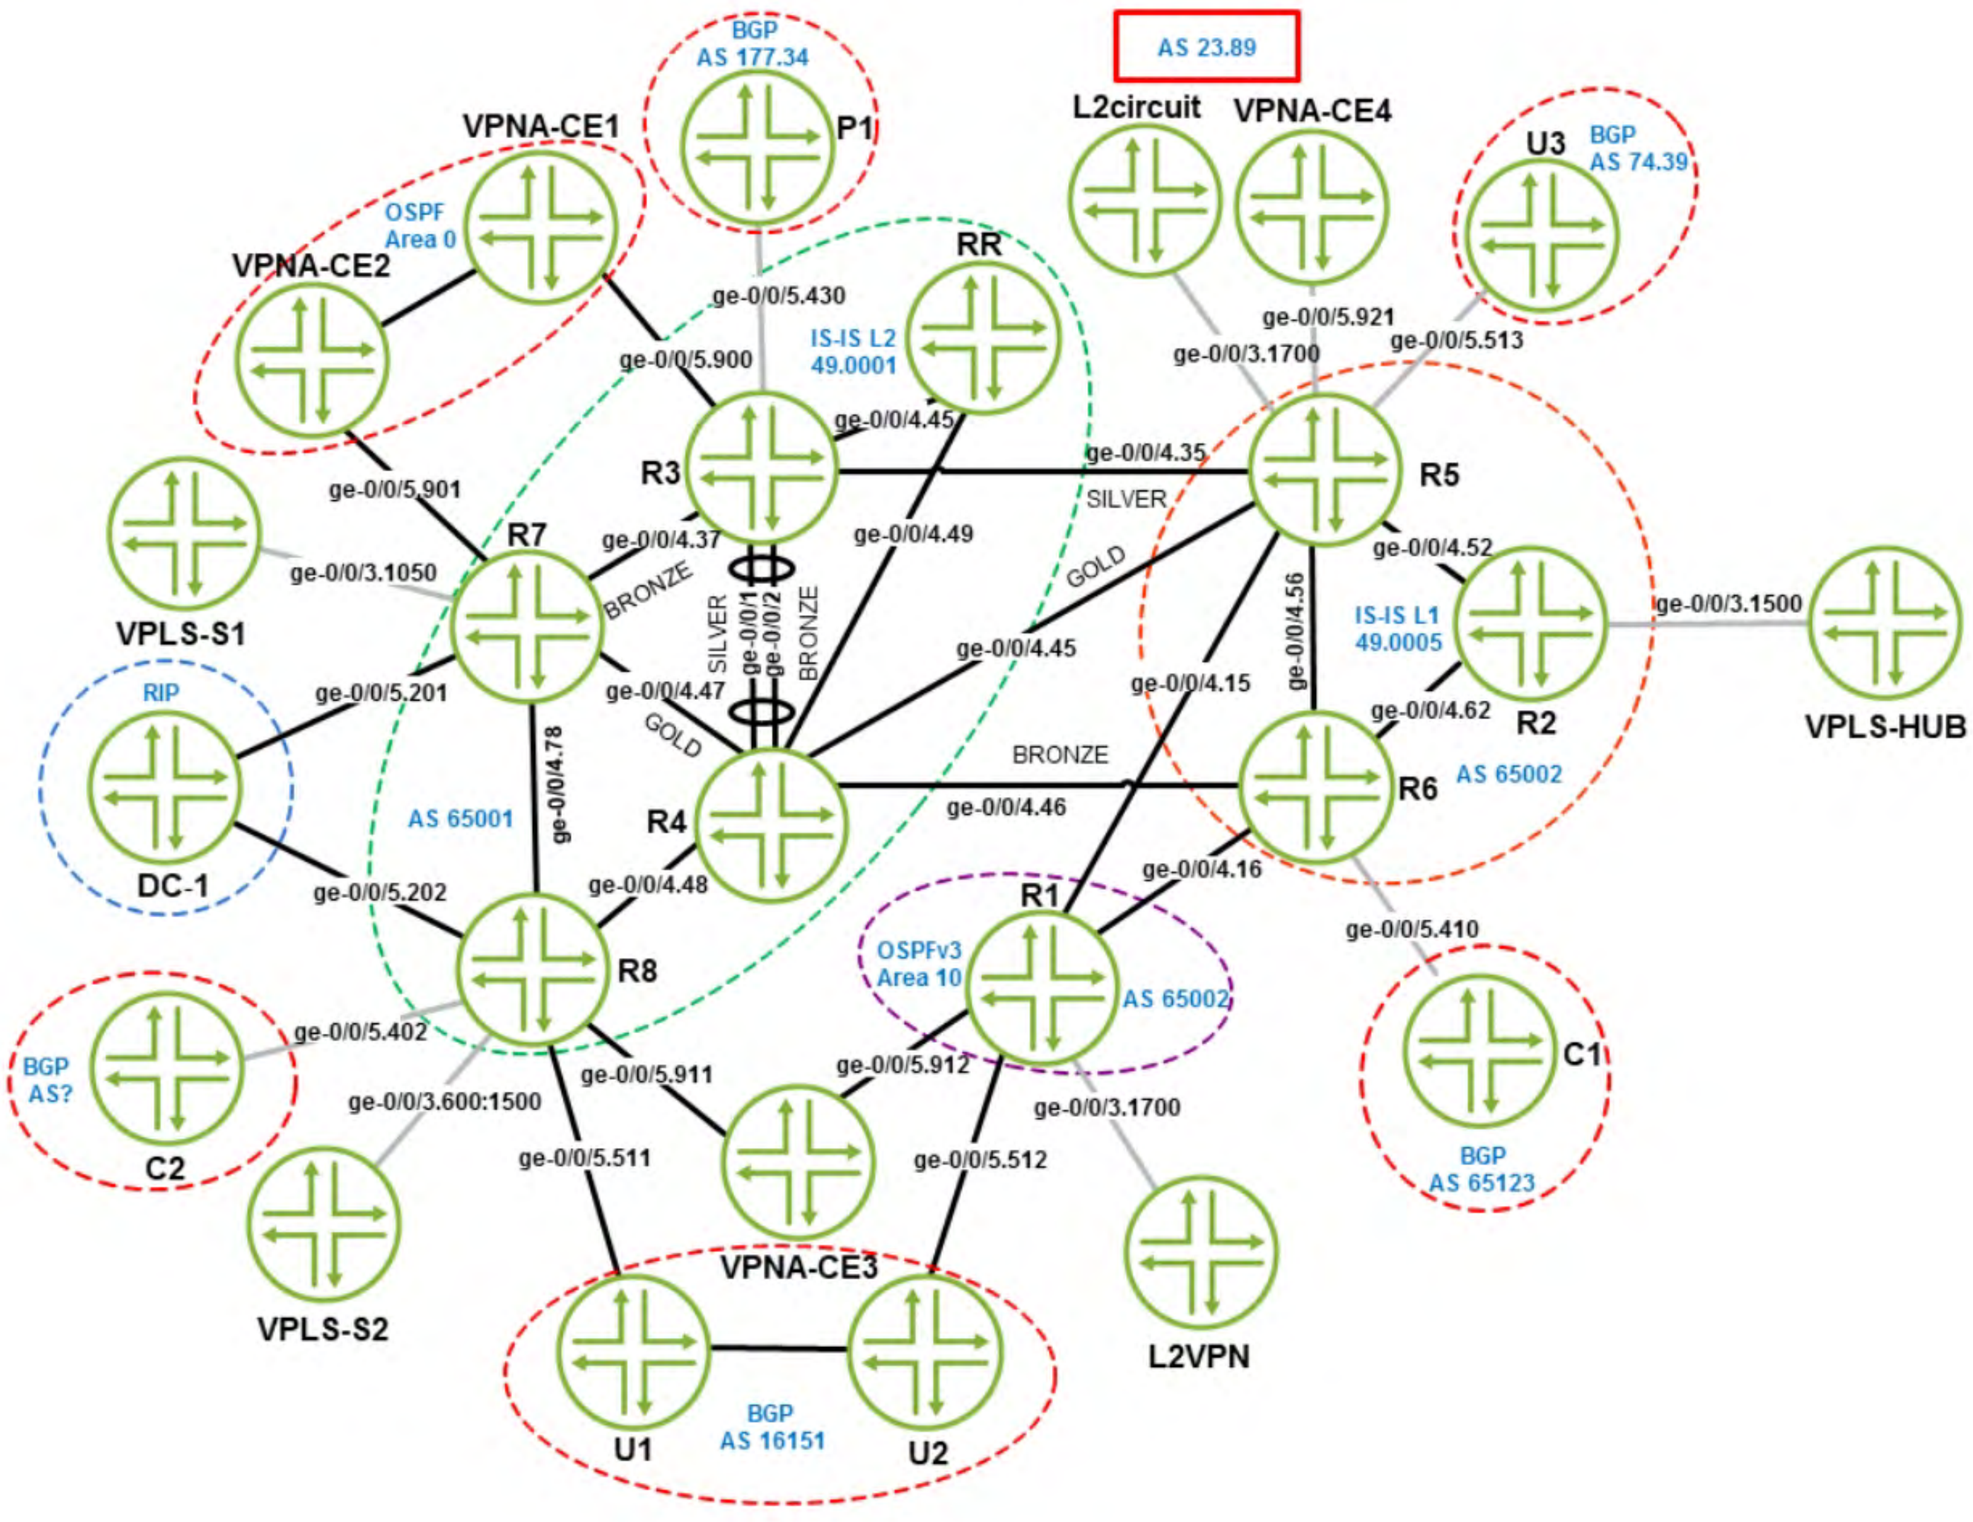
```
Transmitted:
  Packets              :              0                0 pps
  Bytes                :              0                0 bps
  Tail-dropped packets :              0                0 pps
  RED-dropped packets  :              0                0 pps
   Low                 :              0                0 pps
   Medium-low          :              0                0 pps
   Medium-high         :              0                0 pps
   High                :              0                0 pps
  RED-dropped bytes    :              0                0 bps
   Low                 :              0                0 bps
   Medium-low          :              0                0 bps
   Medium-high         :              0                0 bps
   High                :              0                0 bps
```


© 2020 Juniper Networks, Inc. All rights reserved.

Appendix – Super Lab 2

This lab is designed to combine many of the solutions discussed in the previous chapters into a complete practice lab exam. This lab is similar to a real JNCIE-SP lab exam experience. We recommend attempting this lab if you are able to complete the tasks in the previous chapters without any help.

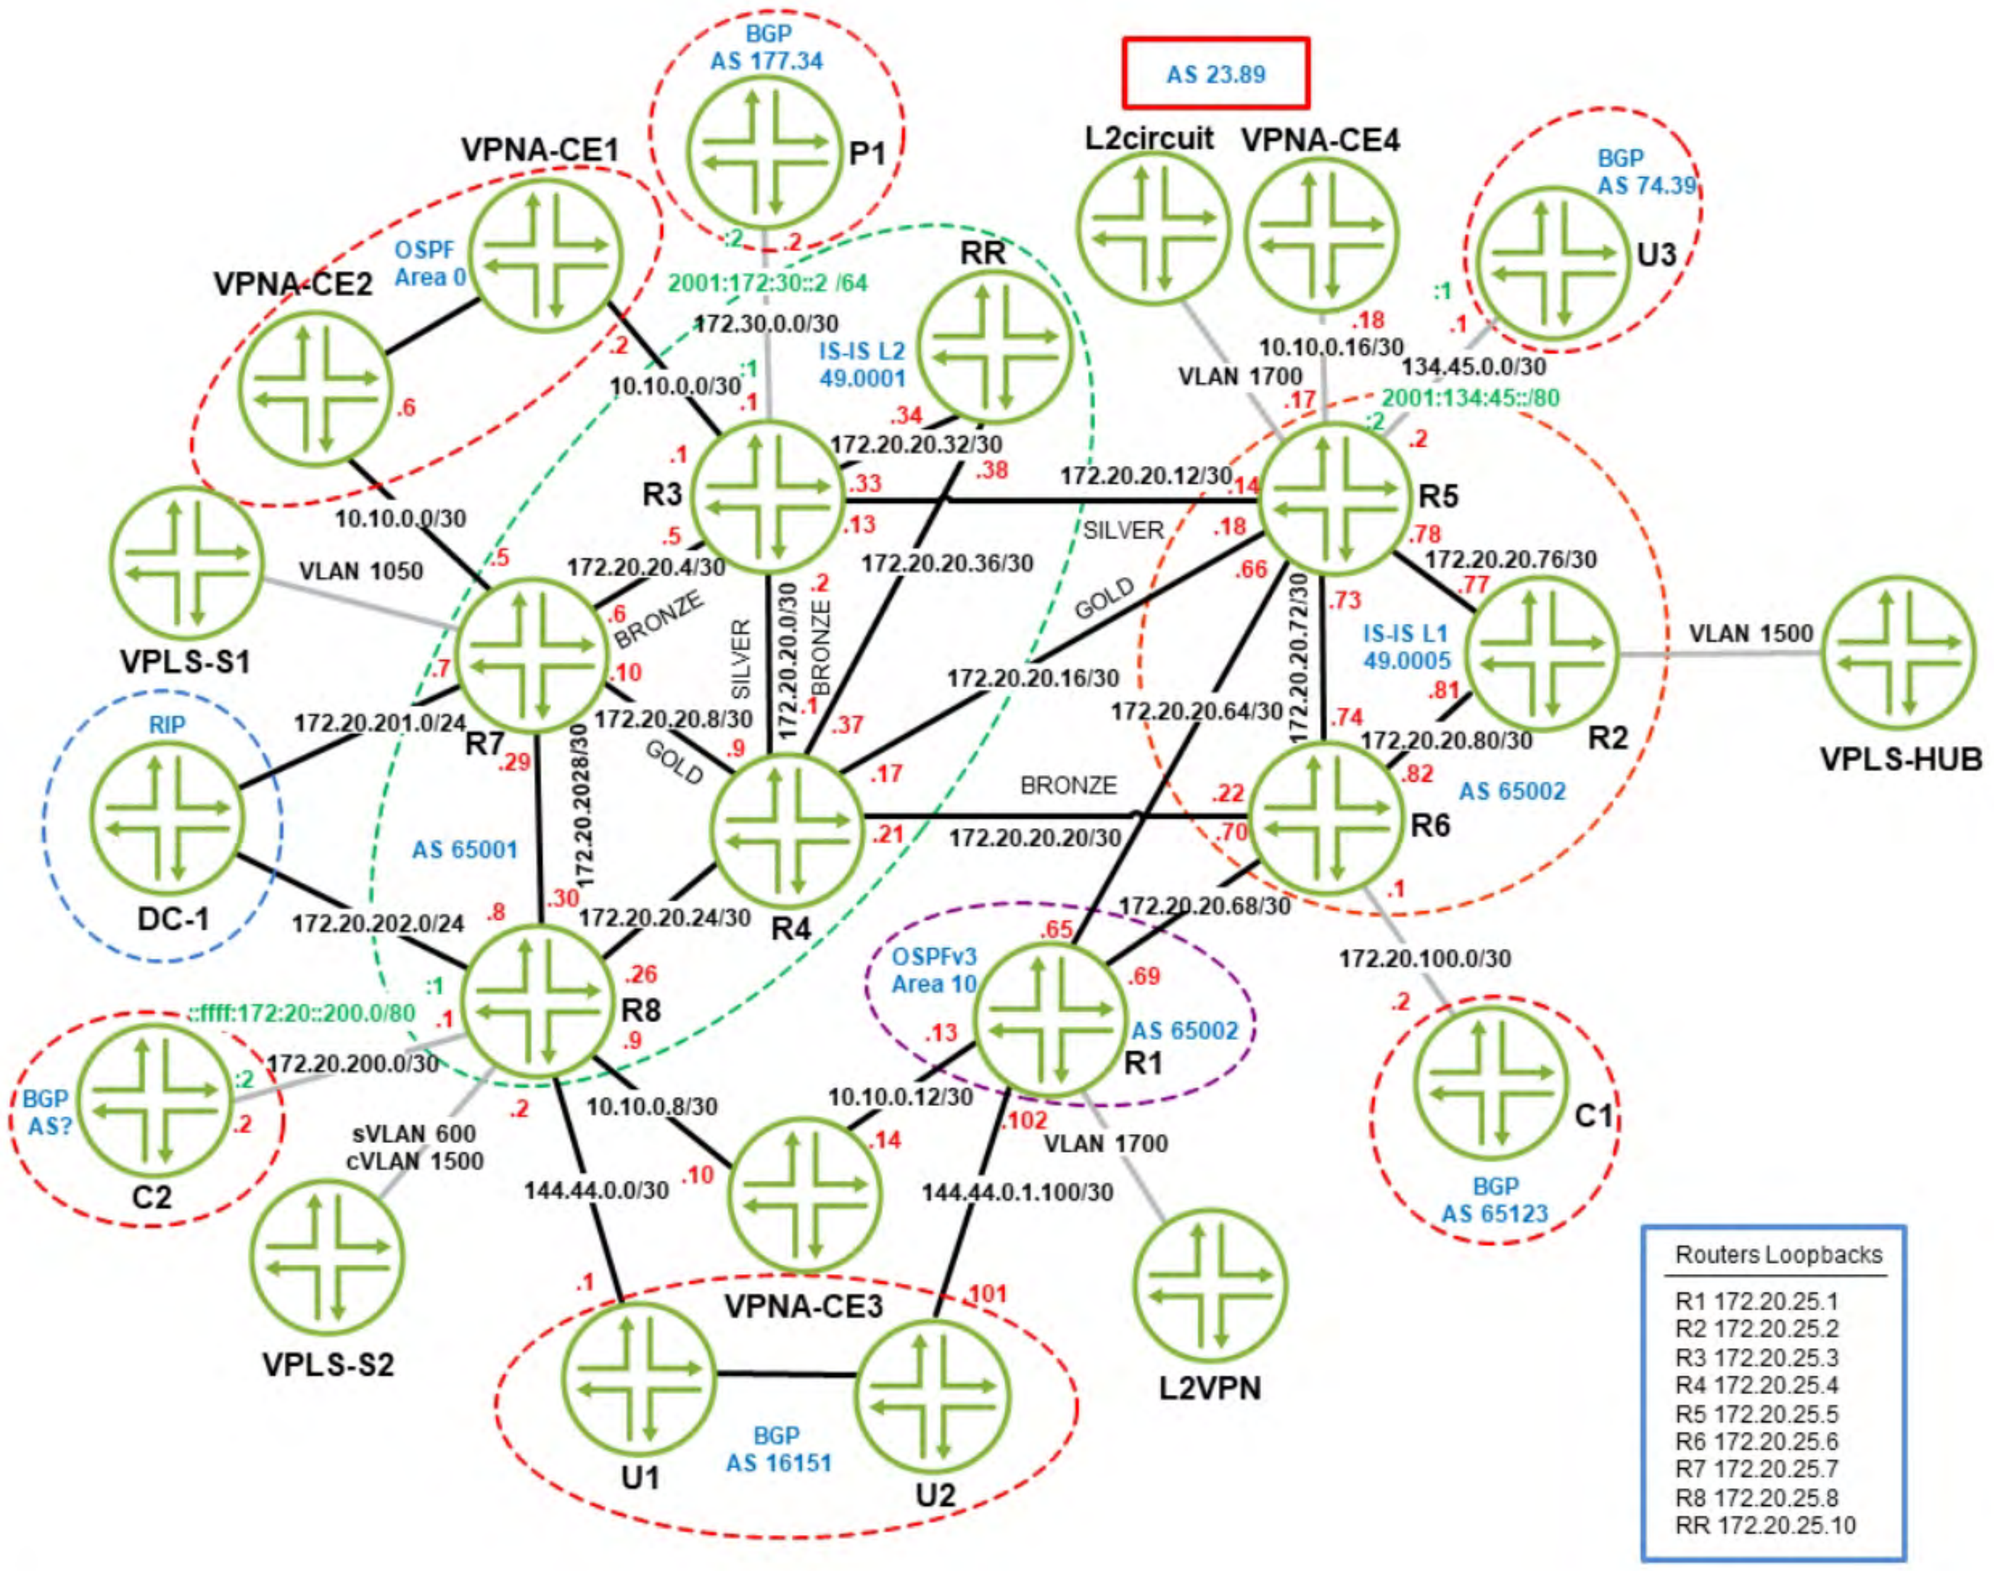
The physical diagram, shown below, visualize the links between the routers. The devices with darker color are the only subject to configuration. The rest of the devices are virtualized on the VR-device, which configuration should not be changed. Use only the interfaces and VLANs on the diagram for connectivity between the routers.



© 2020 Juniper Networks, Inc. All rights reserved.

Physical Diagram

The logical diagram displays all the information needed for configuring the logical connectivity between the devices.



Logical Diagram

Like on the real JNCIE-SP exam, some or all devices have partial configuration. Before you start with solving the Tasks, load the initial configurations provided with the workbook to all 10 physical devices.

Use user *lab* password *lab123* for administrative access to every router.

© 2020 Juniper Networks, Inc. All rights reserved.

Task 1: System Features

First part is focused on initial system configuration, monitoring and securing the Junos operating system. That includes configuring different interface features like bounding physical interface in logical links and applying various address families. You will learn how to configure user accounts and various authentication methods available in Junos. You will configure monitoring and backup system logs, coping automatically the configuration to remote server. Last you will configure a firewall filters to protect the routing-engine of the devices.

Note: It is highly recommended to read the whole lab and verifying the state of the devices before starting with the configuration. Moving from one part to another is also recommended to read all of the tasks for each part you are about the start configuring as next. That will help you to get an idea how the network should look like after at the end.

1) Console access is preconfigured in the initial configuration on all routers. Your task is to increase the security of the console port. Ensure that console session logouts automatically when console cable is disconnected.

    a) All devices including route reflector.

When the cable is disconnected from the console port on Junos device, the user remains logged in. To change the default behavior, you have to use the command line below applied to all routers.

```
[edit]
lab@R1# set system ports console log-out-on-disconnect
```

2) In addition to the SSH protocol, enable Telnet access in case of a need. However, to limit the security risk, limit the Telnet connections to 2 at a time.

    b) All devices including route reflector.

SSH protocol is enabled by the initial configuration, in addition to that the Task requires Telnet protocol to be configured also. Due to unsecure nature of the protocol, limiting the simultaneous session will decrease the risk. The command below will accomplish the requirement.

```
[edit]
lab@R1# set system services telnet connection-limit 2
```

3) Configure all routers to be accessible for file transfer using the FTP protocol. Limit simultaneous connections to only 2.

© 2020 Juniper Networks, Inc. All rights reserved.

c) All devices including route reflector.

SCP protocol allows secure transfer from and to the devices. Similar to the previous Task, the FTP protocol is enabled on all router as a backup method. Additional security measures are needed because of the unencrypted transmission used by FTP protocol.

```
[edit]
lab@R1# set system services ftp connection-limit 2
```

4) Configure all devices to authenticate users attempting to login via RADIUS server located on address 10.10.10.1, using secret *workbook*. The server is configured to accept requests only from the loopback address of the routers.

d) All devices including route reflector.

Junos provides three different methods for user account authentication, authorization and accounting. Those are the local database, using a RADIUS server and using a TACACS+ server. Latter two of which are in essential remotely located databases accessed by Junos device using RADIUS or TACAS+ protocols. Centralized database authentication provides a scalable way for login user configuration. This task requires configuring the devices to communicate successfully with RADIUS server. Parameters like server IP address and shared secret key are mandatory for providing communication with the RADIUS server. In addition to that you must explicitly configure the source IP address of the request send by each of the routers to the loopback IP address.

Note: You must change the source address on each device to correspond to the loopback address.

```
[edit]
lab@R1# set system radius-server 10.10.10.1 secret "$9$uTx31SesYoZDkNdwgJDmPuO1hyK"

[edit]
lab@R1# set system radius-server 10.10.10.1 source-address 172.20.25.1
```

Note: The password here is displayed as it is printed on the terminal. Starting with `$9$` means that a weak algorithm is used to hash the password. Therefore, the password can be derived back from that hash.

5) All user login attempts should fall back to the local used database if not configured on the RADIUS server.

a) All devices including route reflector.

© 2020 Juniper Networks, Inc. All rights reserved.

The task requires the user login attempts to be sent for authentication to the RADIUS server as primary method. However, if the RADIUS rejects the user request, the device should check local database for that user. This can be accomplished when adding the password method as second in the authentication order.

```
[edit]
lab@R1# set system authentication-order radius
[edit]
lab@R1# set system authentication-order password
```

6) Ensure that the reachability to 10.10.10.1/32 is provided even if the rpd process is not yet running. 10.10.1.200 is the gateway for the management segment.

   b) All devices including route reflector.

The **backup-router** command provided routing information during the boot up sequence of Junos, before the rpd process is started yet.

```
[edit]
lab@R1# set system backup-router 10.10.1.200

[edit]
lab@R1# set system backup-router destination 10.10.10.1/32
```

7) Configure on all routers users *noc1* and *noc2*, in such way that they do not fall back to the local database for authentication if the RADIUS server is not accessible. The users must be restricted to only view the configuration and issue operational commands with read-only access. Due to limited expertise of the users, ensure they are not able to request changes in operational mode, like software changes and power-off the devices.

   a) All devices including route reflector.

The task requires configuring of users *noc1* and *noc2* with a restricted access to the devices. Those users must not have access to any commands in configuration mode and to commands requesting changes to the system. In addition, the login attempts by the two users should be rejected if the RADIUS server is not accessible.

```
[edit]
lab@R1# set system login class noc permissions view

[edit]
lab@R1# set system login class noc permissions view-configuration
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set system login class noc deny-commands request

[edit]
lab@R1# set system login user noc1 class noc

[edit]
lab@R1# set system login user noc2 class noc
```

8) In addition to that, restrict the use of those accounts only during the weekdays starting from 8am. to 6:30pm.

You should also restrict the two users to log in only during a specific period in the week.

a) All devices including route reflector.

```
[edit]
lab@R1# set system login class noc allowed-days monday

[edit]
lab@R1# set system login class noc allowed-days tuesday

[edit]
lab@R1# set system login class noc allowed-days wednesday

[edit]
lab@R1# set system login class noc allowed-days thursday

[edit]
lab@R1# set system login class noc allowed-days friday

[edit]
lab@R1# set system login class noc access-start "08:00:00 +0000"

[edit]
lab@R1# set system login class noc access-end "18:30:00 +0000"
```

9) Configure user *support* for cases where additional control over the devices is required. That user should authenticate using password *ops123* if the RADIUS server is inaccessible. The privileges assigned to the user must provide access to the network, reset processes, request changes, clear counters, and view the configuration. However, you are not allowed to configure a custom class for that specific user.

a) All devices including route reflector.

© 2020 Juniper Networks, Inc. All rights reserved.

Junos comes with four predefined login classes: `operator`, `read-only`, `super-user` and `unauthorized`. To accomplish the Task's requirements, you must use the `operator` class.

```
[edit]
lab@R1# set system login user support class operator

[edit]
lab@R1# set system login user support authentication encrypted-password
"$1$Z3BzZuYN$kCx6BfOTWW3s3Z1kSEzjY1"
```

b) Verify the configuration.

The `operator` class included authorization level of access as displayed below.

```
support@R1> show cli authorization
Current user: 'support     ' class 'operator'
Permissions:
    clear       -- Can clear learned network info
    network     -- Can access the network
    reset       -- Can reset/restart interfaces and daemons
    trace       -- Can view trace file settings
    view        -- Can view current values and statistics
Individual command authorization:
    Allow regular expression: none
    Deny regular expression: none
    Allow configuration regular expression: none
    Deny configuration regular expression: none
```

10) During login ensure a banner is displayed warning the users that unauthorized access is not allowed. Use the example below as shown and make sure it is displayed to the user before they are logged in.

```
##########################################

WARNING: Unauthorized access is prohibited!

All sessions are monitored and recorded!

##########################################
```

a) All devices including route reflector.

The system login message is displayed before the user login. Unlike the system login announcement which is displayed after the user login.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set system login message "########################################\nWARNING:
Unauthorized access is prohibited!\nAll sessions are monitored and
recorded!\n########################################"
```

b) Verify the configuration.

```
[edit]
lab@R1#
Using username "lab".
####################################
WARNING: Unauthorized access is prohibited!
All sessions are monitored and recorded!
####################################
Using keyboard-interactive authentication.
Password:
Last login: R1 Dec  4 08:08:32 2016 from 10.10.1.254
--- Junos 15.1F6.9 Kernel 64-bit  JNPR-10.1-20160616.329709_builder_stable_10
```

11) Configure all devices to send the copy of the configuration at each successful commit to a FTP server on address 10.10.10.1. Routers should copy the files in folder */tmp/upload/configurations* using username *jncie* and password *workbook*.

a) All devices including route reflector.

Junos provides two ways for transfer of the configuration to an archive server periodically. Those are, transfer on configured interval of time or every time successful commit is performed. Below is shown how to archive the configuration every time a commit is performed.

```
[edit]
lab@R1# set system archival configuration transfer-on-commit
```

```
[edit]
lab@R1# set system archival configuration archive-sites
"ftp://jncie@10.10.10.1/tmp/upload/configurations" password
"$9$BdTElMg4ZDHmVw2aUH5TBIEyeW"
```

12) Administrators logged in with username *support* should be informed of any configuration change on the console.

a) All devices including route reflector.

© 2020 Juniper Networks, Inc. All rights reserved.

By default, Junos displays only the emergency messages to the user terminal. Users logged in with the *support* account, must receive log messages about configuration changes. The configuration line below accomplishes that.

```
[edit]
lab@R1# set system syslog user support change-log any
```

13) Ensure that all possible messages are logged in the messages syslog file. Limit the message files to not more than 5 files, less than 10 MB each.

    a) All devices including route reflector.

```
[edit]
lab@R1# set system syslog file messages any any

[edit]
lab@R1# set system syslog file messages authorization info

[edit]
lab@R1# set system syslog file messages archive size 10m

[edit]
lab@R1# set system syslog file messages archive files 5
```

14) Configure the interactive-commands syslog file to be transferred once per 12 hours to the FTP server located at address 10.10.10.1. Routers should copy the files in folder */tmp/upload/syslog* using username *jncie* and password *workbook*. To overwrite the default archive size, increase interactive-commands syslog file size to 20 MB.

    a) All devices including route reflector.

By default, Junos rotates the log file when file size reaches 10 m. The archive files are stored on the local storage. However, the task requires increasing the size to 20 m and transferring the files every 12 hours to the FTP server.

```
[edit]
lab@R1# set system syslog file interactive-commands interactive-commands any

[edit]
lab@R1# set system syslog file interactive-commands archive size 20m

[edit]
lab@R1# set system syslog file interactive-commands archive transfer-interval 720
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set system syslog file interactive-commands archive archive-sites
"ftp://jncie@10.10.10.1/tmp/upload/syslog" password "$9$SH8eLNUDkm5F4aGi.56/SreWX-"

[edit]
lab@R1# commit
commit complete
```

15) Configure the aggregated Ethernet interface connecting R3 and R4. Enable active monitoring of each physical interface link. You have to make sure that the bundle always supports a capacity of minimum 1.5Gbps. The aggregated Ethernet interface must not forward traffic if the above requirement is not met.

a) R3

You have to configure aggregated Ethernet interface based on the requirements in the task. Before creating an aggregated interface in Junos, you need to define the number of the ae- interfaces on the chassis. You can run the LACP protocol over the aggregated interface to monitor and control the bundle operation. Junos provide a way to declare the whole ae- interface down if some of the physical links is down as this leads to insufficient bandwidth. Declaring the interface down will force the control plane to find alternative path through the network without dropping traffic due to insufficient bandwidth in the bundle. The configuration lines below show the configuration of aggregated Ethernet interface.

```
[edit]
lab@R3# set chassis aggregated-devices ethernet device-count 1

[edit]
lab@R3# set interfaces ge-0/0/1 gigether-options 802.3ad ae0

[edit]
lab@R3# set interfaces ge-0/0/2 gigether-options 802.3ad ae0

[edit]
lab@R3# set interfaces ae0 unit 0 description "=== connection to R4 ==="

[edit]
lab@R3# set interfaces ae0 aggregated-ether-options minimum-links 2

[edit]
lab@R3# set interfaces ae0 aggregated-ether-options lacp active

[edit]
lab@R3# set interfaces ae0 unit 0 family inet address 172.20.20.1/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# commit
commit complete
```

b) R4

```
[edit]
lab@R4# set chassis aggregated-devices ethernet device-count 1

[edit]
lab@R4# set interfaces ge-0/0/1 gigether-options 802.3ad ae0

[edit]
lab@R4# set interfaces ge-0/0/2 gigether-options 802.3ad ae0


[edit]
lab@R4# set interfaces ae0 unit 0 description "=== connection to R3 ==="

[edit]
lab@R4# set interfaces ae0 aggregated-ether-options minimum-links 2

[edit]
lab@R4# set interfaces ae0 aggregated-ether-options lacp active

[edit]
lab@R4# set interfaces ae0 unit 0 family inet address 172.20.20.2/30

[edit]
lab@R4# commit
commit complete
```

c) Verify the configuration.

The output shows the ae- interface created and the state of the LACP protocol.

```
[edit]
lab@R3# run show lacp interfaces ae0
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/0/1       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/1     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/2       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/2     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-0/0/1                  Current   Fast periodic Collecting distributing
      ge-0/0/2                  Current   Fast periodic Collecting distributing
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# run show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 149, SNMP ifIndex: 530
  Link-level type: Ethernet, MTU: 1514, Speed: 2Gbps, BPDU Error: None, MAC-REWRITE
Error: None, Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled
  Pad to minimum frame size: Disabled
  Minimum links needed: 2, Minimum bandwidth needed: 1bps
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:05:86:49:ba:c0, Hardware address: 00:05:86:49:ba:c0
  Last flapped   : 2016-12-04 09:42:26 UTC (00:00:25 ago)
  Input rate     : 1440 bps (0 pps)
  Output rate    : 2408 bps (2 pps)
...
```

16) Enable support for switching labeled packets, iso and IPv6 traffic on all logical interfaces (except management interface) on all routers except the RR. You must use the most efficient way configuration-wise to accomplish this task.

    a) All devices including route reflector.

The task requires the use of configuration groups to apply the mpls, inet6 and iso families to all interfaces. Configuration groups are inherited by the configuration file if applied to the CLI hierarchy they are intended to configure. In the same way, they can be excluded deeper in the hierarchy for specific statements. Wildcard masks are usually used to match repetitive but diverse statements.

```
[edit]
lab@R1# set groups enable_mpls_iso interfaces <*> unit <*> family iso

[edit]
lab@R1# set groups enable_mpls_iso interfaces <*> unit <*> family inet6

[edit]
lab@R1# set groups enable_mpls_iso interfaces <*> unit <*> family mpls

[edit]
lab@R1# set interfaces apply-groups enable_mpls_iso

[edit]
lab@R1# set interfaces fxp0 apply-groups-except enable_mpls_iso

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) Verify the configuration.

To display the configuration inherited by the apply-group configuration the command below must be used.

```
[edit]
lab@R3# show interfaces ge-0/0/4 | display inheritance
vlan-tagging;
unit 35 {
    description "=== connection to R5 ===";
vlan-id 35;
family inet {
    address 172.20.20.13/30;
}
##
## 'iso' was inherited from group 'enable_mpls_iso'
##
family iso;
##
## 'inet6' was inherited from group 'enable_mpls_iso'
##
family inet6;
##
## 'mpls' was inherited from group 'enable_mpls_iso'
##
family mpls;
   }
```

17) Enable graceful-restart on all routers for all routing protocols except BGP.

c) All devices including route reflector.

Graceful-restart is disabled by default in Junos. To enable it globally, you must include **graceful-restart** under the routing-options hierarchy. This will enable the feature for every routing protocol that supports it. However, you must disable that functionality for the BGP protocol explicitly to complete the task requirement.

```
[edit]
lab@R1# set routing-options graceful-restart

[edit]
lab@R1# set protocols bgp graceful-restart disable
```

18) You must ensure that the control plane of all routers is protected. Configure a firewall filter using the following criteria:

© 2020 Juniper Networks, Inc. All rights reserved.

Last task is the most time consuming because you must configure firewall filters with many terms. In addition, this is the first place where you must also include IPv6 specific configuration.

By default, all filters created without address family are inet filters or in other words IPv4 filters. Filter matching IPv6 traffic must be configured explicitly for family inet6.

- Allow main IPv4 protocols that are already configured and would be required by the subsequences tasks: SSH, BFD, OSPF, RIP, PIM, LDP, RSVP, RADIUS, ping and traceroute.

  a) All devices including route reflector.

SSH protocol is allowed for remote management to the device.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ssh from protocol tcp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ssh from destination-port ssh

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ssh then accept
```

RADIUS replies sent by the server are accepted to the RE.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-radius from protocol udp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-radius from source-port radius

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-radius then accept
```

The BFD protocol is allowed in case it is used for end-to-end monitoring by any of the routing protocol in the network.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bfd from protocol udp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bfd from port 3784
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bfd then accept
```

OSPF and BGP routing protocols traffic are allowed to the RE.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ospf from protocol ospf

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ospf then accept

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bgp from protocol tcp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bgp from port bgp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bgp then accept
```

The RIP protocol is also accepted.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-rip from protocol udp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-rip from port rip

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-rip then accept
```

RSVP and LDP protocols are accepted by the RE filter.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-rsvp from protocol rsvp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-rsvp then accept

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ldp from protocol tcp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ldp from protocol udp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ldp from port ldp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ldp then accept
```

RE accepts PIM protocol.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-pim from protocol pim

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-pim then accept
```

ICMP echo-reply and echo-request are accepted to permit the use of ping network tool.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ping from protocol icmp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ping from icmp-type echo-reply

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ping from icmp-type echo-
request

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ping then accept
```

traceroute network tool uses icmp and udp protocols for tracing the IP path through the network.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-traceroute from protocol udp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-traceroute from destination-
port 33434-33534

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-traceroute then accept

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-traceroute-reply from protocol
icmp


[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-traceroute-reply from icmp-type
time-exceeded
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-traceroute-reply from icmp-type
unreachable

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-traceroute-reply then accept
```

- Allow Telnet access from 172.20.20.0/24 and 172.20.25.0/24 subnets. To limit the size of the firewall filter, ensure that any additional changes to the allowed subnets does not require firewall filter intervention.

a) All devices including route reflector.

To increase the security of the routers, Telnet sessions should be restricted from only controlled number of sources. However, the task requires using prefix-list configuration to limit the size of the firewall filters. Prefix-lists provide a way for grouping subnets and hosts under named list, which can be used in routing policies or firewall filter.

```
[edit]
lab@R1# set policy-options prefix-list telnet-hosts 10.10.1.0/24

[edit]
lab@R1# set policy-options prefix-list telnet-hosts 172.20.20.0/24

[edit]
lab@R1# set policy-options prefix-list telnet-hosts 172.20.25.0/24

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-telnet from source-prefix-list
telnet-hosts

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-telnet from protocol tcp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-telnet from port telnet

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-telnet then accept
```

Because the new term will be added at the end of the filter, it must be inserted before the explicit discard route.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# insert firewall filter protect-re-ipv4 term accept-telnet before term discard-
all
```

- Allow only active FTP to the routers.
  a) All devices including route reflector.

The Task requires permitting active FTP only. The FTP protocol uses two session to operate, control session and data session. In active FTP mode, the client is initiating the control session to the server using port 21. The data session is then sourced by port 20 from the server to the client. Therefore, you must accept a session with destination port 20 ingress to the RE, Accepting only the established connections only to port 20 from the clients will increase security

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ftp-control from protocol tcp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ftp-control from port ftp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ftp-control then accept

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ftp-data from protocol tcp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ftp-data from port ftp-data

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ftp-data from tcp-established

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ftp-data then accept

[edit]
lab@R1# insert firewall filter protect-re-ipv4 term accept-ftp-control before term
discard-all

[edit]
lab@R1# insert firewall filter protect-re-ipv4 term accept-ftp-data before term
discard-all
```

- You must accept BGP sessions only from configured neighbors. Any additional neighbors must be automatically accepted.

Apply-path prefix lists are similar to the configuration groups but instead of populating the configuration, they extract configuration statements and create dynamic list of IP

© 2020 Juniper Networks, Inc. All rights reserved.

addresses. Like the configuration groups they also use wildcards to match part of the configuration representing free text.

```
[edit]
lab@R1# set policy-options prefix-list bgp-neighbors apply-path "protocols bgp group
<*> neighbor <*>"

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bgp from source-prefix-list
bgp-neighbors

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bgp from protocol tcp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bgp from port bgp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bgp then accept

[edit]
lab@R1# insert firewall filter protect-re-ipv4 term accept-bgp before term discard-all
```

You should not forget to apply the filter to the loopback interface. Applying it there will protect the RE from traffic coming through any of the physical interfaces. You must apply the filter for the appropriate family.

```
[edit]
lab@R1# set interfaces lo0 unit 0 family inet filter input protect-re-ipv4
```

- Allow main IPv6 protocols that would be required by the subsequent tasks: OSPFv3, BGP, RIPng, ping and traceroute.

Unlike the IPv4 filters, with IPv6 you have to specify the Layer 4 protocols as a next-header match condition.

```
[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-ospfv3 from next-
header ospf

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-ospfv3 then accept


[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-bgp from source-
prefix-list bgp-neighbors
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-bgp from next-
header tcp

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-bgp from
destination-port bgp

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-bgp then accept
```

The RIP protocol uses UDP port 520. However, IPv6 version, the RIPng uses UDP port 521.

```
[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-ripng from next-
header udp

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-ripng from port
521

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-ripng then accept
```

Neighbor Discovery or ND is a function of IPv6 protocol that replaces the ARP protocol used for IPv4 mapping the MAC address to IP address of the hosts on the broadcast segment. Therefore, it is very important for the operation of IPv6 and must be allowed by the filter protecting the RE. IPv6 Neighbor Discovery uses ICMP messages so the next-header should match the ICMP protocol.

```
[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-nd from next-
header icmp6

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-nd from icmp-type
router-solicit

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-nd from icmp-type
router-advertisement


[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-nd from icmp-type
neighbor-solicit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-nd from icmp-type
neighbor-advertisement

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-nd then accept

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-ping from next-
header icmp6

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-ping from icmp-
type echo-reply

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-ping from icmp-
type echo-request

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-ping then accept

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-traceroute from
next-header udp

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-traceroute from
destination-port 33434-33543

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-traceroute then
accept

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-traceroute-reply
from next-header icmp

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-traceroute-reply
from icmp-type time-exceeded

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-traceroute-reply
from icmp-type destination-unreachable

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-traceroute-reply
then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Remember to apply the filter to the loopback interface. Applying it there will protect the RE from traffic coming through any of the physical interfaces. You must apply the filter in the appropriate family.

```
[edit]
lab@R1# set interfaces lo0 unit 0 family inet6 filter input protect-re-ipv6
```

- Silently discard all other traffic received by the routers. Log the packets discarded by the filter. The log information must be available after a reboot.

    a) All devices including route reflector.

Silently discarding the traffic requires the use of the discard action instead of reject. To log and store the logs locally you must add the syslog knob in the term. Lastly, you should configure a syslog file logging the firewall events.

```
[edit]
lab@R1# set system syslog file audit-trail firewall any

[edit]
lab@R1# set firewall filter protect-re-ipv4 term discard-all then syslog

[edit]
lab@R1# set firewall filter protect-re-ipv4 term discard-all then discard

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term discard-all then syslog

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term discard-all then discard

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 2: Troubleshooting and Configuring IGP

This part is focused on troubleshooting a partially functional IGP network, build with the IS-IS and OSPFv3 routing protocols. You must consider the given list of requirements that the network should meet. It is highly recommended to read the whole part before starting to solve the tasks.

The first task requires fixing the partial working IGP configuration. Using the diagram provided, you must make sure that all adjacencies are established, and routes are exchanged successfully.

1) Ensure that all IS-IS adjacencies are in the UP state. Make sure the following criteria are met:

    a) All adjacencies that are configured with authentication must authenticate successfully. The system ID should be based on the loopback IP address.
    b) Any single interface should have only one adjacency in the appropriate level as shown in the diagram.

First start with examining the current state of the network. List the adjacency table on each router to determine the state to the neighbors.

Router R1 is part of the OSPFv3 domain only, connected to router R5 and R6.

a) IS-IS adjacency state on R2.

At first glance, it seems all adjacencies on router R2 are established. However, the IS-IS level shown on the diagram is different in the output.

```
[edit]
lab@R2# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.52           R5               2  Up                   8  2:6:a:e:5:4
ge-0/0/4.62           R5               2  Up                   8  2:6:a:e:6:4

[edit]
lab@R2# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
ge-0/0/4.52           2   0x1 Disabled        R5.03                10/10
ge-0/0/4.62           2   0x1 Disabled        R5.02                10/10
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) IS-IS adjacency state on R5.

Router R5 has 4 interfaces in IS-IS which corresponds to the diagram, but only 3 of the adjacencies are established successfully. Again, the level of the adjacencies is different than it should be, according to the diagram. Level 3 suggests that the interfaces are configured with both Level 1 and Level 2, which is, by default, in Junos.

```
[edit]
lab@R5# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.35           R3               2  Up                21  2:6:a:e:3:4
ge-0/0/4.45           R4               2  Up                24  2:6:a:e:4:4
ge-0/0/4.52           R2               2  Up                19  2:6:a:e:2:4

[edit]
lab@R5# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
ge-0/0/4.35           3   0x4 R5.00           R5.04                10/10
ge-0/0/4.45           3   0x2 R5.00           R5.02                10/10
ge-0/0/4.52           3   0x3 R5.00           R5.03                10/10
ge-0/0/4.56           3   0x1 R5.00           R5.00                10/10
```

c) IS-IS adjacency state on R6.

Like router R5, R6 has one interface more added to IS-IS as opposed to the number of the adjacencies. Again, the adjacency Level is not correct.

```
[edit]
lab@R6# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.46           R4               2  Up                26  2:6:a:e:4:4
ge-0/0/4.62           R2               2  Up                25  2:6:a:e:2:4

[edit]
lab@R6# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
ge-0/0/4.46           3   0x3 R6.00           R6.03                10/10
ge-0/0/4.56           3   0x1 R6.00           R6.00                10/10
ge-0/0/4.62           3   0x2 R6.00           R6.02                10/10
```

d) IS-IS adjacency state on R3.

© 2020 Juniper Networks, Inc. All rights reserved.

Two out of three adjacencies on router R3 are listed in the adjacency table. However, only one of the adjacencies is in the Up state. The interface table contains 2 interfaces only.

```
[edit]
lab@R3# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 R4               1  Initializing          25  2c:6b:f5:54:61:c3
ae0.0                 R4               2  Up                     8  2c:6b:f5:54:61:c3
ge-0/0/4.35           R5               2  Up                    25  0:50:56:a9:2a:be
ge-0/0/4.37           0172.0020.2507   2  Up                     6  0:50:56:a9:64:6c

[edit]
lab@R3# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ae0.0                 3   0x1 R3.00             R4.02                   10/10
ge-0/0/4.35           2   0x1 Disabled          R5.04                   10/10
ge-0/0/4.37           2   0x1 Disabled          0172.0020.2507.03       10/10
```

e) IS-IS adjacency state on R4.

Router R4 is the router with the most adjacencies according the diagram. You might notice that the System ID of router R7 is not resolved from the LSP.

```
[edit]
lab@R4# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 R3               2  Up                 21  2c:6b:f5:10:1:c3
ge-0/0/4.45           R5               2  Up                  6  2:6:a:e:5:4
ge-0/0/4.46           R5               2  Up                  6  2:6:a:e:6:4
ge-0/0/4.47           0172.0020.2507   2  Up                  7  2:6:a:e:7:4

[edit]
lab@R4# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ae0.0                 3   0x2 R4.00             R4.02                   10/10
ge-0/0/4.45           3   0x1 R4.00             R5.02                   10/10
ge-0/0/4.46           3   0x1 R4.00             R5.03                   10/10
ge-0/0/4.47           3   0x1 R4.00             0172.0020.2507.02       10/10
ge-0/0/4.48           3   0x1 R4.00             R4.00                   10/10
```

f) IS-IS adjacency state on R7.

© 2020 Juniper Networks, Inc. All rights reserved.

Router R5 has 3 interfaces configured in IS-IS which corresponds to the diagram, but only 2 of the adjacencies are established successfully. Similarly, on router R4, the System ID of the adjacencies routers is not resolved correctly.

```
[edit]
lab@R7# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.37           0172.0020.2503   2  Up                   25  2:6:a:e:3:4
ge-0/0/4.47           0172.0020.2504   2  Up                   25  2:6:a:e:4:4

[edit]
lab@R7# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/4.37           2   0x3 Disabled          R7.03                   10/10
ge-0/0/4.47           2   0x2 Disabled          R7.02                   10/10
ge-0/0/4.78           2   0x1 Disabled          R7.00                   10/10
```

g) IS-IS adjacency state on R8.

Router R8 has none of the adjacencies established. In addition, in the IS-IS interface table the DR router ID contains only 0s.

```
[edit]
lab@R8# run show isis adjacency

[edit]
lab@R8# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/4.48           3   0x1 0000.0000.0000.00 0000.0000.0000.00      10/10
ge-0/0/4.78           3   0x1 0000.0000.0000.00 0000.0000.0000.00      10/10
```

Let's fix the problems found according to the task requirements.

h) Fixing the adjacency on R2.

Router R2 should have only Level 1 adjacencies, disabling Level 2 globally for IS-IS will force all adjacencies to establish Level 1 connections only.

```
[edit]
lab@R2# set protocols isis level 2 disable

[edit]
lab@R2# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete



[edit]
lab@R2# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.52           R6               2  Down               0  2:6:a:e:5:4
ge-0/0/4.62           R6               2  Down               0  2:6:a:e:6:4
```

After committing the change, router R2 is not able to successfully establish the adjacencies with routers R5 and R6 anymore.

Looking in IS-IS packets received on the interfaces of routers R5 and R6 may help with finding the problem. Junos provides a tool for monitoring the content of the packets send to and from the local RE.

```
[edit]
lab@R2# run monitor traffic interface ge-0/0/4.52 no-resolve size 1500 detail
Address resolution is OFF.
Listening on ge-0/0/4.52, capture size 1500 bytes

19:22:54.684183  In IS-IS, length 84
      L1 Lan IIH, hlen: 27, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
        source-id: 0172.0020.2505,  holding time: 27s, Flags: [Level 1 only]
        lan-id:    0172.0020.2505.00, Priority: 64, PDU length: 84
          Protocols supported TLV #129, length: 2
            NLPID(s): IPv4 (0xcc), IPv6 (0x8e)
          IPv4 Interface address(es) TLV #132, length: 4
            IPv4 interface address: 172.20.20.78
          IPv6 Interface address(es) TLV #232, length: 16
            IPv6 interface address: fe80::206:a00:340e:504
          Area address(es) TLV #1, length: 4
            Area address (length: 3): 49.0005
          Restart Signaling TLV #211, length: 3
            Flags [none], Remaining holding time 0s
          Authentication TLV #10, length: 16
            simple text password: jncie-sp-level1
^C
2 packets received by filter
0 packets dropped by kernel



[edit]
lab@R2# run monitor traffic interface ge-0/0/4.62 no-resolve size 1500 detail
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Address resolution is OFF.
Listening on ge-0/0/4.62, capture size 1500 bytes

19:23:57.341923  In IS-IS, length 84
      L1 Lan IIH, hlen: 27, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
        source-id: 0172.0020.2505,  holding time: 27s, Flags: [Level 1 only]
        lan-id:    0172.0020.2505.00, Priority: 64, PDU length: 84
          Protocols supported TLV #129, length: 2
            NLPID(s): IPv4 (0xcc), IPv6 (0x8e)
          IPv4 Interface address(es) TLV #132, length: 4
            IPv4 interface address: 172.20.20.82
          IPv6 Interface address(es) TLV #232, length: 16
            IPv6 interface address: fe80::206:a00:3e0e:604
          Area address(es) TLV #1, length: 4
            Area address (length: 3): 49.0005
          Restart Signaling TLV #211, length: 3
            Flags [none], Remaining holding time 0s
          Authentication TLV #10, length: 16
            simple text password: jncie-sp-level1
^C
1 packets received by filter
0 packets dropped by kernel
```

The output suggests that authentication is used on the other side of the adjacencies. The type of authentication used is simple, in other works not encrypted. Hence, the password is visible in the authentication TLV carried by the LSP. Since both adjacencies are configured with the same password, you can configure the authentication at the Level 1 hierarchy.

```
[edit]
lab@R2# set protocols isis level 1 authentication-key "$9$VmsgJTQnCA0ik5FnCpuLxN-
s4mP536AoJ9p0BSyNdbsoJ"

[edit]
lab@R2# set protocols isis level 1 authentication-type simple

[edit]
lab@R2# commit
commit complete
```

Even correcting the missing authentication, the adjacencies are still not established.

```
[edit]
lab@R2# run show isis adjacency
```

Next step is troubleshooting the R2 IS-IS adjacencies by enabling traceoptions.

```
[edit]
lab@R2# set protocols isis traceoptions file ISIS-log
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set protocols isis traceoptions flag hello detail

[edit]
lab@R2# set protocols isis traceoptions flag error detail

[edit]
lab@R2# commit
commit complete
```

The logs are collected in the configured file and can be shown using the command below.

```
[edit]
lab@R2# run show log ISIS-log
Dec  4 10:04:09.199452 Received L1 LAN IIH, source id 0172.0020.2505 on ge-0/0/4.52
Dec  4 10:04:09.199516     intf index 332, snpa 2:6:a:e:5:4
Dec  4 10:04:09.199522     max area 0, circuit type l1, packet length 84
Dec  4 10:04:09.199528     hold time 27, priority 64, circuit id 0172.0020.2505.00
Dec  4 10:04:09.199532     speaks IP
Dec  4 10:04:09.199535     speaks IPV6
Dec  4 10:04:09.199555     IP address 172.20.20.78
Dec  4 10:04:09.199565     IPV6 address fe80::206:a00:340e:504
Dec  4 10:04:09.199575     area address 49.0005 (3 bytes)
Dec  4 10:04:09.199585     restart flags []
Dec  4 10:04:09.199589     16 bytes of authentication data
Dec  4 10:04:09.199596 ERROR: IIH from 0172.0020.2505 with no matching areas, interface
ge-0/0/4.52
Dec  4 10:04:09.199600     local area  49.0001
Dec  4 10:04:09.199604     area address 49.0005 (3 bytes)
```

The output of the log file reveals that the Area ID configured on router R2 is incorrectly set to ID 49.0001. Note that the Area ID is part of the sysid in IS-IS, which is defined by the iso address configured under the loopback interface `family iso` hierarchy. We use the rename command shown below to correct the wrong Area ID.

```
[edit]
lab@R2# rename interfaces lo0.0 family iso address 49.0001.0172.0020.2501.00 to address
49.0005.0172.0020.2501.00

[edit]
lab@R2# commit
commit complete
```

After committing the change, router R2 successfully establishes the Level 1 IS-IS adjacency.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.52           R5               1  Up                    6  2:6:a:e:5:4
ge-0/0/4.62           R5               1  Up                    8  2:6:a:e:6:4

[edit]
lab@R2# delete protocols isis traceoptions

[edit]
lab@R2# commit
commit complete
```

i) Fixing the adjacency on R5.

Star configuring the interfaces with the correct IS-IS levels as required by the diagram.

```
[edit]
lab@R5# set protocols isis interface ge-0/0/4.35 level 1 disable

[edit]
lab@R5# set protocols isis interface ge-0/0/4.45 level 1 disable

[edit]
lab@R5# set protocols isis interface ge-0/0/4.52 level 2 disable

[edit]
lab@R5# set protocols isis interface ge-0/0/4.56 level 2 disable

[edit]
lab@R5# commit
commit complete
```

Enabling the IS-IS traceoptions helps with solving the missing adjacency.

```
[edit]
lab@R5# set protocols isis traceoptions file ISIS_log

[edit]
lab@R5# set protocols isis traceoptions flag error detail

[edit]
lab@R5# set protocols isis traceoptions flag hello detail

[edit]
lab@R5# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

The file captures an error, that IS-IS hello packets from R6 are ignored due to duplicated sysid.

```
[edit]
lab@R5# run show log ISIS_log
Dec  4 09:35:24 trace_on: Tracing to "/var/log/ISIS_log" started
Dec  4 09:35:33.647422 ERROR: ISIS ignored a bad packet: IIH with duplicate sysid on
interface ge-0/0/4.56
```

Indeed, checking the initial configuration proves that both routers R5 and R6 have the same sysid configured on their loopback interfaces.

```
[edit]
lab@R6# rename interfaces lo0.0 family iso address 49.0005.0172.0020.2505.00 to address
49.0005.0172.0020.2506.00

[edit]
lab@R6# commit
commit complete
```

Renaming the sysid on R6 with the correct one derived from the loopback IP address. After committing the configuration, the adjacency table contains all 4 neighbors in the Up state.

```
[edit]
lab@R5# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/4.35           R3             2  Up                   20
ge-0/0/4.45           R4             2  Up                   19
ge-0/0/4.52           R2             1  Up                   25  2:6:a:e:2:4
ge-0/0/4.56           R6             1  Up                    8  2:6:a:e:6:4

[edit]
lab@R5# delete protocols isis traceoptions

[edit]
lab@R5# commit
commit complete
```

j) Fixing the adjacency on R6.

Similar like R5, configure only the required IS-IS levels for each interface on R6 as shown in the diagram.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols isis interface ge-0/0/4.46 level 1 disable


[edit]
lab@R6# set protocols isis interface ge-0/0/4.56 level 2 disable

[edit]
lab@R6# set protocols isis interface ge-0/0/4.62 level 2 disable

[edit]
lab@R6# commit
commit complete

[edit]
lab@R6# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.46           R4               2  UP                   23
ge-0/0/4.56           R5               1  Up                      26  2:6:a:e:5:4
ge-0/0/4.62           R2               1  Up                      22  2:6:a:e:2:4
```

k) Fixing the adjacency on R3.

By disabling Level 1 ensures that there can only be Level 2 adjacencies on router R3. The issue with the adjacency to R7 is solved in the next steps.

```
[edit]
lab@R3# set protocols isis level 1 disable

[edit]
lab@R3# commit
commit complete

[edit]
lab@R3# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 R4               2  Up                   0  0:5:86:a7:de:c0
ge-0/0/4.35           R5               2  Up                  20
ge-0/0/4.37           0172.0020.2507   2  Up                   8  2:6:a:e:7:4
```

l) Fixing the adjacency on R4.

Initial configuration on router R4 includes Level 1 authentication preconfigured. However, R4 is only a Level 2 router, which requires disabling Level 1 completely.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# show protocols isis
no-ipv4-routing;
level 1 {
    authentication-key "$9$8pAx-wHkPQz6oJi.PQ3nrevMxdUDimfzVwT369OBvWLxVw"; ## SECRET-
DATA
    authentication-type simple;
}
level 2 {
    authentication-key "$9$m5z6hclMWxp0EylMLXDikq5F1IEreW36vLxNY2k.PTn/"; ## SECRET-
DATA
    authentication-type simple;
}
interface ge-0/0/4.45;
interface ge-0/0/4.46;
interface ge-0/0/4.47 {
    level 2 {
        hello-authentication-key "$9$OI1OIhrWLNYgJevgJGDmPBIE"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface ge-0/0/4.48;
interface ae0.0;

[edit]
lab@R4# delete protocols isis level 1

[edit]
lab@R4# set protocols isis level 1 disable

[edit]
lab@R4# commit
commit complete
```

The router R4 adjacency with R7 looks like it cannot derive the system name from the LSP. However, the state of the adjacency is Up.

```
[edit]
lab@R4# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ae0.0                 R3               2  Up                23  0:5:86:49:ba:c0
ge-0/0/4.45           R5               2  Up                24
ge-0/0/4.46           R6               2  Up                19
ge-0/0/4.47           0172.0020.2507   2  Up                 8  2:6:a:e:7:4
```

m)Fixing the adjacency on R7.

© 2020 Juniper Networks, Inc. All rights reserved.

The adjacency table on router R7 provided an indication of the problem with the IS-IS communication between R4 and R3 with R7. The fact that the adjacencies are Up but the System name is not resolved, means the Hellos are successfully exchanged but not the LSPs between the neighbors. Enable traceoptions to find the reason.

```
[edit]
lab@R7# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/4.37           0172.0020.2503   2  Up                   25  2:6:a:e:3:4
ge-0/0/4.47           0172.0020.2504   2  Up                   21  2:6:a:e:4:4

[edit]
lab@R7# set protocols isis traceoptions file ISIS_log

[edit]
lab@R7# set protocols isis traceoptions flag error detail

[edit]
lab@R7# set protocols isis traceoptions flag lsp detail

[edit]
lab@R7# commit
commit complete

[edit]
lab@R7# run show log ISIS_log
Dec  4 10:24:10 trace_on: Tracing to "/var/log/ISIS_log" started

Dec  4 10:25:14.438952 Received L2 LSP 0172.0020.2506.00-00, on interface ge-0/0/4.37
Dec  4 10:25:14.439013     from 0172.0020.2503
Dec  4 10:25:14.439023     sequence 0xe, checksum 0, lifetime 0
Dec  4 10:25:14.439028     max area 0, length 45
Dec  4 10:25:14.439036     no partition repair, no database overload
Dec  4 10:25:14.439041     IS type 3, metric type 0
Dec  4 10:25:14.439049     16 bytes of authentication data
Dec  4 10:25:14.439128 ERROR: LSP authentication information with bad length 16
Dec  4 10:25:14.439681 ERROR: ISIS ignored a bad packet: L2 LSP
id 0172.0020.2506.00-00 from 0172.0020.2503 on interface ge-0/0/4.37 failed
authentication
Dec  4 10:25:14.443267 Received L2 LSP 0172.0020.2506.00-00, on interface ge-0/0/4.47
Dec  4 10:25:14.443275     from 0172.0020.2504
Dec  4 10:25:14.443278     sequence 0xe, checksum 0, lifetime 0
Dec  4 10:25:14.443282     max area 0, length 45
Dec  4 10:25:14.443286     no partition repair, no database overload
Dec  4 10:25:14.443290     IS type 3, metric type 0
Dec  4 10:25:14.443295     16 bytes of authentication data
Dec  4 10:25:14.443300 ERROR: LSP authentication information with bad length 16
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Dec  4 10:25:14.443304 Notification not sent due to 5-second throttle as per rfc4444:
isisAuthenticationFailure
Dec  4 10:25:14.443309 ERROR: ISIS ignored a bad packet: L2 LSP id 0172.0020.2506.00-00
from 0172.0020.2504 on interface ge-0/0/4.47 failed authentication
```

The logs prove that the LSP authentication is failing. The authentication in IS-IS protocol can be configured under the IS-IS level section and under the interfaces. Authentication applied under the interface affects only the Hello messages exchanged over that interface. Whereas, the authentication configured under the level authenticates both, the LSP and the Hello packets. If configured, the interface authentication is always used.

Fixing the authentication configuration under the IS-IS level will establish successfully the adjacencies.

Because simple authentication is used, looking into the LSP packets will reveal the password configured on both ends.

```
[edit]
lab@R7# run monitor traffic interface ge-0/0/4.37 no-resolve size 1500 detail
Address resolution is OFF.
Listening on ge-0/0/4.37, capture size 1500 bytes

10:22:32.073687 Out IS-IS, length 98
        L2 CSNP, hlen: 33, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
          source-id:    0172.0020.2507.00, PDU length: 98
          start lsp-id: 0000.0000.0000.00-00
          end lsp-id:   ffff.ffff.ffff.ff-ff
            LSP entries TLV #9, length: 48
              lsp-id: 0172.0020.2507.00-00, seq: 0x0000000e, lifetime:   732s, chksum:
0xa429
              lsp-id: 0172.0020.2507.02-00, seq: 0x0000000b, lifetime:   490s, chksum:
0xca87
              lsp-id: 0172.0020.2507.03-00, seq: 0x0000000b, lifetime:   732s, chksum:
0x7dd5
            Authentication TLV #10, length: 13
              simple text password: jncie-level2

10:22:21.636391  In IS-IS, length 45
        L2 LSP, hlen: 27, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
          lsp-id: 0172.0020.2504.02-00, seq: 0x00000036, lifetime:     0s
          chksum: 0x0000(unverified), PDU length: 45, Flags: [ L1L2 IS ]
            Authentication TLV #10, length: 16
              simple text password: jncie-sp-level2

10:22:21.652052  In IS-IS, length 168
        L2 LSP, hlen: 27, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lsp-id: 0172.0020.2503.00-00, seq: 0x00000040, lifetime:  1198s
chksum: 0x547c (correct), PDU length: 168, Flags: [ L1L2 IS ]
  Area address(es) TLV #1, length: 4
    Area address (length: 3): 49.0001
  LSP Buffersize TLV #14, length: 2
    LSP Buffersize: 1492
  Protocols supported TLV #129, length: 1
    NLPID(s): IPv6 (0x8e)
  Traffic Engineering Router ID TLV #134, length: 4
    Traffic Engineering Router ID: 172.20.25.3
  IPv4 Interface address(es) TLV #132, length: 4
    IPv4 interface address: 172.20.25.3
  Hostname TLV #137, length: 2
    Hostname: R3
  IS Reachability TLV #2, length: 23
    IsNotVirtual
    IS Neighbor: 0172.0020.2507.03, Default Metric: 10, Internal
    IS Neighbor: 0172.0020.2505.00, Default Metric: 10, Internal
  Extended IS Reachability TLV #22, length: 60
    IS Neighbor: 0172.0020.2507.03, Metric: 10, sub-TLVs present (16)
      IPv4 interface address subTLV #6, length: 4, 172.20.20.5
      Link Local/Remote Identifier subTLV #4, length: 8, 0x0000014d,
0x00000000
    IS Neighbor: 0172.0020.2505.00, Metric: 10, sub-TLVs present (22)
      IPv4 interface address subTLV #6, length: 4, 172.20.20.13
      IPv4 neighbor address subTLV #8, length: 4, 172.20.20.14
      Link Local/Remote Identifier subTLV #4, length: 8, 0x0000014c,
0x0000014d
  Router Capability TLV #242, length: 5
   Router ID:     172.20.25.3,Flags: 0x00000001
  Authentication TLV #10, length: 16
    simple text password: jncie-sp-level2

[edit]
lab@R7# set protocols isis level 2 authentication-key jncie-sp-level2

[edit]
lab@R7# delete protocols isis traceoptions

[edit]
lab@R7# commit
commit complete
```

Setting the correct password for Level 2 on R7 ensures that the adjacencies are successfully established.

```
[edit]
lab@R7# run show isis adjacency
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Interface              System         L State       Hold (secs) SNPA
ge-0/0/4.37            R3             2  Up                   21  2:6:a:e:3:4
ge-0/0/4.47            R4             2  Up                   23  2:6:a:e:4:4
```

n) Fixing the adjacency on R8.

R8 interfaces to R7 and R4 are listed in the IS-IS interface table but the DR field has only 0s for the two adjacencies. However, this is enough to conclude that the interfaces are configured under IS-IS protocol and are assigned with `family iso`. Next, you must check if the loopback interface is configured with a system ID.

```
[edit]
lab@R8# show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-re-ipv4;
        }
        address 172.20.25.8/32;
    }
    family inet6 {
        filter {
            input protect-re-ipv6;
        }
        address 2001:fedc::8/128;
    }
}
```

IS-IS requires a NET address configured under the loopback based on the Area ID and a unique system ID.

```
[edit]
lab@R8# set interfaces lo0.0 family iso address 49.0001.0172.0020.2508.00

[edit]
lab@R8# commit
commit complete
```

After adding the missing configuration, the adjacencies are successfully established.

```
[edit]
lab@R8# run show isis adjacency
Interface              System         L State       Hold (secs) SNPA
ge-0/0/4.48            R4             2  Up                   22  2:6:a:e:4:4
ge-0/0/4.78            R7             2  Up                   22  2:6:a:e:7:4
```


© 2020 Juniper Networks, Inc. All rights reserved.

c) Interfaces between area 49.0001 and 49.0005 should not elect DIS.

Similar to OSPF routing protocol, IS-IS protocol uses a method to limit the number of adjacencies and the traffic in broadcast segments. IS-IS elects one of the routes as Designated Intermediate System (DIS). Configuring point-to-point on the broadcast interface in IS-IS, disables the DIS election and eliminates the additional control traffic.

a) R3

```
[edit]
lab@R3# set protocols isis interface ge-0/0/4.35 point-to-point

[edit]
lab@R3# commit
commit complete
```

b) R4

```
[edit]
lab@R4# set protocols isis interface ge-0/0/4.45 point-to-point

[edit]
lab@R4# set protocols isis interface ge-0/0/4.46 point-to-point

[edit]
lab@R4# commit
commit complete
```

c) R5

```
[edit]
lab@R5# set protocols isis interface ge-0/0/4.35 point-to-point

[edit]
lab@R5# set protocols isis interface ge-0/0/4.45 point-to-point

[edit]
lab@R5# commit
commit complete
```

d) R6

```
[edit]
lab@R6# set protocols isis interface ge-0/0/4.46 point-to-point
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# commit
commit complete
```

e) Verify the configuration

The adjacency table on router R5 is missing the SNPA on the interfaces configured as point-to-point. This a clear sign that the DIS is not elected. SNPA represents the MAC address of the next-hop in LAN segments.

```
[edit]
lab@R5# run show isis adjacency
Interface             System         L State       Hold (secs) SNPA
ge-0/0/4.35           R3             2  Up                   20
ge-0/0/4.45           R4             2  Up                   19
ge-0/0/4.52           R2             1  Up                   25  2:6:a:e:2:4
ge-0/0/4.56           R6             1  Up                    8  2:6:a:e:6:4
```

d) Loopback interface addresses must be present as internal routes in the respected areas. R2 loopback should appear as an external address.

Although, the loopback interface is mandatory to be configured for the correct operation of IS-IS protocol, you still must explicitly advertise the IPv4 and IPv6 loopback addresses. If you configure the loopback interface under the `protocol isis` hierarchy directly, the loopback address is advertised as internal IS-IS route. Rather than using a routing policy to advertise the route, then it will appear as external IS-IS route.

a) All IS-IS routers except R2.

Passive option prevents the IS-IS protocol from bringing up adjacency over the interface.

```
[edit]
lab@R3# set protocols isis interface lo0.0 passive

[edit]
lab@R3# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) R2

The requirement to advertise R2 loopback needs the use of routing policy applied to the IS-IS protocol. Because IS-IS supports IPv6 routing by default, adding the loopback interface under IS-IS section in the previous step, will advertise the IPv6 address together with the IPv4. To achieve the same effect on router R2 you can match the interface name in the routing policy. This way both addresses are advertised in IS-IS as external routes.

```
[edit]
lab@R2# set policy-options policy-statement to.isis term loopback from interface lo0.0

[edit]
lab@R2# set policy-options policy-statement to.isis term loopback then accept

[edit]
lab@R2# set protocols isis export to.isis

[edit]
lab@R2# commit
commit complete
```

c) Verify the configuration.

IPv4 and IPv6 addresses are advertised in IS-IS to all routers.

R2

```
[edit]
lab@R2# run show route 172.20.25.0/28 table inet.0

inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.2/32     *[Direct/0] 16:46:24
                    > via lo0.0
172.20.25.5/32     *[IS-IS/15] 13:21:25, metric 10
                    > to 172.20.20.78 via ge-0/0/4.52
172.20.25.6/32     *[IS-IS/15] 13:21:18, metric 10
                    > to 172.20.20.82 via ge-0/0/4.62
```

```
[edit]
lab@R2# run show route 2001:fedc::/64
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet6.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2001:fedc::2/128   *[Direct/0] 17:21:47
                    > via lo0.0
2001:fedc::5/128   *[IS-IS/15] 13:56:48, metric 10
                    > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
2001:fedc::6/128   *[IS-IS/15] 13:56:41, metric 10
                    > to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
```

R3

```
[edit]
lab@R3# run show route 172.20.25.0/28 table inet.0

inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.3/32     *[Direct/0] 16:56:55
                    > via lo0.0
```

```
[edit]
lab@R3# run show route 2001:fedc::/64

inet6.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2001:fedc::3/128   *[Direct/0] 17:23:02
                    > via lo0.0
2001:fedc::4/128   *[IS-IS/18] 00:42:26, metric 10
                    > to fe80::2e6b:f5ff:fe54:61c1 via ae0.0
2001:fedc::5/128   *[IS-IS/18] 13:58:07, metric 10
                    > to fe80::250:5600:23a9:2abe via ge-0/0/4.35
2001:fedc::6/128   *[IS-IS/18] 00:42:26, metric 20
                    > to fe80::250:5600:23a9:2abe via ge-0/0/4.35
                      to fe80::2e6b:f5ff:fe54:61c1 via ae0.0
2001:fedc::7/128   *[IS-IS/18] 13:57:56, metric 10
                    > to fe80::250:5600:25a9:646c via ge-0/0/4.37
2001:fedc::8/128   *[IS-IS/18] 00:42:26, metric 20
                      to fe80::250:5600:25a9:646c via ge-0/0/4.37
                    > to fe80::2e6b:f5ff:fe54:61c1 via ae0.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

R4

```
[edit]
lab@R4# run show route 172.20.25.0/28 table inet.0

inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.4/32     *[Direct/0] 16:58:13
                    > via lo0.0
```

```
[edit]
lab@R4# run show route 2001:fedc::/64

inet6.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2001:fedc::3/128   *[IS-IS/18] 00:43:08, metric 10
                    > to fe80::2e6b:f5ff:fe10:1c1 via ae0.0
2001:fedc::4/128   *[Direct/0] 17:23:39
                    > via lo0.0
2001:fedc::5/128   *[IS-IS/18] 13:58:49, metric 10
                    > to fe80::250:5600:2da9:2abe via ge-0/0/4.45
2001:fedc::6/128   *[IS-IS/18] 13:58:42, metric 10
                    > to fe80::250:5600:2ea9:bd1 via ge-0/0/4.46
2001:fedc::7/128   *[IS-IS/18] 13:58:38, metric 10
                    > to fe80::250:5600:2fa9:646c via ge-0/0/4.47
2001:fedc::8/128   *[IS-IS/18] 13:58:33, metric 10
                    > to fe80::250:5600:30a9:775f via ge-0/0/4.48
```

The reason that R3 and R4 only see their local IPv4 loopback address is because IS-IS IPv4 routing is disabled on these routers in the initial configuration. Deleting the command and reverting to the default IS-IS operation solves the Task.

R3

```
[edit]
lab@R3# delete protocols isis no-ipv4-routing

[edit]
lab@R3# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

R4

```
[edit]
lab@R4# delete protocols isis no-ipv4-routing

[edit]
lab@R4# commit
commit complete
```

Review the routing tables after making the change to verify that the additional loopback addresses are learned.

R3

```
[edit]
lab@R3# run show route 172.20.25.0/28 table inet.0

inet.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.3/32     *[Direct/0] 17:20:15
                    > via lo0.0
172.20.25.4/32     *[IS-IS/18] 00:00:17, metric 10
                    > to 172.20.20.2 via ae0.0
172.20.25.5/32     *[IS-IS/18] 00:00:22, metric 10
                    > to 172.20.20.14 via ge-0/0/4.35
172.20.25.6/32     *[IS-IS/18] 00:00:17, metric 20
                    > to 172.20.20.14 via ge-0/0/4.35
                      to 172.20.20.2 via ae0.0
172.20.25.7/32     *[IS-IS/18] 00:00:22, metric 10
                    > to 172.20.20.6 via ge-0/0/4.37
172.20.25.8/32     *[IS-IS/18] 00:00:17, metric 20
                      to 172.20.20.6 via ge-0/0/4.37
                    > to 172.20.20.2 via ae0.0
```

R4

```
[edit]
lab@R4# run show route 172.20.25.0/28 table inet.0

inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.20.25.3/32     *[IS-IS/18] 00:01:12, metric 10
                    > to 172.20.20.1 via ae0.0
172.20.25.4/32     *[Direct/0] 17:21:06
                    > via lo0.0
172.20.25.5/32     *[IS-IS/18] 00:01:12, metric 10
                    > to 172.20.20.18 via ge-0/0/4.45
172.20.25.6/32     *[IS-IS/18] 00:01:12, metric 10
                    > to 172.20.20.22 via ge-0/0/4.46
172.20.25.7/32     *[IS-IS/18] 00:01:12, metric 10
                    > to 172.20.20.10 via ge-0/0/4.47
172.20.25.8/32     *[IS-IS/18] 00:01:12, metric 10
                    > to 172.20.20.26 via ge-0/0/4.48
```

2) Configure R1 to exchange successfully OSPF routing information with routers R5 and R6. Make sure that IPv4 and IPv6 loopback IP addresses between the OSPF domain and IS-IS domains are mutually reachable.

Router R1 must be connected to the network using the OSPFv3 routing protocol. The loopback IPv4 and IPv6 addresses of IS-IS routers and OSPF domain should be mutually exchanged.

a) R1

Configuring router R1 for OSPFv3 for IPv4 and IPv6 routing. The loopback interface is added as under the OSPF process.

```
[edit]
lab@R1# set protocols ospf3 realm ipv4-unicast area 0.0.0.10 interface lo0.0 passive

[edit]
lab@R1# set protocols ospf3 realm ipv4-unicast area 0.0.0.10 interface ge-0/0/4.16
interface-type p2p

[edit]
lab@R1# set protocols ospf3 realm ipv4-unicast area 0.0.0.10 interface ge-0/0/4.15
interface-type p2p

[edit]
lab@R1# set protocols ospf3 area 0.0.0.10 interface ge-0/0/4.16 interface-type p2p

[edit]
lab@R1# set protocols ospf3 area 0.0.0.10 interface lo0.0 passive

[edit]
lab@R1# set protocols ospf3 area 0.0.0.10 interface ge-0/0/4.15 interface-type p2p
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

b) R5

In addition to enabling OSPFv3 routing protocol on routers R5 and R6, routing policy is needed for redistribution of the loopback IP addresses from IS-IS to OSPF.

The routing policy is advertising IPv4 loopback addresses from protocol IS-IS to OSPF.

```
[edit]
lab@R5# set protocols ospf3 realm ipv4-unicast export to.ospf3.ipv4

[edit]
lab@R5# set protocols ospf3 realm ipv4-unicast area 0.0.0.10 interface ge-0/0/4.15
interface-type p2p


[edit]
lab@R5# set protocols ospf3 export to.ospf3.ipv6

[edit]
lab@R5# set protocols ospf3 area 0.0.0.10 interface ge-0/0/4.15 interface-type p2p
```

Amending the external preference of IS-IS routes lower than the OSPF external preference is needed to prevent suboptimal routing.

```
[edit]
lab@R5# set protocols isis level 1 external-preference 140
```

The routing policy is advertising only the IS-IS loopback routes to the OSPF domain.

```
[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term isis.loopback.routes
from protocol isis

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term isis.loopback.routes
from route-filter 172.20.25.0/24 prefix-length-range /32-/32

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term isis.loopback.routes
then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term local.loopback from
protocol direct


[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term local.loopback from
route-filter 172.20.25.0/24 prefix-length-range /32-/32

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term local.loopback then
accept

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term isis.loopback.routes
from protocol isis

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term isis.loopback.routes
from route-filter 2001:fedc::/64 prefix-length-range /128-/128

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term isis.loopback.routes
then accept

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term local.loopback from
protocol direct

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term local.loopback from
route-filter 2001:fedc::/64 prefix-length-range /128-/128

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term local.loopback then
accept

[edit]
lab@R5# commit
commit complete
```

c) R6

Similar configuration is applied to router R6.

```
[edit]
lab@R6# set protocols ospf3 realm ipv4-unicast export to.ospf3.ipv4
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols ospf3 realm ipv4-unicast area 0.0.0.10 interface ge-0/0/4.16
interface-type p2p

[edit]
lab@R6# set protocols ospf3 export to.ospf3.ipv6

[edit]
lab@R6# set protocols ospf3 area 0.0.0.10 interface ge-0/0/4.16 interface-type p2p

[edit]
lab@R6# set protocols isis level 1 external-preference 140

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term isis.loopback.routes
from protocol isis

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term isis.loopback.routes
from route-filter 172.20.25.0/24 prefix-length-range /32-/32

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term isis.loopback.routes
then accept

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term local.loopback from
protocol direct

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term local.loopback from
route-filter 172.20.25.0/24 prefix-length-range /32-/32

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term local.loopback then
accept

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term isis.loopback.routes
from protocol isis

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term isis.loopback.routes
from route-filter 2001:fedc::/64 prefix-length-range /128-/128

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term isis.loopback.routes
then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term local.loopback from
protocol direct

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term local.loopback from
route-filter 2001:fedc::/64 prefix-length-range /128-/128

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term local.loopback then
accept

[edit]
lab@R6# commit
commit complete
```

d) Verify the configuration.

OSPFv3 neighbor table on router R1 shows the established adjacency using the IPv6 link local addresses.

```
[edit]
lab@R1# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
172.20.25.5     ge-0/0/4.15           Full      128     36
  Neighbor-address fe80::206:a00:f0e:504
172.20.25.6     ge-0/0/4.16           Full      128     35
  Neighbor-address fe80::206:a00:100e:604
```

Router R1 successfully receives the routes advertised from routers R5 and R6.

```
[edit]
lab@R1# run show route 172.20.25.0/28 table inet.0

inet.0: 526 destinations, 1018 routes (526 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.0/28     *[BGP/170] 00:21:54, localpref 100, from 172.20.25.5
                      AS path: (65001) I, validation-state: unverified
                    > to 172.20.20.66 via ge-0/0/4.15
                    [BGP/170] 00:21:54, localpref 100, from 172.20.25.6
                      AS path: (65001) I, validation-state: unverified
                    > to 172.20.20.66 via ge-0/0/4.15
                      to 172.20.20.70 via ge-0/0/4.16
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.20.25.1/32     *[Direct/0] 23:17:45
                    > via lo0.0
172.20.25.2/32     *[OSPF3/150] 14:22:47, metric 20, tag 0
                    > to 172.20.20.70 via ge-0/0/4.16
172.20.25.3/32     *[OSPF3/150] 00:21:54, metric 10, tag 0
                    > to 172.20.20.66 via ge-0/0/4.15
172.20.25.4/32     *[OSPF3/150] 00:18:09, metric 10, tag 0
                    > to 172.20.20.66 via ge-0/0/4.15
                      to 172.20.20.70 via ge-0/0/4.16
172.20.25.5/32     *[OSPF3/150] 16:08:12, metric 0, tag 0
                    > to 172.20.20.66 via ge-0/0/4.15
172.20.25.6/32     *[OSPF3/150] 16:07:37, metric 0, tag 0
                    > to 172.20.20.70 via ge-0/0/4.16
172.20.25.7/32     *[OSPF3/150] 00:18:09, metric 20, tag 0
                    > to 172.20.20.66 via ge-0/0/4.15
                      to 172.20.20.70 via ge-0/0/4.16
172.20.25.8/32     *[OSPF3/150] 00:18:09, metric 20, tag 0
                      to 172.20.20.66 via ge-0/0/4.15
                    > to 172.20.20.70 via ge-0/0/4.16
172.20.25.10/32    *[OSPF3/150] 00:18:09, metric 20, tag 0
                      to 172.20.20.66 via ge-0/0/4.15
                    > to 172.20.20.70 via ge-0/0/4.16
```

OSPFv3 neighbor table on router R5.

```
[edit]
lab@R5# run show ospf3 neighbor
ID              Interface             State     Pri   Dead
172.20.25.1     ge-0/0/4.15           Full      128     36
  Neighbor-address fe80::206:a00:f0e:104
```

```
[edit]
lab@R5# run show route 172.20.25.1/32

inet.0: 28 destinations, 32 routes (28 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.1/32     *[OSPF3/10] 00:31:48, metric 1
                    > to 172.20.20.65 via ge-0/0/4.15
```

The reason for lowering the IS-IS external preference is clearly visible with the output below. Failing to do so, would prevent router R5 and R6 to advertise the R2 loopback IP address.

```
[edit]
lab@R5# run show route 172.20.25.2/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 558 destinations, 635 routes (558 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.2/32     *[IS-IS/140] 00:12:38, metric 30
                     > to 172.20.20.74 via ge-0/0/4.56
                    [OSPF3/150] 01:01:03, metric 20, tag 0
                     > to 172.20.20.65 via ge-0/0/4.15
```

```
[edit]
lab@R6# run show route 172.20.25.1/32

inet.0: 558 destinations, 634 routes (558 active, 0 holddown, 0 hidden)
Restart Complete
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both
172.20.25.1/32     *[OSPF3/10] 00:26:16, metric 1
                     > to 172.20.20.69 via ge-0/0/4.16
```

```
[edit]
lab@R6# run show route 172.20.25.2/32

inet.0: 558 destinations, 634 routes (558 active, 0 holddown, 0 hidden)
Restart Complete
@ = Routing Use Only, # = Forwarding Use Only
+ = Active Route, - = Last Active, * = Both
172.20.25.2/32     *[IS-IS/140] 00:27:44, metric 20
                     > to 172.20.20.81 via ge-0/0/4.62
                    [OSPF3/150] 00:10:40, metric 30, tag 0
                     > to 172.20.20.69 via ge-0/0/4.16
```

The output below shows that the IPv6 addresses are advertised too.

```
[edit]
lab@R1# run show route 2001:fedc::/64

inet6.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2001:fedc::1/128   *[Direct/0] 23:21:29
                     > via lo0.0
2001:fedc::2/128   *[OSPF3/150] 14:33:19, metric 20, tag 0
                     > to fe80::250:5600:fa9:2abe via ge-0/0/4.15
2001:fedc::3/128   *[OSPF3/150] 00:25:38, metric 10, tag 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to fe80::250:5600:fa9:2abe via ge-0/0/4.15
2001:fedc::4/128    *[OSPF3/150] 00:21:53, metric 10, tag 0
                    > to fe80::250:5600:fa9:2abe via ge-0/0/4.15
                      to fe80::250:5600:10a9:bd1 via ge-0/0/4.16
2001:fedc::5/128    *[OSPF3/150] 16:11:56, metric 0, tag 0
                    > to fe80::250:5600:fa9:2abe via ge-0/0/4.15
2001:fedc::6/128    *[OSPF3/150] 16:11:21, metric 0, tag 0
                    > to fe80::250:5600:10a9:bd1 via ge-0/0/4.16
2001:fedc::7/128    *[OSPF3/150] 00:25:38, metric 20, tag 0
                    > to fe80::250:5600:fa9:2abe via ge-0/0/4.15
2001:fedc::8/128    *[OSPF3/150] 00:21:53, metric 20, tag 0
                      to fe80::250:5600:fa9:2abe via ge-0/0/4.15
                    > to fe80::250:5600:10a9:bd1 via ge-0/0/4.16
```

3) Advertise a default route to R1 so that it can reach the rest of the network. Ensure that a default route is advertised only when the interfaces to Level 2 area are operational. In such a way, that if the R6 interface to R4 fails the default route is not advertised anymore to R1.

a) R5

The task requires advertising a default route to the OSPF domain, but only if the links to the Level 2 area are operational. To solve this, you must use a route policy condition.

The route policy condition is tracking the routers R3 and R4 loopback addresses. If those addresses are not present in the table, the condition will be false.

```
[edit]
lab@R5# set policy-options policy-statement ipv4-condition term 1 from route-filter
172.20.25.3/32 exact

[edit]
lab@R5# set policy-options policy-statement ipv4-condition term 1 from route-filter
172.20.25.4/32 exact

[edit]
lab@R5# set policy-options policy-statement ipv4-condition term 1 then accept

[edit]
lab@R5# set policy-options policy-statement ipv4-condition term 2 then reject
```

An additional routing policy condition is configured for tracking the IPv6 loopback addresses.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set policy-options policy-statement ipv6-condition term 1 from route-filter
2001:fedc::3/128 exact

[edit]
lab@R5# set policy-options policy-statement ipv6-condition term 1 from route-filter
2001:fedc::4/128 exact

[edit]
lab@R5# set policy-options policy-statement ipv6-condition term 1 then accept

[edit]
lab@R5# set policy-options policy-statement ipv6-condition term 2 then reject
```

The routing policy condition is configured as a condition for the aggregate routes.

```
[edit]
lab@R5# set routing-options rib inet6.0 aggregate route ::/0 policy ipv6-condition

[edit]
lab@R5# set routing-options aggregate route 0.0.0.0/0 policy ipv4-condition
```

Adding the aggregate route to the already applied routing policy for OSPF.

```
[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term ip4.default from
protocol aggregate

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term ip4.default from route-
filter 0.0.0.0/0 exact

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term ip4.default then accept

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term 1 from protocol
aggregate

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term 1 from route-filter ::/0
exact

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term 1 then accept

[edit]
lab@R5# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) R6

Similar configuration is applied to router R6.

```
[edit]
lab@R6# set policy-options policy-statement ipv4-condition term 1 from route-filter
172.20.25.3/32 exact

[edit]
lab@R6# set policy-options policy-statement ipv4-condition term 1 from route-filter
172.20.25.4/32 exact

[edit]
lab@R6# set policy-options policy-statement ipv4-condition term 1 then accept

[edit]
lab@R6# set policy-options policy-statement ipv4-condition term 2 then reject

[edit]
lab@R6# set policy-options policy-statement ipv6-condition term 1 from route-filter
2001:fedc::3/128 exact

[edit]
lab@R6# set policy-options policy-statement ipv6-condition term 1 from route-filter
2001:fedc::4/128 exact

[edit]
lab@R6# set policy-options policy-statement ipv6-condition term 1 then accept

[edit]
lab@R6# set policy-options policy-statement ipv6-condition term 2 then reject

[edit]
lab@R6# set routing-options rib inet6.0 aggregate route ::/0 policy ipv6-condition

[edit]
lab@R6# set routing-options aggregate route 0.0.0.0/0 policy ipv4-condition

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term ip4.default from
protocol aggregate

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term ip4.default from route-
filter 0.0.0.0/0 exact
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term ip4.default then accept

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term 1 from protocol
aggregate

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term 1 from route-filter ::/0
exact

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term 1 then accept
```

```
[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term 1 from protocol
aggregate

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term 1 from route-filter ::/0
exact

[edit]
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term 1 then accept

[edit]
lab@R6# commit
commit complete
```

c) Verify the configuration

The IPv4 default route is received by R1 from routers R5 and R6.

```
[edit]
lab@R1# run show route table inet.0 0/0 exact

inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF3/150] 00:01:16, metric 0, tag 0
                    > to 172.20.20.66 via ge-0/0/4.15
                      to 172.20.20.70 via ge-0/0/4.16
```

The IPv6 default route is received by R1 from routers R5 and R6.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show route table inet6.0 ::/0 exact

inet6.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

::/0                *[OSPF3/150] 00:52:54, metric 0, tag 0
                     > to fe80::250:5600:fa9:2abe via ge-0/0/4.15
                       to fe80::250:5600:10a9:bd1 via ge-0/0/4.16
```

4) Configure R7 and R8 to exchange RIP and RIPng routes from DC-1. A summary route representing all IPv4 and IPv6 loopback addresses should be advertised to DC-1 only when the connection to R3 and R4 is present. You are not allowed to amend the `routing-options` hierarchy on R7 and R8 to accomplish this task.

a) R3

Next, you must advertise a summary routes of IPv4 and IPv6 loopback addresses to the DC-1 router. However, the task requires those summaries to be advertised only if routers R7 and R8 have connectivity to either R3 and R4. You are not allowed to change the `routing-options` section on routers R7 and R8. This means you must ensure that the summary routes are received on R7 and R8 from other point in the network. Routers R3 and R4 are a good location from where the summary routes can be advertised. This will align with the requirement for routers R3 and R4 presence also.

```
[edit]
lab@R3# set policy-options policy-statement to.isis term aggregate.level2 from protocol
aggregate

[edit]
lab@R3# set policy-options policy-statement to.isis term aggregate.level2 from route-
filter 172.20.25.0/28 exact

[edit]
lab@R3# set policy-options policy-statement to.isis term aggregate.level2 to level 2


[edit]
lab@R3# set policy-options policy-statement to.isis term aggregate.level2 then accept

[edit]
lab@R3# set policy-options policy-statement to.isis term aggregate.level2.ipv6 from
protocol aggregate
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set policy-options policy-statement to.isis term aggregate.level2.ipv6 from
route-filter 2001:fedc::/64 exact

[edit]
lab@R3# set policy-options policy-statement to.isis term aggregate.level2.ipv6 to level
2

[edit]
lab@R3# set policy-options policy-statement to.isis term aggregate.level2.ipv6 then
accept

[edit]
lab@R3# set routing-options rib inet6.0 aggregate route 2001:fedc::/64

[edit]
lab@R3# set routing-options aggregate route 172.20.25.0/28


[edit]
lab@R3# set protocols isis export to.isis

[edit]
lab@R3# commit
commit complete
```

b) R4

```
[edit]
lab@R4# set policy-options policy-statement to.isis term aggregate.level2 from protocol
aggregate

[edit]
lab@R4# set policy-options policy-statement to.isis term aggregate.level2 from route-
filter 172.20.25.0/28 exact


[edit]
lab@R4# set policy-options policy-statement to.isis term aggregate.level2 to level 2

[edit]
lab@R4# set policy-options policy-statement to.isis term aggregate.level2 then accept

[edit]
lab@R4# set policy-options policy-statement to.isis term aggregate.level2.ipv6 from
protocol aggregate

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# set policy-options policy-statement to.isis term aggregate.level2.ipv6 from
route-filter 2001:fedc::/64 exact

[edit]
lab@R4# set policy-options policy-statement to.isis term aggregate.level2.ipv6 to level
2

[edit]
lab@R4# set policy-options policy-statement to.isis term aggregate.level2.ipv6 then
accept

[edit]
lab@R4# set routing-options rib inet6.0 aggregate route 2001:fedc::/64

[edit]
lab@R4# set routing-options aggregate route 172.20.25.0/28


[edit]
lab@R4# set protocols isis export to.isis

[edit]
lab@R4# commit
commit complete
```

c) R7

Configuring the interface for connectivity with DC-1 router. The IPv6 family is applied by the apply-group from the first part. The IPv6 link local addresses are needed for the RIPng connectivity.

```
[edit]
lab@R7# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R7# set interfaces ge-0/0/5 unit 201 description "=== connection to DC-1 ==="

[edit]
lab@R7# set interfaces ge-0/0/5 unit 201 vlan-id 201

[edit]
lab@R7# set interfaces ge-0/0/5 unit 201 family inet address 172.20.201.7/24
```

Applying import policy to the RIP protocol.

```
[edit]
lab@R7# set protocols rip import from.rip
```


© 2020 Juniper Networks, Inc. All rights reserved.

Applying export policy to the RIP protocol.

```
[edit]
lab@R7# set protocols rip group DC-1 export to.rip

[edit]
lab@R7# set protocols rip group DC-1 neighbor ge-0/0/5.201
```

Applying import policy to the RIPng protocol.

```
[edit]
lab@R7# set protocols ripng import from.ripng
```

Applying export policy to the RIP protocol.

```
[edit]
lab@R7# set protocols ripng group DC-1 export to.ripng

[edit]
lab@R7# set protocols ripng group DC-1 neighbor ge-0/0/5.201
```

The import policy applied RIP rejects the readvertised internal summary route in the first term to prevent suboptimal routing.

```
[edit]
lab@R7# set policy-options policy-statement from.rip term aggregate.reject from
protocol rip

[edit]
lab@R7# set policy-options policy-statement from.rip term aggregate.reject from route-
filter 172.20.25.0/28 exact


[edit]
lab@R7# set policy-options policy-statement from.rip term aggregate.reject then reject

[edit]
lab@R7# set policy-options policy-statement from.rip term accept.rest from protocol rip

[edit]
lab@R7# set policy-options policy-statement from.rip term accept.rest then accept
```

Similar approach is used to create the import policy for IPv6 RIPng.

```
[edit]
lab@R7# set policy-options policy-statement from.ripng term aggregate.reject from
protocol ripng
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set policy-options policy-statement from.ripng term aggregate.reject from
route-filter 2001:fedc::/64 exact

[edit]
lab@R7# set policy-options policy-statement from.ripng term aggregate.reject then
reject

[edit]
lab@R7# set policy-options policy-statement from.ripng term accept.rest from protocol
ripng

[edit]
lab@R7# set policy-options policy-statement from.ripng term accept.rest then accept

[edit]
lab@R7# set policy-options policy-statement from.ripng term accept.rest then accept
```

The export policy matches the aggregate route advertised by routers R3 and R4.

```
[edit]
lab@R7# set policy-options policy-statement to.rip term aggregate from protocol isis


[edit]
lab@R7# set policy-options policy-statement to.rip term aggregate from route-filter
172.20.25.0/28 exact

[edit]
lab@R7# set policy-options policy-statement to.rip term aggregate then accept

[edit]
lab@R7# set policy-options policy-statement to.ripng term aggregate.ipv6 from protocol
isis route-filter 2001:fedc::/64 exact

[edit]
lab@R7# set policy-options policy-statement to.ripng term aggregate.ipv6 then accept
```

The routing preference of the RIP protocol is lowered to make the received routes always active on routers R7 and R8.

```
[edit]
lab@R7# set protocols rip group DC-1 preference 17
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set protocols ripng group DC-1 preference 17

[edit]
lab@R7# commit
commit complete
```

d) R8

Similar configuration is applied to router R8.

```
[edit]
lab@R8# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R8# set interfaces ge-0/0/5 unit 202 description "=== connection to DC-1 ==="

[edit]
lab@R8# set interfaces ge-0/0/5 unit 202 vlan-id 202

[edit]
lab@R8# set interfaces ge-0/0/5 unit 202 family inet address 172.20.202.8/24

[edit]
lab@R8# set protocols rip import from.rip

[edit]
lab@R8# set protocols rip group DC-1 export to.rip

[edit]
lab@R8# set protocols rip group DC-1 neighbor ge-0/0/5.202

[edit]
lab@R8# set protocols ripng import from.ripng

[edit]
lab@R8# set protocols ripng group DC-1 export to.ripng

[edit]
lab@R8# set protocols ripng group DC-1 neighbor ge-0/0/5.202

[edit]
lab@R8# set policy-options policy-statement from.rip term aggregate.reject from
protocol rip

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# set policy-options policy-statement from.rip term aggregate.reject from route-
filter 172.20.25.0/28 exact

[edit]
lab@R8# set policy-options policy-statement from.rip term aggregate.reject then reject

[edit]
lab@R8# set policy-options policy-statement from.rip term accept.rest from protocol rip

[edit]
lab@R8# set policy-options policy-statement from.rip term accept.rest then accept

[edit]
lab@R8# set policy-options policy-statement from.ripng term aggregate.reject from
protocol ripng

[edit]
lab@R8# set policy-options policy-statement from.ripng term aggregate.reject from
route-filter 2001:fedc::/64 exact

[edit]
lab@R8# set policy-options policy-statement from.ripng term aggregate.reject then
reject

[edit]
lab@R8# set policy-options policy-statement from.ripng term accept.rest from protocol
ripng

[edit]
lab@R8# set policy-options policy-statement from.ripng term accept.rest then accept

[edit]
lab@R8# set policy-options policy-statement from.ripng term accept.rest then accept


[edit]
lab@R8# set policy-options policy-statement to.rip term aggregate from protocol isis

[edit]
lab@R8# set policy-options policy-statement to.rip term aggregate from route-filter
172.20.25.0/28 exact

[edit]
lab@R8# set policy-options policy-statement to.rip term aggregate then accept

[edit]
lab@R8# set policy-options policy-statement to.ripng term aggregate.ipv6 from protocol
isis route-filter 2001:fedc::/64 exact
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set policy-options policy-statement to.ripng term aggregate.ipv6 then accept

[edit]
lab@R8# set protocols rip group DC-1 preference 17

[edit]
lab@R8# set protocols ripng group DC-1 preference 17

[edit]
lab@R8# commit
commit complete
```

e) Verify the configuration.

The RIP neighbor is connected successfully, and routes are received on routers R7 and R8.

```
[edit]
lab@R7# run show rip neighbor
                    Local  Source          Destination     Send   Receive   In
Neighbor            State  Address         Address         Mode   Mode     Met
--------            -----  -------         -----------     ----   -------  ---
ge-0/0/5.201           Up 172.20.201.7    224.0.0.9       mcast  both       1
```

```
[edit]
lab@R7# run show route protocol rip

inet.0: 42 destinations, 42 routes (42 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.15.0.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.1.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.2.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.3.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.4.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.5.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.6.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
10.15.7.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.8.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.9.0/24       *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.10.0/24      *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.11.0/24      *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.12.0/24      *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.13.0/24      *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.14.0/24      *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
10.15.15.0/24      *[RIP/100] 01:09:29, metric 2, tag 100
                    > to 172.20.201.201 via ge-0/0/5.201
224.0.0.9/32       *[RIP/100] 01:09:31, metric 1
                      MultiRecv
```

The RIP neighbor is connected successfully, and routes are received by routers R7 and R8.

```
[edit]
lab@R7# run show ripng neighbor
                     Source                                  Dest                 In
Neighbor     State   Address                                 Address   Send Recv  Met
--------     -----   -------                                 -------   ---- ----  ---
ge-0/0/5.201    Up fe80::250:5600:c9a9:44ce                 ff02::9   yes yes     1

[edit]
lab@R7# run show route protocol ripng

inet.0: 41 destinations, 41 routes (41 active, 0 holddown, 0 hidden)
Restart Complete

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

inet6.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2015:4321::/64     *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:1::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
2015:4321:2::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:3::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:4::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:5::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:6::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:7::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:8::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:9::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:a::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:b::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:c::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:d::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:e::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
2015:4321:f::/64   *[RIPng/100] 01:10:00, metric 2, tag 0
                    > to fe80::206:a00:c90e:908 via ge-0/0/5.201
ff02::9/128        *[RIPng/100] 00:03:27, metric 1
                      MultiRecv
```

The IPv4 and IPv6 summary routes are advertised to RIP and RIPng protocols.

Note: Because the interfaces are using the link local address, you must get your neighbor address for the next couple commands based on the output you received when validating the neighbor session above.

```
[edit]
lab@R7# run show route advertising-protocol rip 172.20.201.7

inet.0: 42 destinations, 42 routes (42 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.0/28     *[IS-IS/165] 00:02:48, metric 20
                    > to 172.20.20.5 via ge-0/0/4.37
                      to 172.20.20.9 via ge-0/0/4.47
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show route advertising-protocol ripng fe80::206:a00:c90e:705

inet6.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2001:fedc::/64     *[IS-IS/165] 00:05:10, metric 20
                    > to fe80::250:5600:25a9:3002 via ge-0/0/4.37
                      to fe80::250:5600:2fa9:7243 via ge-0/0/4.47
```

5) The routes received from DC-1 must be present in all IS-IS enabled routers individually.

a) R5

The task requires all RIP routes to be present in the routing tables of all IS-IS routers. However, those routes are missing from the IS-IS Level 1 routers.

```
[edit]
lab@R5# run show route 10.15.7.0/24
```

b) R3 and R4

The reason for this is that we do not have a policy to redistribute these RIP routes into IS-IS from R7 and R8.

c) R7

```
[edit]
lab@R7# set policy-options policy-statement to.isis term rip.routes from protocol rip

[edit]
lab@R7# set policy-options policy-statement to.isis term rip.routes then accept

[edit]
lab@R7# set policy-options policy-statement to.isis term ripng.routes from protocol
ripng

[edit]
lab@R7# set policy-options policy-statement to.isis term ripng.routes then accept

[edit]
lab@R7# set protocols isis export to.isis
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# commit
commit complete
```

d) R8

```
 [edit]
lab@R8# set policy-options policy-statement to.isis term rip.routes from protocol rip

[edit]
lab@R8# set policy-options policy-statement to.isis term rip.routes then accept
```

```
[edit]
lab@R8# set policy-options policy-statement to.isis term ripng.routes from protocol
ripng

[edit]
lab@R8# set policy-options policy-statement to.isis term ripng.routes then accept
```

```
[edit]
lab@R8# set protocols isis export to.isis

[edit]
lab@R8# commit
commit complete
```

e) Verify the configuration.

After committing the change, the IPv4 and IPv6 RIP routes are sent to IS-IS Level 1 routers.

```
[edit]
lab@R5# run show route 10.15.7.0/24

inet.0: 46 destinations, 50 routes (46 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.15.7.0/24        *[IS-IS/165] 00:04:50, metric 22
                     > to 172.20.20.13 via ge-0/0/4.35
                       to 172.20.20.17 via ge-0/0/4.45
```

```
[edit]
lab@R5# run show route 2015:4321::/64

inet6.0: 34 destinations, 38 routes (34 active, 0 holddown, 0 hidden)
```


© 2020 Juniper Networks, Inc. All rights reserved.

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

2015:4321::/64     *[IS-IS/165] 00:05:13, metric 30
                    > to fe80::250:5600:23a9:3002 via ge-0/0/4.35
                      to fe80::250:5600:2da9:7243 via ge-0/0/4.45
```

6) Router R7 must be the gateway by default for the traffic to and from the DC-1 data center.

a) R8

To complete this task, you can increase the metric on router R8 for the RIP routes redistributed to and from IS-IS. To do that you have to modify your redistribution policy into IS-IS.

```
[edit]
lab@R8# set policy-options policy-statement to.isis term rip.routes then metric 2000

[edit]
lab@R8# set policy-options policy-statement to.isis term ripng.routes then metric 2000
```

RIP protocol allows setting the metric for all routes sent to a certain neighbor.

```
[edit]
lab@R8# set protocols rip group DC-1 metric-out 10

[edit]
lab@R8# set protocols ripng group DC-1 metric-out 10

[edit]
lab@R8# commit
commit complete
```

b) Verify the configuration.

After this change, the summary route advertised by R8 to DC-1 is with a higher metric.

```
[edit]
lab@R8# run show route advertising-protocol rip 172.20.202.8

inet.0: 41 destinations, 41 routes (41 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.0/28     *[IS-IS/18] 01:31:35, metric 20
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    > to 172.20.20.25 via ge-0/0/4.48
```

7) Make sure that routers R5 and R6 advertise only the Level 1 loopback addresses to Level 2 IS-IS network together with the R1 loopback address.
8) Router R2 must have all individual RIP and RIPng prefixes and all the loopback addresses of the rest of the routers. Limit the possibility of R2 to send traffic to destinations not explicitly present in the routing table.

a) R5

The next two Tasks require the configuration of routing policies to control the routing information between IS-IS Level 1 and Level 2.

An IS-IS export routing policy is configured on routers R5 and R6.

```
[edit]
lab@R5# set policy-options policy-statement to.isis term 1 from protocol ospf3

[edit]
lab@R5# set policy-options policy-statement to.isis term 1 from route-filter
172.20.25.1/32 exact

[edit]
lab@R5# set policy-options policy-statement to.isis term 1 then accept


[edit]
lab@R5# set policy-options policy-statement to.isis term 2 from protocol ospf3

[edit]
lab@R5# set policy-options policy-statement to.isis term 2 from route-filter
2001:fedc::1/128 exact

[edit]
lab@R5# set policy-options policy-statement to.isis term 2 then accept
```

Only the loopback IP addresses are matched by term 3. Avoid specifying the protocol ensures that the direct and IS-IS routes are matched. Actually, that term also eliminates the need of term 2, which was configured in one of the previous tasks.

```
[edit]
lab@R5# set policy-options policy-statement to.isis term 3 from route-filter
172.20.25.0/24 prefix-length-range /32-/32

[edit]
lab@R5# set policy-options policy-statement to.isis term 3 then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement to.isis term 4 from route-filter
2001:fedc::/64 prefix-length-range /128-/128

[edit]
lab@R5# set policy-options policy-statement to.isis term 4 then accept
```

The term below is named `rip.routes.to.level1` although it matches the IS-IS Level 2 routes. Since the RIP routes are advertised in IS-IS by routers R7 and R8, they become IS-IS external routes.

```
[edit]
lab@R5# set policy-options policy-statement to.isis term rip.routes.to.level1 from
protocol isis

[edit]
lab@R5# set policy-options policy-statement to.isis term rip.routes.to.level1 from
level 2


[edit]
lab@R5# set policy-options policy-statement to.isis term rip.routes.to.level1 to level
1

[edit]
lab@R5# set policy-options policy-statement to.isis term rip.routes.to.level1 then
accept

[edit]
lab@R5# set policy-options policy-statement to.isis term explicit.reject then reject

[edit]
lab@R5# set protocols isis export to.isis

[edit]
lab@R5# commit
commit complete
```

b) R6

The same configuration is applied to router R6.

```
[edit]
lab@R6# set policy-options policy-statement to.isis term 1 from protocol ospf3

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R6# set policy-options policy-statement to.isis term 1 from route-filter
172.20.25.1/32 exact

[edit]
lab@R6# set policy-options policy-statement to.isis term 1 then accept

[edit]
lab@R6# set policy-options policy-statement to.isis term 2 from protocol ospf3

[edit]
lab@R6# set policy-options policy-statement to.isis term 2 from route-filter
2001:fedc::1/128 exact

[edit]
lab@R6# set policy-options policy-statement to.isis term 2 then accept

[edit]
lab@R6# set policy-options policy-statement to.isis term 3 from route-filter
172.20.25.0/24 prefix-length-range /32-/32


[edit]
lab@R6# set policy-options policy-statement to.isis term 3 then accept

[edit]
lab@R6# set policy-options policy-statement to.isis term 4 from route-filter
2001:fedc::/64 prefix-length-range /128-/128

[edit]
lab@R6# set policy-options policy-statement to.isis term 4 then accept

[edit]
lab@R6# set policy-options policy-statement to.isis term rip.routes.to.level1 from
protocol isis

[edit]
lab@R6# set policy-options policy-statement to.isis term rip.routes.to.level1 from
level 2

[edit]
lab@R6# set policy-options policy-statement to.isis term rip.routes.to.level1 to level
1

[edit]
lab@R6# set policy-options policy-statement to.isis term rip.routes.to.level1 then
accept

[edit]
lab@R6# set policy-options policy-statement to.isis term explicit.reject then reject
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols isis export to.isis

[edit]
lab@R6# commit
commit complete
```

c) R2

You must disable the acceptance of the attached bit set in the LSP advertised by routers R5 and R6. This will limit router R2 to send traffic to destinations not present in the routing table.

```
[edit]
lab@R2# set protocols isis ignore-attached-bit

[edit]
lab@R2# commit
commit complete
```

d) Verify the configuration.

As a result, router R2 does not generate a default route despite the attached bit set.

```
[edit]
lab@R2# run show route 10.15.0.0/20

inet.0: 39 destinations, 39 routes (39 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.15.0.0/24       *[IS-IS/165] 00:01:20, metric 32
                    > to 172.20.20.78 via ge-0/0/4.52
                      to 172.20.20.82 via ge-0/0/4.62
10.15.1.0/24       *[IS-IS/165] 00:01:20, metric 32
                    > to 172.20.20.78 via ge-0/0/4.52
                      to 172.20.20.82 via ge-0/0/4.62
10.15.2.0/24       *[IS-IS/165] 00:01:20, metric 32
                    > to 172.20.20.78 via ge-0/0/4.52
                      to 172.20.20.82 via ge-0/0/4.62
10.15.3.0/24       *[IS-IS/165] 00:01:20, metric 32
                    > to 172.20.20.78 via ge-0/0/4.52
                      to 172.20.20.82 via ge-0/0/4.62
10.15.4.0/24       *[IS-IS/165] 00:01:20, metric 32
                    > to 172.20.20.78 via ge-0/0/4.52
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                      to 172.20.20.82 via ge-0/0/4.62
10.15.5.0/24        *[IS-IS/165] 00:01:20, metric 32
                     > to 172.20.20.78 via ge-0/0/4.52
                      to 172.20.20.82 via ge-0/0/4.62
10.15.6.0/24        *[IS-IS/165] 00:01:20, metric 32
                     > to 172.20.20.78 via ge-0/0/4.52
                      to 172.20.20.82 via ge-0/0/4.62
10.15.7.0/24        *[IS-IS/165] 00:01:20, metric 32
                     > to 172.20.20.78 via ge-0/0/4.52
                      to 172.20.20.82 via ge-0/0/4.62
10.15.8.0/24        *[IS-IS/165] 00:01:20, metric 32
                      to 172.20.20.78 via ge-0/0/4.52
                     > to 172.20.20.82 via ge-0/0/4.62
10.15.9.0/24        *[IS-IS/165] 00:01:20, metric 32
                      to 172.20.20.78 via ge-0/0/4.52
                     > to 172.20.20.82 via ge-0/0/4.62
10.15.10.0/24       *[IS-IS/165] 00:01:20, metric 32
                      to 172.20.20.78 via ge-0/0/4.52
                     > to 172.20.20.82 via ge-0/0/4.62
10.15.11.0/24       *[IS-IS/165] 00:01:20, metric 32
                      to 172.20.20.78 via ge-0/0/4.52
                     > to 172.20.20.82 via ge-0/0/4.62
10.15.12.0/24       *[IS-IS/165] 00:01:20, metric 32
                      to 172.20.20.78 via ge-0/0/4.52
                     > to 172.20.20.82 via ge-0/0/4.62
10.15.13.0/24       *[IS-IS/165] 00:01:20, metric 32
                      to 172.20.20.78 via ge-0/0/4.52
                     > to 172.20.20.82 via ge-0/0/4.62
10.15.14.0/24       *[IS-IS/165] 00:01:20, metric 32
                      to 172.20.20.78 via ge-0/0/4.52
                     > to 172.20.20.82 via ge-0/0/4.62
10.15.15.0/24       *[IS-IS/165] 00:01:20, metric 32
                      to 172.20.20.78 via ge-0/0/4.52
                     > to 172.20.20.82 via ge-0/0/4.62
```

```
[edit]
lab@R2# run show route 2015:4321::/32

inet6.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2015:4321::/64      *[IS-IS/165] 00:02:04, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                       to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:1::/64    *[IS-IS/165] 00:02:04, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                       to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:2::/64    *[IS-IS/165] 00:02:04, metric 40
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                       to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:3::/64   *[IS-IS/165] 00:02:04, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                       to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:4::/64   *[IS-IS/165] 00:02:04, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                       to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:5::/64   *[IS-IS/165] 00:02:04, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                       to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:6::/64   *[IS-IS/165] 00:02:04, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                       to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:7::/64   *[IS-IS/165] 00:02:04, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                       to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:8::/64   *[IS-IS/165] 00:02:04, metric 40
                       to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                     > to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:9::/64   *[IS-IS/165] 00:02:04, metric 40
                       to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                     > to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:a::/64   *[IS-IS/165] 00:02:04, metric 40
                       to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                     > to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:b::/64   *[IS-IS/165] 00:02:04, metric 40
                       to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                     > to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:c::/64   *[IS-IS/165] 00:02:04, metric 40
                       to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                     > to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:d::/64   *[IS-IS/165] 00:02:04, metric 40
                       to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                     > to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:e::/64   *[IS-IS/165] 00:02:04, metric 40
                       to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                     > to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
2015:4321:f::/64   *[IS-IS/165] 00:02:04, metric 40
                       to fe80::250:5600:34a9:2abe via ge-0/0/4.52
                     > to fe80::250:5600:3ea9:bd1 via ge-0/0/4.62
```

```
[edit]
lab@R2# run show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence Checksum Lifetime Attributes
R2.00-00                  0x57   0xb38f     1054 L1
R2.02-00                  0x52   0x9772      996 L1
R2.03-00                  0x54   0xd232     1153 L1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
R5.00-00                              0x81    0xf9f4       979 L1 L2 Attached
R5.02-00                              0x52    0x8079       999 L1 L2
R6.00-00                              0x59     0xae4      1035 L1 L2 Attached
  6 LSPs

IS-IS level 2 link-state database:
  0 LSPs
```

9) The physical design of the network provides two end-to-end symmetrical paths from router R2 to router R7. To better utilize the combined bandwidth, make sure that IPv4 routes prefers the south path, through routers R4 and R6. Whereas the IPv6 traffic takes the north path, through routers R3 and R5. This requirement should only concern the traffic originated from routers R2 and R7.

Although IS-IS provides routing for IPv4 and IPv6 protocols by default, both families use a single SPF algorithm to calculated paths in the network. However, you can configure IS-IS to calculate an alternate topology per family. Multitopology allows you to use per-family SPF calculations.

a) R2

An additional IPv6 unicast topology is configured on all routers on the path between R2 and R7. To force the two SPF algorithms to use different paths, the metric is tuned accordingly for the different topologies.

```
[edit]
lab@R2# set protocols isis topologies ipv6-unicast

[edit]
lab@R2# set protocols isis interface ge-0/0/4.52 level 1 metric 1000

[edit]
lab@R2# set protocols isis interface ge-0/0/4.62 level 1 ipv6-unicast-metric 1000

[edit]
lab@R2# commit
commit complete
```

b) R5

```
[edit]
lab@R5# set protocols isis topologies ipv6-unicast

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set protocols isis interface ge-0/0/4.52 level 1 metric 1000

[edit]
lab@R5# commit
commit complete
```

c) R6

```
[edit]
lab@R6# set protocols isis topologies ipv6-unicast

[edit]
lab@R6# set protocols isis interface ge-0/0/4.62 level 1 ipv6-unicast-metric 1000


[edit]
lab@R6# commit
commit complete
```

d) R3

```
[edit]
lab@R3# set protocols isis topologies ipv6-unicast

[edit]
lab@R3# set protocols isis interface ge-0/0/4.37 level 2 metric 1000

[edit]
lab@R3# commit
commit complete
```

e) R4

```
[edit]
lab@R4# set protocols isis topologies ipv6-unicast

[edit]
lab@R4# set protocols isis interface ge-0/0/4.47 level 2 ipv6-unicast-metric 1000

[edit]
lab@R4# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

f) R7

```
[edit]
lab@R7# set protocols isis topologies ipv6-unicast

[edit]
lab@R7# set protocols isis interface ge-0/0/4.37 level 2 metric 1000

[edit]
lab@R7# set protocols isis interface ge-0/0/4.47 level 2 ipv6-unicast-metric 1000

[edit]
lab@R7# commit
commit complete
```

g) R8

```
[edit]
lab@R8# set protocols isis topologies ipv6-unicast

[edit]
lab@R8# commit
commit complete
```

h) RR

```
[edit]
lab@route-reflector# set protocols isis topologies ipv6-unicast

[edit]
lab@route-reflector# commit
commit complete
```

i) Verify the configuration

After committing the change, router R2 calculates diverse paths per IPv4 and IPv6 families.

```
[edit]
lab@R2# run show route 10.15.0.0/20

inet.0: 39 destinations, 39 routes (39 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.15.0.0/24        *[IS-IS/165] 00:05:13, metric 32
                     > to 172.20.20.82 via ge-0/0/4.62
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
10.15.1.0/24        *[IS-IS/165] 00:05:13, metric 32
                     > to 172.20.20.82 via ge-0/0/4.62
10.15.2.0/24        *[IS-IS/165] 00:05:13, metric 32
                     > to 172.20.20.82 via ge-0/0/4.62
10.15.3.0/24        *[IS-IS/165] 00:05:13, metric 32
                     > to 172.20.20.82 via ge-0/0/4.62
...
```

```
[edit]
lab@R2# run show route 2015:4321::/32

inet6.0: 28 destinations, 28 routes (28 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2015:4321::/64      *[IS-IS/165] 00:03:42, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
2015:4321:1::/64    *[IS-IS/165] 00:03:42, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
2015:4321:2::/64    *[IS-IS/165] 00:03:42, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
2015:4321:3::/64    *[IS-IS/165] 00:03:42, metric 40
                     > to fe80::250:5600:34a9:2abe via ge-0/0/4.52
...
```

10) Make sure that R5 and R6 receive routes with metric higher than 63 from Level 2 routers R3 and R4.

By default, the IS-IS metric or cost can have value up to 63 with a combined path cost of 1023. If that is not sufficient, IS-IS can be configured to support much wider cost values. All routers that are part of the IS-IS Level 2 are configured for wide metrics only.

a) R3

```
[edit]
lab@R3# set protocols isis level 2 wide-metrics-only

[edit]
lab@R3# commit
commit complete
```

b) R4

```
[edit]
lab@R4# set protocols isis level 2 wide-metrics-only
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# commit
commit complete
```

a) R5

```
[edit]
lab@R5# set protocols isis level 2 wide-metrics-only

[edit]
lab@R5# commit
commit complete
```

a) R6

```
[edit]
lab@R6# set protocols isis level 2 wide-metrics-only

[edit]
lab@R6# commit
commit complete
```

b) R7

```
[edit]
lab@R7# set protocols isis level 2 wide-metrics-only

[edit]
lab@R7# commit
commit complete
```

c) R8

```
[edit]
lab@R8# set protocols isis level 2 wide-metrics-only

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

d) Verify the configuration

The output below shows that router R5 calculates a combined metric value above 1023.

Note: To test this you can deactivate the ge-0/0/5 interface on R7 to ensure the route with the higher metric is chosen as the path to reach the RIP network. The following output is from R5 after disabling the DC facing (RIP) interface on R7.

```
[edit]
lab@R5# run show route 10.15.7.0/24

inet.0: 46 destinations, 51 routes (46 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.15.7.0/24        *[IS-IS/18] 00:00:07, metric 2020
                     > to 172.20.20.17 via ge-0/0/4.45
```

11) Ensure that R7 and R8 will advertise overload bit set if DC-1 advertises more than 50 prefixes.

By default, there is no limit to the number of prefixes that can be exported into IS-IS. To protect the routing protocol from accidental or malicious events, limit on the exported routes can be applied. If that limit is exceeded, the router sets overload bit and the router is not used for transit traffic.

a) R7 and R8

```
[edit]
lab@R7# set protocols isis level 2 prefix-export-limit 50

[edit]
lab@R7# commit
commit complete
```

Task 3: Troubleshooting and Configuring BGP

This part is focused on configuring and troubleshooting the BGP protocol. You have to consider the given list of requirements that the network should meet. It is highly recommended to read the whole part before starting to solve the tasks.

1) Configure the necessary connectivity to the dedicated route-reflector device. It already has a partial configuration however you have to ensure that appropriate IGP configuration is in place to provide connectivity to and from the route-reflector.

© 2020 Juniper Networks, Inc. All rights reserved.

a) Route-reflector

Starting with the connectivity towards the route-reflector device from routers R3 and R4.

```
[edit]
lab@route-reflector# set interfaces lo0.0 family iso address 49.0005.0172.0020.2510.00

[edit]
lab@route-reflector# set protocols isis level 2 authentication-key jncie-sp-level2

[edit]
lab@route-reflector# set protocols isis level 2 authentication-type simple

[edit]
lab@route-reflector# set protocols isis level 2 wide-metrics-only

[edit]
lab@route-reflector# commit
commit complete
```

b) R3

```
[edit]
lab@R3# set interfaces ge-0/0/4.39 description "=== connection to route-reflector ==="

[edit]
lab@R3# set interfaces ge-0/0/4.39 vlan-id 39 family inet address 172.20.20.33/30

[edit]
lab@R3# set protocols isis interface ge-0/0/4.39

[edit]
lab@R3# commit
commit complete
```

c) R4

```
[edit]
lab@R4# set interfaces ge-0/0/4.49 description "=== connection to route-reflector ==="

[edit]
lab@R4# set interfaces ge-0/0/4.49 vlan-id 49 family inet address 172.20.20.37/30
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set protocols isis interface ge-0/0/4.49

[edit]
lab@R4# commit
commit complete
```

d) Verify the configuration

The router-reflector router successfully establishes adjacency with the routers R3 and R4.

```
[edit]
lab@route-reflector# run show isis adjacency
Interface             System           L State       Hold (secs) SNPA
ge-0/0/1.39           R3               2  Up                   6  2:6:a:e:3:4
ge-0/0/1.49           R4               2  Up                   8  2:6:a:e:4:4
```

2) Design IBGP network consisting of two confederations. AS65001 should include all L2 IS-IS routers, where AS65002 must include the rest of the routers. Configure redundant IBGP sessions between the two confederations.

BGP confederations is another way for limiting the IBGP full-mesh scaling problem. The network is split into smaller sub-ASes effectively lowering the number of the IBGP sessions required. Usually, the sub-AS numbers are taken from the private AS number range. Special external BGP sessions are used to connect the sub-AS, called cBGP. Redundant cBGP sessions are required by the task for successful completion.

3) Routers in confederations 65001 should use the route-reflector to peer with each other. The routers in the other confederation should peer in full mesh.

The task requires a router-reflector to be used for sub-AS 65001 and full-mesh sessions for AS65002.

a) R1

Apply the confederation configuration on R1 as part of sub-AS 65002. It peers using full-mesh sessions with the rest of the routers that are in IS-IS Level 1 area. The sub-AS number is configured as the autonomous-system. The global AS number representing the whole confederation is configured under the confederation section. Latter part is the same on all routers part from the same confederation. Last, the two sub-AS are configured as members of the confederation.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set routing-options autonomous-system 65002

[edit]
lab@R1# set routing-options confederation 23.89

[edit]
lab@R1# set routing-options confederation members 65001

[edit]
lab@R1# set routing-options confederation members 65002

[edit]
lab@R1# set protocols bgp group internal type internal

[edit]
lab@R1# set protocols bgp group internal local-address 172.20.25.1

[edit]
lab@R1# set protocols bgp group internal neighbor 172.20.25.5

[edit]
lab@R1# set protocols bgp group internal neighbor 172.20.25.6

[edit]
lab@R1# set protocols bgp group internal neighbor 172.20.25.2

[edit]
lab@R1# commit
commit complete
```

b) R2

Similar configuration is applied to R2.

```
[edit]
lab@R2# set routing-options autonomous-system 65002

[edit]
lab@R2# set routing-options confederation 23.89

[edit]
lab@R2# set routing-options confederation members 65001

[edit]
lab@R2# set routing-options confederation members 65002

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R2# set protocols bgp group internal type internal

[edit]
lab@R2# set protocols bgp group internal local-address 172.20.25.2

[edit]
lab@R2# set protocols bgp group internal neighbor 172.20.25.1

[edit]
lab@R2# set protocols bgp group internal neighbor 172.20.25.5

[edit]
lab@R2# set protocols bgp group internal neighbor 172.20.25.6

[edit]
lab@R2# commit
commit complete
```

c) R5

In addition to that, R5 is configured also with two cBGP sessions to routers R3 and R4 for redundancy as required by the task.

```
[edit]
lab@R5# set routing-options autonomous-system 65002

[edit]
lab@R5# set routing-options confederation 23.89

[edit]
lab@R5# set routing-options confederation members 65001

[edit]
lab@R5# set routing-options confederation members 65002

[edit]
lab@R5# set protocols bgp group internal type internal

[edit]
lab@R5# set protocols bgp group internal local-address 172.20.25.5

[edit]
lab@R5# set protocols bgp group internal neighbor 172.20.25.1

[edit]
lab@R5# set protocols bgp group internal neighbor 172.20.25.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols bgp group internal neighbor 172.20.25.6
```

cBGP sessions require explicit multihop configuration when the source and the destination are indirectly connected.

```
[edit]
lab@R5# set protocols bgp group sub-AS65001 multihop

[edit]
lab@R5# set protocols bgp group sub-AS65001 local-address 172.20.25.5

[edit]
lab@R5# set protocols bgp group sub-AS65001 peer-as 65001

[edit]
lab@R5# set protocols bgp group sub-AS65001 neighbor 172.20.25.3

[edit]
lab@R5# set protocols bgp group sub-AS65001 neighbor 172.20.25.4

[edit]
lab@R5# commit
commit complete
```

d) R6

Similar to router R5 configuration is applied to router R6.

```
[edit]
lab@R6# set routing-options autonomous-system 65002

[edit]
lab@R6# set routing-options confederation 23.89

[edit]
lab@R6# set routing-options confederation members 65001

[edit]
lab@R6# set routing-options confederation members 65002

[edit]
lab@R6# set protocols bgp group internal type internal
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols bgp group internal local-address 172.20.25.6

[edit]
lab@R6# set protocols bgp group internal neighbor 172.20.25.5

[edit]
lab@R6# set protocols bgp group internal neighbor 172.20.25.1

[edit]
lab@R6# set protocols bgp group internal neighbor 172.20.25.2

[edit]
lab@R6# set protocols bgp group sub-AS65001 multihop

[edit]
lab@R6# set protocols bgp group sub-AS65001 local-address 172.20.25.6

[edit]
lab@R6# set protocols bgp group sub-AS65001 peer-as 65001

[edit]
lab@R6# set protocols bgp group sub-AS65001 neighbor 172.20.25.3

[edit]
lab@R6# set protocols bgp group sub-AS65001 neighbor 172.20.25.4



[edit]
lab@R6# commit
commit complete
```

e) R3

R3 is part of sub-AS 65001. It peers with routers R5 and R6 with cBGP sessions to interconnect the two sub-AS confederations. Unlike the routers in sub-AS 65002, the routes from sub-AS 65001 should peer with the route-reflector.

```
[edit]
lab@R3# set routing-options autonomous-system 65001

[edit]
lab@R3# set routing-options confederation 23.89

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set routing-options confederation members 65001

[edit]
lab@R3# set routing-options confederation members 65002

[edit]
lab@R3# set protocols bgp group internal type internal

[edit]
lab@R3# set protocols bgp group internal local-address 172.20.25.3
```

A single IBGP session is configured to the router-reflector only.

```
 [edit]
lab@R3# set protocols bgp group internal neighbor 172.20.25.10 passive
```

Two multihop cBGP sessions are configured to the other sub-AS routers R5 and R6.

```
[edit]
lab@R3# set protocols bgp group sub-AS65002 multihop

[edit]
lab@R3# set protocols bgp group sub-AS65002 local-address 172.20.25.3

[edit]
lab@R3# set protocols bgp group sub-AS65002 peer-as 65002

[edit]
lab@R3# set protocols bgp group sub-AS65002 neighbor 172.20.25.5

[edit]
lab@R3# set protocols bgp group sub-AS65002 neighbor 172.20.25.6

[edit]
lab@R3# commit
commit complete
```

f) R4

Router R4 has similar configuration to router R3.

```
[edit]
lab@R4# set routing-options autonomous-system 65001

[edit]
lab@R4# set routing-options confederation 23.89

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R4# set routing-options confederation members 65001

[edit]
lab@R4# set routing-options confederation members 65002

[edit]
lab@R4# set protocols bgp group internal type internal

[edit]
lab@R4# set protocols bgp group internal local-address 172.20.25.4

[edit]
lab@R4# set protocols bgp group internal neighbor 172.20.25.10 passive

[edit]
lab@R4# set protocols bgp group sub-AS65002 multihop

[edit]
lab@R4# set protocols bgp group sub-AS65002 local-address 172.20.25.4

[edit]
lab@R4# set protocols bgp group sub-AS65002 peer-as 65002

[edit]
lab@R4# set protocols bgp group sub-AS65002 neighbor 172.20.25.5


[edit]
lab@R4# set protocols bgp group sub-AS65002 neighbor 172.20.25.6

[edit]
lab@R4# commit
commit complete
```

g) R7

Configure R7 as part of sub-AS 65001. Only a single IBGP session with the router-reflector is needed on R7.

```
[edit]
lab@R7# set routing-options autonomous-system 65001


[edit]
lab@R7# set routing-options confederation 23.89

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# set routing-options confederation members 65001

[edit]
lab@R7# set routing-options confederation members 65002

[edit]
lab@R7# set protocols bgp group internal type internal

[edit]
lab@R7# set protocols bgp group internal local-address 172.20.25.7

[edit]
lab@R7# set protocols bgp group internal neighbor 172.20.25.10 passive

[edit]
lab@R7# commit
commit complete
```

h) R8

Router R8 has similar configuration for router R7.

```
[edit]
lab@R8# set routing-options autonomous-system 65001

[edit]
lab@R8# set routing-options confederation 23.89

[edit]
lab@R8# set routing-options confederation members 65001

[edit]
lab@R8# set routing-options confederation members 65002

[edit]
lab@R8# set protocols bgp group internal type internal

[edit]
lab@R8# set protocols bgp group internal local-address 172.20.25.8

[edit]
lab@R8# set protocols bgp group internal neighbor 172.20.25.10 passive

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

i) Route-reflector

Last, the router-reflector itself is configured as part of the confederation. In addition to that, a single IBGP cluster is configured to reflect the routing information to all other routers that are part of sub-AS 65001.

```
[edit]
lab@route-reflector# set routing-options autonomous-system 65001

[edit]
lab@route-reflector# set routing-options confederation 23.89

[edit]
lab@route-reflector# set routing-options confederation members 65001

[edit]
lab@route-reflector# set routing-options confederation members 65002

[edit]
lab@route-reflector# set protocols bgp group internal type internal

[edit]
lab@route-reflector# set protocols bgp group internal local-address 172.20.25.10

[edit]
lab@route-reflector# set protocols bgp group internal cluster 1.1.1.1

[edit]
lab@route-reflector# set protocols bgp group internal neighbor 172.20.25.3

[edit]
lab@route-reflector# set protocols bgp group internal neighbor 172.20.25.4

[edit]
lab@route-reflector# set protocols bgp group internal neighbor 172.20.25.7

[edit]
lab@route-reflector# set protocols bgp group internal neighbor 172.20.25.8

[edit]
lab@route-reflector# commit
commit complete
```

j) Verify the configuration

© 2020 Juniper Networks, Inc. All rights reserved.

After committing the configuration, R3 has 3 established BGP sessions. Two cBGP sessions with sub-AS 65002 routers R5 and R6. The third is the IBGP session with the router-reflector.

```
[edit]
lab@R3# run show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                         0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.25.5           65002         24         23       0       0       10:25 0/0/0/0
0/0/0/0
172.20.25.6           65002         23         23       0       0       10:21 0/0/0/0
0/0/0/0
172.20.25.10          65001          5          4       0       0        1:35 0/0/0/0
0/0/0/0
```

R5 has a couple of more IBGP session because within sub-AS 65002 a route-reflector is not used. It requires IBGP session to all other internal routers in the sub-AS.

```
[edit]
lab@R5# run show bgp summary
Groups: 2 Peers: 5 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                         0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.25.1           65002         31         29       0       0       13:11 0/0/0/0
0/0/0/0
172.20.25.2           65002         31         29       0       0       13:10 0/0/0/0
0/0/0/0
172.20.25.3           65001         26         25       0       0       11:38 0/0/0/0
0/0/0/0
172.20.25.4           65001         25         24       0       0       11:05 0/0/0/0
0/0/0/0
172.20.25.6           65002         30         28       0       0       12:37 0/0/0/0
0/0/0/0
```

```
[edit]
lab@R2# run show bgp summary
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Groups: 1 Peers: 3 Down peers: 0
Table           Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                        0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.25.1           65002         38         38       0       0       17:02 0/0/0/0
0/0/0/0
172.20.25.5           65002         31         31       0       0       13:53 0/0/0/0
0/0/0/0
172.20.25.6           65002         30         29       0       0       13:11 0/0/0/0
0/0/0/0
```

4) Advertise an IPv4 summary route 172.20.25.0/28 from routers R3 and R4. The summary route should be advertised to all routers in your network.

a) R3

Next, you have to advertise summary route representing the loopback IPv4 addresses of all routers. The task requires doing that only from routers R3 and R4. An aggregate summary route is already present on router R3 and R4 needed for the previous task. Routing policy is configured and applied to the IBGP and cBGP neighbors.

```
[edit]
lab@R3# set policy-options policy-statement local.aggregate term v4.aggregate from
protocol aggregate

[edit]
lab@R3# set policy-options policy-statement local.aggregate term v4.aggregate from
route-filter 172.20.25.0/28 exact

[edit]
lab@R3# set policy-options policy-statement local.aggregate term v4.aggregate then
accept

[edit]
lab@R3# set protocols bgp group internal export local.aggregate

[edit]
lab@R3# set protocols bgp group sub-AS65002 export local.aggregate


[edit]
lab@R3# commit
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
commit complete
```

b) R4

Similar configuration is applied to a router R4.

```
[edit]
lab@R4# set policy-options policy-statement local.aggregate term v4.aggregate from
protocol aggregate

[edit]
lab@R4# set policy-options policy-statement local.aggregate term v4.aggregate from
route-filter 172.20.25.0/28 exact

[edit]
lab@R4# set policy-options policy-statement local.aggregate term v4.aggregate then
accept

[edit]
lab@R4# set protocols bgp group internal export local.aggregate

[edit]
lab@R4# set protocols bgp group sub-AS65002 export local.aggregate

[edit]
lab@R4# commit
commit complete
```

c) Route-reflector, R5, and R6

Because the router-reflector, R5 and R6 have the IPv4 summary route learned by IS-IS in their routing tables, the same route received via IBGP is marked as inactive. To force BGP to advertise the route despite that, you have to configure advertise-inactive knob.

Repeat this on the RR, R5, and R6 devices.

```
[edit]
lab@route-reflector# set protocols bgp advertise-inactive

[edit]
lab@route-reflector# commit
commit complete
```

d) Verify the configuration

© 2020 Juniper Networks, Inc. All rights reserved.

The outputs below show the summary route is advertised to router R5 from R3 and then further advertised to router R2.

```
[edit]
lab@R5# run show route receive-protocol bgp 172.20.25.3 172.20.25.0/28

inet.0: 49 destinations, 57 routes (49 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.20.25.0/28          172.20.25.3                  100        (65001) I
```

```
[edit]
lab@R5# run show route advertising-protocol bgp 172.20.25.2 172.20.25.0/28

inet.0: 49 destinations, 57 routes (49 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.20.25.0/28          172.20.25.4                  100        (65001) I
```

The router-reflector has the summary route in the routing table from 3 sources. The IS-IS learned route is active but even so, it is reflected to the other clients.

```
[edit]
lab@route-reflector# run show route 172.20.25.0/28 exact

inet.0: 40 destinations, 42 routes (40 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.0/28     *[IS-IS/18] 01:29:53, metric 20
                    > to 172.20.20.37 via ge-0/0/1.49
                      to 172.20.20.33 via ge-0/0/1.39
                    [BGP/170] 01:28:26, localpref 100, from 172.20.25.3
                      AS path: I, validation-state: unverified
                    > to 172.20.20.33 via ge-0/0/1.39
                    [BGP/170] 01:46:22, localpref 100, from 172.20.25.4
                      AS path: I, validation-state: unverified
                    > to 172.20.20.37 via ge-0/0/1.49
```

5) Routers R3 and R4 should also advertise IPv6 summary route 2001:fedc::/64 representing all IPv6 loopback addresses to the routers in AS65002. However, the summary must not be advertised to routers R1 and R2. Native IPv6 IBGP sessions are not allowed on the route-reflector.

© 2020 Juniper Networks, Inc. All rights reserved.

This task requires advertising IPv6 routes using a single BGP session to routers R5 and R6. However, with a single BGP sessions the next-hop attribute of the IPv6 routes is set by Junos to an IPv4-mapped embedded IPv6 address. Because the other routers in the network cannot resolve this address, it will stay hidden and not usable. You must change the next-hop address to the loopback IPv6 address already advertised by IS-IS.

a) R3

```
[edit]
lab@R3# set protocols bgp group internal family inet6 unicast
```

Adding the family inet6 to the session overrides the default inet family. This requires explicitly configuring the IBGP session with family inet also.

```
[edit]
lab@R3# set protocols bgp group internal family inet unicast

[edit]
lab@R3# set protocols bgp group sub-AS65002 family inet unicast

[edit]
lab@R3# set protocols bgp group sub-AS65002 family inet6 unicast

[edit]
lab@R3# set policy-options policy-statement local.aggregate term v6.aggregate from
protocol aggregate

[edit]
lab@R3# set policy-options policy-statement local.aggregate term v6.aggregate from
route-filter 2001:fedc::/64 exact
```

Setting the next-hop attribute to the local IPv6 loopback address.

```
[edit]
lab@R3# set policy-options policy-statement local.aggregate term v6.aggregate then
next-hop 2001:fedc::3

[edit]
lab@R3# set policy-options policy-statement local.aggregate term v6.aggregate then
accept

[edit]
lab@R3# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) R4

Router R4 is applied with similar configuration like on router R3.

```
[edit]
lab@R4# set protocols bgp group internal family inet unicast

[edit]
lab@R4# set protocols bgp group internal family inet6 unicast

[edit]
lab@R4# set protocols bgp group sub-AS65002 family inet unicast

[edit]
lab@R4# set protocols bgp group sub-AS65002 family inet6 unicast

[edit]
lab@R4# set policy-options policy-statement local.aggregate term v6.aggregate from
protocol aggregate

[edit]
lab@R4# set policy-options policy-statement local.aggregate term v6.aggregate from
route-filter 2001:fedc::/64 exact

[edit]
lab@R4# set policy-options policy-statement local.aggregate term v6.aggregate then
next-hop 2001:fedc::4

[edit]
lab@R4# set policy-options policy-statement local.aggregate term v6.aggregate then
accept

[edit]
lab@R4# commit
commit complete
```

c) R7

Family inet6 is added to the rest of the routers in AS 65001, including the router-reflector.

```
[edit]
lab@R7# set protocols bgp group internal family inet unicast

[edit]
lab@R7# set protocols bgp group internal family inet6 unicast
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# commit
commit complete
```

d) R8

```
[edit]
lab@R8# set protocols bgp group internal family inet unicast

[edit]
lab@R8# set protocols bgp group internal family inet6 unicast

[edit]
lab@R8# commit
commit complete
```

e) R5

The task requires the IPv6 summary route to be advertised to the routers R5 and R6. However, the rest of the routers within sub-AS 65002 should not receive this route. To accomplish that, you should enable family inet6 only to cBGP sessions on R5 and R6.

```
[edit]
lab@R5# set protocols bgp group sub-AS65001 family inet unicast

[edit]
lab@R5# set protocols bgp group sub-AS65001 family inet6 unicast

[edit]
lab@R5# commit
commit complete
```

f) R6

```
[edit]
lab@R6# set protocols bgp group sub-AS65001 family inet unicast

[edit]
lab@R6# set protocols bgp group sub-AS65001 family inet6 unicast

[edit]
lab@R6# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

g) Route-reflector

```
[edit]
lab@route-reflector# set protocols bgp group internal family inet unicast

[edit]
lab@route-reflector# set protocols bgp group internal family inet6 unicast

[edit]
lab@route-reflector# commit
commit complete
```

h) Verify the configuration.

The IPv6 summary route is present on R6 inet6.0 table with a next hop of R3 IPv6 loopback address. However, the route is not advertised further to router R2.

```
 [edit]
lab@R5# run show route receive-protocol bgp 172.20.25.3 2001:fedc::/64

inet6.0: 33 destinations, 36 routes (33 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  2001:fedc::/64          2001:fedc::3                 100        (65001) I
```

```
[edit]
lab@R5# run show route advertising-protocol bgp 172.20.25.2 2001:fedc::/64
```

6) Make sure that IBGP session on the route-reflector are not initiated on port 179.

a) All devices in AS65001 including the route-reflector (R3, R4, R7, R8).

The task requires that the route-reflector router waits for another router to send BGP session initial TCP request. Configuring the session as passive will accomplish that.

Repeat this on the R3, R4, R7, and R8 devices.

```
 [edit]
lab@R3# set protocols bgp group internal neighbor 172.20.25.10 passive

[edit]
lab@R3# commit
commit complete
```

b) Verify the configuration.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@route-reflector# run show bgp neighbor 172.20.25.3
Peer: 172.20.25.3+179 AS 65001 Local: 172.20.25.10+52316 AS 65001
  Group: internal                Routing-Instance: master
...
```

```
[edit]
lab@route-reflector# run show bgp neighbor 172.20.25.4
Peer: 172.20.25.4+179 AS 65001 Local: 172.20.25.10+53181 AS 65001
  Group: internal                Routing-Instance: master
...
```

7) Configure an EBGP session with the Upstream (U1, U2 and U3) peers. U1 and U2 are part of the same autonomous system, whereas U3 is in a different AS. However, there is connectivity and route exchange between the two autonomous-systems. Peers U1 and U2 provide only IPv4 connectivity. Whereas peer U3 provides IPv6 service also, make sure that you configure two separate sessions.

a) R1

Start with the configuration of the interfaces to U1 and U2.

```
[edit]
lab@R1# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R1# set interfaces ge-0/0/5 unit 512 description "=== connection to U2 ==="


[edit]
lab@R1# set interfaces ge-0/0/5 unit 512 vlan-id 512

[edit]
lab@R1# set interfaces ge-0/0/5 unit 512 family inet address 144.44.0.102/30
```

Configure the EBGP session with peer U2.

```
[edit]
lab@R1# set protocols bgp group U2 peer-as 16151

[edit]
lab@R1# set protocols bgp group U2 neighbor 144.44.0.101

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# commit
commit complete
```

b) R8

```
[edit]
lab@R8# set interfaces ge-0/0/5 unit 511 description "=== connection to U1 ==="

[edit]
lab@R8# set interfaces ge-0/0/5 unit 511 vlan-id 511

[edit]
lab@R8# set interfaces ge-0/0/5 unit 511 family inet address 144.44.0.2/30
```

Configure the EBGP session with peer U1.

```
[edit]
lab@R8# set protocols bgp group U1 peer-as 16151

[edit]
lab@R8# set protocols bgp group U1 neighbor 144.44.0.1

[edit]
lab@R8# commit
commit complete
```

c) R5

The task also requires that R5 should have two separate sessions to U3, one for IPv4 and other for IPv6. The interface is configured with IPv4 and IPv6 families and addresses as shown on the diagram. Two separate BGP groups are configured, one for IPv4 family and other for IPv6 family.

```
[edit]
lab@R5# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R5# set interfaces ge-0/0/5 unit 513 description "=== connection to U3 ==="

[edit]
lab@R5# set interfaces ge-0/0/5 unit 513 vlan-id 513

[edit]
lab@R5# set interfaces ge-0/0/5 unit 513 family inet address 134.45.0.2/30

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set interfaces ge-0/0/5 unit 513 family inet6 address 2001:134:45::2/80

[edit]
lab@R5# set protocols bgp group U3-v4 peer-as 74.39

[edit]
lab@R5# set protocols bgp group U3-v4 neighbor 134.45.0.1

[edit]
lab@R5# set protocols bgp group U3-v6 peer-as 74.39

[edit]
lab@R5# set protocols bgp group U3-v6 neighbor 2001:134:45::1

[edit]
lab@R5# commit
commit complete
```

d) Verify the configuration.

The BGP summary table on routers R8, R1 and R5 shows that EBGP sessions are established successfully and prefixes are received.

```
[edit]
lab@R8# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     461        460          0          0          0          0
inet6.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
144.44.0.1            16151        384          4       0       0        1:46
460/460/460/0        0/0/0/0
172.20.25.10          65001         22        398       0       1        8:56 Establ
  inet.0: 0/1/1/0
  inet6.0: 0/0/0/0
```

```
[edit]
lab@R1# run show bgp summary
Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     922        460          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
144.44.0.101            16151        387        385          0          0        3:14
1/460/460/0             0/0/0/0
172.20.25.2             65002        110        500          0          0       48:52 0/0/0/0
0/0/0/0
172.20.25.5             65002        485        494          0          0       45:44
459/461/461/0           0/0/0/0
172.20.25.6             65002        107        492          0          0       45:05 0/1/1/0
0/0/0/0
```

```
[edit]
lab@R5# run show bgp summary
Groups: 4 Peers: 7 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       464        460          0          0          0          0
inet6.0
                        90         88          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
134.45.0.1           4849703        382        391          0          0        1:13
460/460/460/0         0/0/0/0
172.20.25.1            65002        496        485          0          0       46:08 0/1/1/0
0/0/0/0
172.20.25.2            65002        104        486          0          0       46:07 0/0/0/0
0/0/0/0
172.20.25.3            65001         27        863          0          0       11:03 Establ
  inet.0: 0/1/1/0
  inet6.0: 0/1/1/0
172.20.25.4            65001         27        863          0          0       10:58 Establ
  inet.0: 0/1/1/0
  inet6.0: 0/1/1/0
172.20.25.6            65002        108        484          0          0       45:34 0/1/1/0
0/0/0/0
2001:134:45::1       4849703         83          3          0          0        1:05 Establ
  inet6.0: 88/88/88/0
```

8) Configure an EBGP session with Provider (P1) peer. You should configure a single session capable of receiving IPv4 and IPv6 prefixes. You are not allowed to use the *accept-remote-nexthop* option.

Unlike the connection to U3, where two separate EBGP sessions are used to configure R3 to peer with P1, you must configure a single EBGP session advertising both IPv4 and IPv6 prefixes.

a) R3

Configure the interface to connect with peer P1 according the diagram.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R3# set interfaces ge-0/0/5 unit 430 description "=== connection to P1 ==="

[edit]
lab@R3# set interfaces ge-0/0/5 unit 430 vlan-id 430

[edit]
lab@R3# set interfaces ge-0/0/5 unit 430 family inet address 172.30.0.1/30

[edit]
lab@R3# set interfaces ge-0/0/5 unit 430 family inet6 address 2001:172:30::1/64
```

By default, Junos will use IPv4-Mapped Embedded IPv6 Address as next-hop attribute for IPv6 prefixes advertised in IPv4 BGP sessions. Since the address is different from the address with which the EBGP session is established, BGP will not accept those prefixes. One way to solve that is using the accept-remote-nexthop knob but the Task explicitly does not allow using it. Other way to force the BGP process to accept those routes is configuring the sessions as multihop. This will allow the BGP to accept prefixes with any next-hop.

```
[edit]
lab@R3# set protocols bgp group P1 multihop
```

Although, the prefixes will be accepted, you must still amend the next-hop to the direct interface address to fix the data path also. Import routing policy is used to accomplish that.

```
[edit]
lab@R3# set protocols bgp group P1 import resolve.v6

[edit]
lab@R3# set protocols bgp group P1 family inet unicast

[edit]
lab@R3# set protocols bgp group P1 family inet6 unicast

[edit]
lab@R3# set protocols bgp group P1 peer-as 177.34

[edit]
lab@R3# set protocols bgp group P1 neighbor 172.30.0.2
```


© 2020 Juniper Networks, Inc. All rights reserved.

The import routing policy amending the next-hop of IPv6 prefixes received from peer P1 is defined below.

```
[edit]
lab@R3# set policy-options policy-statement resolve.v6 term 1 from family inet6

[edit]
lab@R3# set policy-options policy-statement resolve.v6 term 1 from next-hop
::ffff:172.30.0.2

[edit]
lab@R3# set policy-options policy-statement resolve.v6 term 1 then next-hop
2001:172:30::2

[edit]
lab@R3# set policy-options policy-statement resolve.v6 term 1 then accept



[edit]
lab@R3# commit
commit complete
```

b) Verify the configuration.

Number of IPv4 and IPv6 prefixes are received from peer P1.

```
[edit]
lab@R3# run show bgp summary
Groups: 3 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     935         16          0          0          0          0
inet6.0
                     104         16          0          0          0          0
Peer                     AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.25.5           65002        872         37       0       1       14:38 Establ
  inet.0: 0/460/460/0
  inet6.0: 0/88/88/0
172.20.25.6           65002        794         36       0       1       14:15 Establ
  inet.0: 0/459/459/0
  inet6.0: 0/0/0/0
172.20.25.10          65001         33         36       0       1       13:54 Establ
  inet.0: 0/0/0/0
  inet6.0: 0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.30.0.2          11599906          3          1        0        0          11 Establ
  inet.0: 16/16/16/0
  inet6.0: 16/16/16/0
```

The routes have their next hop changed to the direct interface of peer P1.

```
[edit]
lab@R3# run show route receive-protocol bgp 172.30.0.2 table inet6.0

inet6.0: 137 destinations, 137 routes (49 active, 0 holddown, 88 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  2001:f0f4:f691:d::/64
*                         2001:172:30::2                          11599906 I
  2001:f0f4:f691:d1::/64
*                         2001:172:30::2                          11599906 I
  2001:f0f4:f691:d2::/64
*                         2001:172:30::2                          11599906 I
  2001:f0f4:f691:d3::/64
*                         2001:172:30::2                          11599906 I
  2001:f0f4:f691:d4::/64
*                         2001:172:30::2                          11599906 I
  2001:f0f4:f691:d5::/64
*                         2001:172:30::2                          11599906 I
```

The output below is taken after the `multihop` feature is configured, but before the import routing policy changed the next hop.

```
[edit]
lab@R3# run show route receive-protocol bgp 172.30.0.2 table inet6.0

inet6.0: 138 destinations, 138 routes (50 active, 0 holddown, 88 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  2001:f0f4:f691:d::/64
*                         ::ffff:172.30.0.2                       11599906 I
  2001:f0f4:f691:d1::/64
*                         ::ffff:172.30.0.2                       11599906 I
  2001:f0f4:f691:d2::/64
*                         ::ffff:172.30.0.2                       11599906 I
  2001:f0f4:f691:d3::/64
*                         ::ffff:172.30.0.2                       11599906 I
  2001:f0f4:f691:d4::/64
*                         ::ffff:172.30.0.2                       11599906 I
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

9) Customer (C1) is using IPv4 and IPv6 services provided by your network. Two separate BGP sessions must be used. However, due to hardware limitation at the customer side the IPv6 interface can use only link-local addresses. You must discover the peer IPv6 link-local address and establish the EBGP session to the customer.

a) R6

BGP can establish session over the link local IPv6 addresses, eliminating the need of configuring IPv6 address for family inet6. The task speculates that you must find way to discover the IPv6 link local address.

```
[edit]
lab@R6# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R6# set interfaces ge-0/0/5 unit 410 vlan-id 410

[edit]
lab@R6# set interfaces ge-0/0/5 unit 410 family inet address 172.20.100.1/30

[edit]
lab@R6# commit
commit complete
```

One way to discover the link local IPv6 address is pinging the Multicast interface local scope address. Any device on the same segment will respond with its link local address.

```
[edit]
lab@R6# run ping ff02::1 interface ge-0/0/5.410
PING6(56=40+8+8 bytes) fe80::250:5601:9aa9:2fc0 --> ff02::1
16 bytes from fe80::250:5601:9aa9:5bf3, icmp_seq=0 hlim=64 time=114.826 ms
16 bytes from fe80::250:5601:9aa9:5bf3, icmp_seq=1 hlim=64 time=4.906 ms
16 bytes from fe80::250:5601:9aa9:5bf3, icmp_seq=2 hlim=64 time=4.650 ms
^C
--- ff02::1 ping6 statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 4.650/41.461/114.826/51.877 ms
```

Note: An IPv6 link local address is based on a prefix (FE80) and an interface identifier (EUI-64). For this reason, you will likely discover a different link local address that must be used as the one seen in the example above. You might also need to change the peering address of R6 configured on VR-DEVICE2 routing-instance C1.

```
[edit]
lab@R6# set protocols bgp group C1-v4 peer-as 65120
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols bgp group C1-v4 neighbor 172.20.100.2
```

To configure BGP session using the link local IPv6 address, you must specify the local interface. This specification is needed because any IPv6 enabled interface has a link local address automatically configured, and BGP must know which to use for the specific session.

```
[edit]
lab@R6# set protocols bgp group C1-v6 local-interface ge-0/0/5.410

[edit]
lab@R6# set protocols bgp group C1-v6 peer-as 65120

[edit]
lab@R6# set protocols bgp group C1-v6 neighbor fe80::250:5601:9aa9:5bf3



[edit]
lab@R6# commit
commit complete
```

b) Verify the configuration.

After committing the configuration, router R6 successfully established two EBGP sessions with peer C1.

```
[edit]
lab@R6# run show bgp summary
Groups: 4 Peers: 7 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     512        491          0          0          0          0
inet6.0
                      23         21          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.25.1           65002        518        133       0       0       56:49 1/1/1/0
0/0/0/0
172.20.25.2           65002        126        133       0       0       56:45 0/0/0/0
0/0/0/0
172.20.25.3           65001         53        811       0       0       22:00 Establ
  inet.0: 16/17/17/0
  inet6.0: 16/17/17/0
172.20.25.4           65001         52        814       0       0       21:56 Establ
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  inet.0: 0/1/1/0
  inet6.0: 0/1/1/0
172.20.25.5             65002        511        133       0       0       56:54
458/477/477/0          0/0/0/0
172.20.100.2            65120          9        387       0       0        2:50
16/16/16/0             0/0/0/0
fe80::250:5601:9aa9:5bf3       65120          7          7       0       0        2:04
Establ
  inet6.0: 5/5/5/0
```

10) Customer (C2) is also using IPv4 and IPv6 services over two separate EBGP sessions. The autonomous-system number of the customer is not provided. You need to discover the EBGP AS and configure the EBGP session using the md5 authentication password jncie.

To configure the EBGP sessions to peer with C2, you must first discover the peer AS number because the task speculates that it is not known.

a) R8

```
[edit]
lab@R8# set interfaces ge-0/0/5 unit 420 description "=== connection to C2 ===="

[edit]
lab@R8# set interfaces ge-0/0/5 unit 420 vlan-id 420

[edit]
lab@R8# set interfaces ge-0/0/5 unit 420 family inet address 172.20.200.1/30

[edit]
lab@R8# set interfaces ge-0/0/5 unit 420 family inet6 address ::ffff:172.20.200.1/80

[edit]
lab@R8# set protocols bgp group C2 neighbor 172.20.200.2
```

Configure a random chosen AS number temporary for the C2 peer session on router R8.

```
[edit]
lab@R8# set protocols bgp group C2 peer-as 65000

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Enabling traceoptions to capture the BGP log messages.

```
[edit]
lab@R8# set protocols bgp traceoptions file BGP_log

[edit]
lab@R8# set protocols bgp traceoptions flag open detail

[edit]
lab@R8# commit
commit complete
```

Look into the collected logs you can find. BGP open error message reveals the correct AS number configured on peer C2.

```
[edit]
lab@R8# run show log BGP_log
Dec  4 13:40:45 trace_on: Tracing to "/var/log/BGP_log" started
Dec  4 13:41:05.352405 bgp_send: sending 61 bytes to 172.20.200.2 (External AS 65000)
Dec  4 13:41:05.352555
Dec  4 13:41:05.352555 BGP SEND 172.20.200.1+54261 -> 172.20.200.2+179
Dec  4 13:41:05.352567 BGP SEND message type 1 (Open) length 61
Dec  4 13:41:05.352590 BGP SEND version 4 as 23456 holdtime 90 id 172.20.25.8 parmlen
32
Dec  4 13:41:05.352595 BGP SEND MP capability AFI=1, SAFI=1
Dec  4 13:41:05.352598 BGP SEND MP capability AFI=2, SAFI=1
Dec  4 13:41:05.352601 BGP SEND Refresh capability, code=128
Dec  4 13:41:05.352604 BGP SEND Refresh capability, code=2
Dec  4 13:41:05.352610 BGP SEND 4 Byte AS-Path capability (65), as_num 1507417
Dec  4 13:41:05.353738
Dec  4 13:41:05.353738 BGP RECV 172.20.200.2+179 -> 172.20.200.1+54261
Dec  4 13:41:05.353748 BGP RECV message type 1 (Open) length 71
Dec  4 13:41:05.353753 BGP RECV version 4 as 65123 holdtime 90 id 172.20.200.2 parmlen
42
Dec  4 13:41:05.353760 BGP RECV MP capability AFI=1, SAFI=1
Dec  4 13:41:05.353764 BGP RECV MP capability AFI=2, SAFI=1
Dec  4 13:41:05.353767 BGP RECV Refresh capability, code=128
Dec  4 13:41:05.353770 BGP RECV Refresh capability, code=2
Dec  4 13:41:05.353810 BGP RECV Restart capability, code=64, time=120,
flags=Notification
Dec  4 13:41:05.353815 BGP RECV 4 Byte AS-Path capability (65), as_num 65123
Dec  4 13:41:05.353820 BGP RECV Long-Lived Graceful Restart capability, code=71
Dec  4 13:41:05.353887 bgp_process_open:3871: NOTIFICATION sent to 172.20.200.2
(External AS 65000): code 2 (Open Message Error) subcode 2 (bad peer AS number),
Reason: peer 172.20.200.2 (External AS 65000) claims 65123, 65000 configured
Dec  4 13:41:05.353893 bgp_send: sending 21 bytes to 172.20.200.2 (External AS 65000
```


© 2020 Juniper Networks, Inc. All rights reserved.

After the correct AS number is discovered, the configuration is amended accordingly. The password given with the task description is used for the EBGP session to peer C2.

```
[edit]
lab@R8# set protocols bgp group C2 family inet unicast

[edit]
lab@R8# set protocols bgp group C2 family inet6 unicast

[edit]
lab@R8# set protocols bgp group C2 authentication-key "$9$ealMLNGUHmPQs2Zj"

[edit]
lab@R8# set protocols bgp group C2 peer-as 65123

[edit]
lab@R8# set protocols bgp group C2 neighbor 172.20.200.2

[edit]
lab@R8# delete protocols bgp traceoptions


[edit]
lab@R8# commit
commit complete
```

b) Verify the configuration.

R2 successfully established a single EBGP session with peer C2, routing both IPv4 and IPv6 traffic. Amending the next hop is not needed here because IPv4-Mapped Embedded IPv6 Address is configured on the interface.

```
[edit]
lab@R8# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     477        476          0          0          0          0
inet6.0
                       5          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
144.44.0.1            16151        432         52       0       0       22:48
460/460/460/0         0/0/0/0
172.20.25.10          65001         68        448       0       1       29:58 Establ
  inet.0: 0/1/1/0
  inet6.0: 0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.20.200.2          65123          3          1       0       0          12 Establ
  inet.0: 16/16/16/0
  inet6.0: 5/5/5/0
```

These next few tasks are mostly based on BGP route manipulation according to the given requirements. This is one of the most important parts of the JNCIE-SP exams.

11) All IPv4 and IPv6 prefixes received from Upstream, Partner, and Customer peers must be active and reachable on all devices in your AS. Unless they are filtered by policy in some of the subsequent tasks.

First, you must make sure all the routes are resolved within the network beyond the border routers. There are a few ways you can achieved that, however the next-hop set by a routing policy is the easiest and recommended. The IPv6 traffic requires manual setting the next hop to the local loopback IPv6 address on each border router peering with an external IPv6 neighbor. That is needed because internal BGP is using a single session for both, IPv4 and IPv6, which means that Junos will set an IPv4-Mapped IPv6 address by default.

a) R3

```
[edit]
lab@R3# set policy-options policy-statement nhs term v6 from protocol bgp
```

The routing policy term related to the IPv6 routes matches on inet6.0 routing table.

```
[edit]
lab@R3# set policy-options policy-statement nhs term v6 from rib inet6.0
```

As-path matching only routes advertised by the direct neighbor AS of P1.

```
[edit]
lab@R3# set policy-options policy-statement nhs term v6 from as-path P1.adjacent
```

The next-hop address is manually set to the local loopback IPv6 address.

```
[edit]
lab@R3# set policy-options policy-statement nhs term v6 then next-hop 2001:fedc::3

[edit]
lab@R3# set policy-options policy-statement nhs term v6 then accept

[edit]
lab@R3# set policy-options policy-statement nhs term v4 from as-path P1.adjacent

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# set policy-options policy-statement nhs term v4 then next-hop self
```

To match 4byte AS numbers in regular expressions, you must escape the dot as it is configured below.

```
[edit]
lab@R3# set policy-options as-path P1.adjacent "177\.34 .*"
```

The routing policy is applied in the export direction to the IBGP and cBGP groups.

```
[edit]
lab@R3# set protocols bgp group internal export nhs

[edit]
lab@R3# set protocols bgp group sub-AS65002 export nhs

[edit]
lab@R3# commit
commit complete
```

b) R1

```
[edit]
lab@R1# set policy-options policy-statement nhs term 1 then next-hop self

[edit]
lab@R1# set protocols bgp group internal export nhs

[edit]
lab@R1# commit
commit complete
```

c) R5

```
[edit]
lab@R5# set policy-options policy-statement nhs term v6 from protocol bgp


[edit]
lab@R5# set policy-options policy-statement nhs term v6 from rib inet6.0

[edit]
lab@R5# set policy-options policy-statement nhs term v6 from as-path U3.adjacent

[edit]
lab@R5# set policy-options policy-statement nhs term v6 then next-hop 2001:fedc::5

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set policy-options policy-statement nhs term v6 then accept

[edit]
lab@R5# set policy-options policy-statement nhs term v4 from as-path U3.adjacent

[edit]
lab@R5# set policy-options policy-statement nhs term v4 then next-hop self

[edit]
lab@R5# set policy-options as-path U3.adjacent "74\.39 .*"

[edit]
lab@R5# set protocols bgp group internal export nhs

[edit]
lab@R5# set protocols bgp group sub-AS65001 export nhs

[edit]
lab@R5# commit
commit complete
```

d) R8

```
[edit]
lab@R8# set policy-options policy-statement nhs term v6 from protocol bgp

[edit]
lab@R8# set policy-options policy-statement nhs term v6 from rib inet6.0

[edit]
lab@R8# set policy-options policy-statement nhs term v6 then next-hop 2001:fedc::8

[edit]
lab@R8# set policy-options policy-statement nhs term v6 then accept

[edit]
lab@R8# set policy-options policy-statement nhs term v4 then next-hop self


[edit]
lab@R8# set protocols bgp group internal export nhs

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

e) R6

```
[edit]
lab@R6# set policy-options policy-statement nhs term v6 from protocol bgp

[edit]
lab@R6# set policy-options policy-statement nhs term v6 from rib inet6.0

[edit]
lab@R6# set policy-options policy-statement nhs term v6 from as-path C1.origin

[edit]
lab@R6# set policy-options policy-statement nhs term v6 then next-hop 2001:fedc::6

[edit]
lab@R6# set policy-options policy-statement nhs term v6 then accept

[edit]
lab@R6# set policy-options policy-statement nhs term v4 from as-path C1.origin

[edit]
lab@R6# set policy-options policy-statement nhs term v4 then next-hop self

[edit]
lab@R6# set policy-options as-path C1.origin ".* 65120"

[edit]
lab@R6# set protocols bgp group internal export nhs

[edit]
lab@R6# set protocols bgp group sub-AS65001 export nhs

[edit]
lab@R6# commit
commit complete
```

f) Verify the configuration.

The route-reflector receives and reflects the IPv4 and IPv6 routes from all border routers as active.

```
[edit]
lab@route-reflector# run show bgp summary
Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     970        508          0          0          0          0
inet6.0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    209          5          0          0          0          0
Peer                     AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.25.3           65001       1112       1031       0       0       36:15 Establ
  inet.0: 492/493/493/0
  inet6.0: 0/110/110/0
172.20.25.4           65001        824       1323       0       0       36:11 Establ
  inet.0: 0/1/1/0
  inet6.0: 0/94/94/0
172.20.25.7           65001         80       1151       0       0       36:08 Establ
  inet.0: 0/0/0/0
  inet6.0: 0/0/0/0
172.20.25.8           65001        841       1243       0       0       36:01 Establ
  inet.0: 16/476/476/0
  inet6.0: 5/5/5/0
```

R2 has none of the IPv6 routes installed because it has no BGP session enabled for family inet6.

```
[edit]
lab@R2# run show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     509        508          0          0          0          0
Peer                     AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.25.1           65002        563        169       0       0     1:16:30 0/0/0/0
0/0/0/0
172.20.25.5           65002        933        162       0       0     1:13:21
492/493/493/0        0/0/0/0
172.20.25.6           65002        175        160       0       0     1:12:39
16/16/16/0           0/0/0/0
```

12) Advertise the IBGP IPv4 summary route to all external peers except to P1. You are not allowed to use route policy to accomplish this.

On most border routers, the local IPv4 summary route learned by BGP is not active because the same route is advertised by the IS-IS protocol. To force BGP, advertise the inactive routes you have to use `advertise-inactive` option. We have already enabled this option on R5 and R6 in a previous task.

a) R3

```
[edit]
lab@R3# set protocols bgp advertise-inactive
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# commit
commit complete
```

b) R8

```
[edit]
lab@R8# set protocols bgp advertise-inactive

[edit]
lab@R8# commit
commit complete
```

c) Verify the configuration.

IPv4 summary route is successfully advertised, despite is inactive in the routing table.

```
[edit]
lab@R5# run show route advertising-protocol bgp 134.45.0.1 172.20.25.0/28

inet.0: 547 destinations, 1958 routes (542 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.20.25.0/28          Self                                    I
```

```
[edit]
lab@R6# run show route advertising-protocol bgp 172.20.100.2 172.20.25.0/28

inet.0: 541 destinations, 1495 routes (541 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.20.25.0/28          Self                                    I
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 144.44.0.1 172.20.25.0/28

inet.0: 559 destinations, 576 routes (554 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.20.25.0/28          Self                                    I
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.20.200.2 172.20.25.0/28

inet.0: 559 destinations, 576 routes (554 active, 0 holddown, 0 hidden)
Restart Complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.20.25.0/28          Self                                    I
```

```
[edit]
lab@R1# run show route advertising-protocol bgp 144.44.0.101 172.20.25.0/28

inet.0: 522 destinations, 982 routes (522 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.20.25.0/28          Self                                       I
```

13) Router R1 is chosen as backup gateway to AS 16151 therefore the routes advertised from peer U1 are prepended to enforce that in the outbound direction. Make sure that for the inbound direction R1 is also less preferred than the path to router R8.

a) R1

MED or metric BGP attribute is compared between routes received from the same peering AS. Since, peers U1 and U2 are peering with same ASs, you can increase the metric on the routes send from router R1 out to U2. Setting the metric directly on the neighbor configuration under BGP will ensure that all routes to that neighbor have their metric increased.

```
[edit]
lab@R1# set protocols bgp group U2 metric-out 100

[edit]
lab@R1# commit
commit complete
```

b) Verify the configuration.

After the change, router R1 is advertising metric value 100 to peer U1.

```
[edit]
lab@R1# run show route advertising-protocol bgp 144.44.0.101 172.20.25.0/28

inet.0: 522 destinations, 982 routes (522 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.20.25.0/28          Self                 100                I
```


© 2020 Juniper Networks, Inc. All rights reserved.

14) Make sure that router R8 is the preferred exit gateway for the same prefixes received from all Upstream peers. This requirement should not affect R8 routing table. In case the connectivity to U1 is lost, R1 and R5 should prefer their local connections to the upstream connections.

a) R8

The local-preference attribute is amongst the first compared in the Junos best path selection algorithm, a higher value is preferred. Setting the local-preference to 200 for the external routes advertised by router R8 will make it the preferred gateway to the Upstream destinations. To avoid this change to have effect locally on router R8, apply the policy in the export direction to the route reflector.

```
[edit]
lab@R8# set policy-options policy-statement to.internal term lp.higher then local-
preference 200

[edit]
lab@R8# set protocols bgp group internal export to.internal

[edit]
lab@R8# commit
commit complete
```

b) R1

The second part of the task speculates about a case is which R8 is missing in the network. Then R1 and R5 should prefer their local connection to the U2 and U3 respectively. The BGP algorithm prefers routes received from external peers to routes received from external peer through internal session. However, we know from the previous task, that peer U2 advertises the routes to R1 prepended. Because of that, the path selection algorithm will not reach the point where external routes are compared. Comparing the as-path length is higher in the evaluation check. Therefore, you have to force BGP to skip this step and continue down in the algorithm by using the knob below.

```
[edit]
lab@R1# set protocols bgp path-selection as-path-ignore

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

c) Verify the configuration.

Router R8 advertises the peer U1 routes with an increased local-preference to the internal BGP neighbors. The local-preference of routes in R8 routing table are not affected.

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.20.25.10


inet.0: 555 destinations, 556 routes (555 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 0.0.0.0/4               Self                         200        16151 61671 I
* 0.0.0.0/5               Self                         200        16151 61671 27075 I
* 1.64.0.0/10             Self                         200        16151 61671 I
* 1.84.160.0/20           Self                         200        16151 33112 I
* 1.96.0.0/11             Self                         200        16151 33112 63164
40776 51777 I
* 1.161.192.0/21          Self                         200        16151 33112 30404
32138 45045 I
```

```
[edit]
lab@R8# run show route

inet.0: 555 destinations, 556 routes (555 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

0.0.0.0/4          *[BGP/170] 00:39:40, localpref 100
                      AS path: 16151 61671 I, validation-state: unverified
                    > to 144.44.0.1 via ge-0/0/5.511
0.0.0.0/5          *[BGP/170] 00:39:40, localpref 100
                      AS path: 16151 61671 27075 I, validation-state: unverified
                    > to 144.44.0.1 via ge-0/0/5.511
1.64.0.0/10        *[BGP/170] 00:39:40, localpref 100
                      AS path: 16151 61671 I, validation-state: unverified
                    > to 144.44.0.1 via ge-0/0/5.511
```

15) Prefixes advertised by C1 should be uniquely marked in AS65001 only. Also, C1 routes must not be advertised to any EBGP peer except to peer C2. In reverse direction, peer C1 should receive only the C2 routes and the local summary routes.

a) R6

© 2020 Juniper Networks, Inc. All rights reserved.

The task simply requires configuring routing policy controlling the routes received and advertised to peer C1.

```
[edit]
lab@R6# set policy-options policy-statement to.AS65001 term C1.routes from protocol bgp
```

The routing-policy matches only routes originated by C1 AS.

```
[edit]
lab@R6# set policy-options policy-statement to.AS65001 term C1.routes from as-path
C1.origin
```

Unique community is applied to the routes advertised to sub-AS 65001 routers.

```
[edit]
lab@R6# set policy-options policy-statement to.AS65001 term C1.routes then community
add C1.routes
```

The term below is rejecting the IPv4 routes advertised by Upstreams to peer C1.

```
[edit]
lab@R6# set policy-options policy-statement to.C1-v4 term reject.Upstream from protocol
bgp

[edit]
lab@R6# set policy-options policy-statement to.C1-v4 term reject.Upstream from as-path
U.adjacent

[edit]
lab@R6# set policy-options policy-statement to.C1-v4 term reject.Upstream then reject
```

The term below is accepting the IPv6 routes advertised by C2 to peer C1.

```
[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term C2.routes from protocol bgp

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term C2.routes from rib inet6.0

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term C2.routes from as-path
C2.origin

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term C2.routes then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

Explicitly advertising the IPv6 summary route to peer C1.

```
[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term local.aggregate from protocol
bgp

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term local.aggregate from route-
filter 2001:fedc::/64 exact

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term local.aggregate then accept
```

Rejecting the rest of the IPv6 routes advertised to peer C1.

```
[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term reject.rest-v6 from rib
inet6.0

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term reject.rest-v6 then reject

[edit]
lab@R6# set policy-options community C1.routes members 65120:100

[edit]
lab@R6# set policy-options as-path C2.origin ".* 65123"
```

The regular expression below is matching on routes having adjacent AS number of either U1 and U2 or either U3. The same result can be achieved using separate regular expressions, but practice using advanced regular expressions helps for solving complicated tasks. Due to the use of confederations, you can have the most left AS set to some of the sub-AS numbers.

```
[edit]
lab@R6# set policy-options as-path U.adjacent ".? (16151|74\.39) .*"
```

The routing policies are applied to the BGP neighbors to which they are intended.

```
[edit]
lab@R6# set protocols bgp group C1-v4 export to.C1-v4

[edit]
lab@R6# set protocols bgp group C1-v6 export to.C1-v6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols bgp group sub-AS65001 export to.AS65001

[edit]
lab@R6# commit
commit complete
```

b) R1

```
[edit]
lab@R1# set policy-options policy-statement to.U2 term C1.reject from protocol bgp

[edit]
lab@R1# set policy-options policy-statement to.U2 term C1.reject from as-path C1.origin

[edit]
lab@R1# set policy-options policy-statement to.U2 term C1.reject then reject

[edit]
lab@R1# set policy-options as-path C1.origin ".* 65120"

[edit]
lab@R1# set protocols bgp group U2 export to.U2

[edit]
lab@R1# commit
commit complete
```

c) R5

You must consider the case when router R6 sessions to sub-AS are down and configure router R5 to attach community to C1 routers. In addition, you must reject the C1 advertised routes egress peer U3.

```
[edit]
lab@R5# set policy-options policy-statement to.AS65001 term C1.routes from protocol bgp

[edit]
lab@R5# set policy-options policy-statement to.AS65001 term C1.routes from as-path
C1.origin

[edit]
lab@R5# set policy-options policy-statement to.AS65001 term C1.routes then community
add C1.routes

[edit]
lab@R5# set policy-options policy-statement to.U3-v4 term C1.reject from protocol bgp
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement to.U3-v4 term C1.reject from as-path
C1.origin

[edit]
lab@R5# set policy-options policy-statement to.U3-v4 term C1.reject then reject

[edit]
lab@R5# set policy-options as-path C1.origin ".* 65120"

[edit]
lab@R5# set policy-options community C1.routes members 65120:100

[edit]
lab@R5# set protocols bgp group sub-AS65001 export to.AS65001

[edit]
lab@R5# set protocols bgp group U3-v4 export to.U3-v4


[edit]
lab@R5# commit
commit complete
```

d) Verify the configuration.

After applying the routing policies, you should see the summary route for the internal network and the routes learned from C2 advertised to peer C1.

```
[edit]
lab@R6# run show route advertising-protocol bgp 172.20.100.2

inet.0: 558 destinations, 1094 routes (558 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 101.15.128.0/24         Self                                    65123 I
* 101.15.129.0/24         Self                                    65123 I
* 101.15.130.0/24         Self                                    65123 I
* 101.15.131.0/24         Self                                    65123 I
* 101.15.132.0/24         Self                                    65123 I
* 101.15.133.0/24         Self                                    65123 I
* 101.15.134.0/24         Self                                    65123 I
* 101.15.135.0/24         Self                                    65123 I
* 101.15.136.0/24         Self                                    65123 I
* 101.15.137.0/24         Self                                    65123 I
* 101.15.138.0/24         Self                                    65123 I
* 101.15.139.0/24         Self                                    65123 I
* 101.15.140.0/24         Self                                    65123 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 101.15.141.0/24          Self                                  65123 I
* 101.15.142.0/24          Self                                  65123 I
* 101.15.143.0/24          Self                                  65123 I
* 105.24.64.0/24           Self                                  11599906 I
* 105.25.65.0/24           Self                                  11599906 I
* 105.26.66.0/24           Self                                  11599906 I
* 105.27.67.0/24           Self                                  11599906 I
* 105.28.68.0/24           Self                                  11599906 I
* 105.29.69.0/24           Self                                  11599906 I
* 105.30.70.0/24           Self                                  11599906 I
* 105.31.71.0/24           Self                                  11599906 I
* 105.32.72.0/24           Self                                  11599906 I
* 105.33.73.0/24           Self                                  11599906 I
* 105.34.74.0/24           Self                                  11599906 I
* 105.35.75.0/24           Self                                  11599906 I
* 105.36.76.0/24           Self                                  11599906 I
* 105.36.77.0/24           Self                                  11599906 I
* 105.36.78.0/24           Self                                  11599906 I
* 105.36.79.0/24           Self                                  11599906 I
  172.20.25.0/28           Self                                  I
```

The C1 routes are assigned with community on routers part of sub-AS 65001.

```
[edit]
lab@R8# run show route 101.10.32.0/24 extensive

inet.0: 555 destinations, 572 routes (555 active, 0 holddown, 0 hidden)
Restart Complete
101.10.32.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 101.10.32.0/24 -> {indirect(1048574)}
Page 0 idx 1, (group U1 type External) Type 1 val 0xd97b600 (adv_entry)
   Advertised metrics:
     Nexthop: Self
     AS path: [65001] 65120 I
     Communities: 65120:100
Page 0 idx 2, (group C2 type External) Type 1 val 0xd97b440 (adv_entry)
   Advertised metrics:
     Nexthop: Self
     AS path: [65001] 65120 I
     Communities: 65120:100
Path 101.10.32.0 from 172.20.25.10 Vector len 4.  Val: 1 2
        *BGP    Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xbe78e30
                Next-hop reference count: 65
                Source: 172.20.25.10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Next hop type: Router, Next hop index: 607
Next hop: 172.20.20.25 via ge-0/0/4.48, selected
Session Id: 0x1a6
Protocol next hop: 172.20.25.3
Indirect next hop: 0xc29d300 1048574 INH Session ID: 0x1ac
State: <Active Int Ext>
Local AS: 65001 Peer AS: 65001
Age: 5:11       Metric2: 20
Validation State: unverified
ORR Generation-ID: 0
Task: BGP_65001.172.20.25.10
Announcement bits (3): 0-KRT 2-BGP_RT_Background 3-Resolve tree 1
AS path: (65002) 65120 I  (Originator)
Cluster list:  1.1.1.1
Originator ID: 172.20.25.3
Communities: 65120:100
Accepted
Localpref: 100
Router ID: 172.20.25.10
Indirect next hops: 1
        Protocol next hop: 172.20.25.3 Metric: 20
        Indirect next hop: 0xc29d300 1048574 INH Session ID: 0x1ac
        Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 172.20.20.25 via ge-0/0/4.48
                Session Id: 0x1a6
                172.20.25.3/32 Originating RIB: inet.0
                  Metric: 20    Node path count: 1
                  Forwarding nexthops: 1
                        Nexthop: 172.20.20.25 via ge-0/0/4.48
                        Session Id: 1a6
```

16) Only advertise a summary route representing all C2 IPv4 routes to the internal routers and make sure it is advertised only to the Upstream neighbors and C1. Ensure all the routes received from C2 as originated from its own autonomous-system. In reverse direction, C2 must receive the C1 routes and Upstream routes.

The task simply requires configuring routing policy controlling the routes received and send to peer C1.

a) R8

Configuring routing policy ensures that only C2 AS originated routes are received from peer C2.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set policy-options policy-statement from.C2 term C2.origin from protocol bgp

[edit]
lab@R8# set policy-options policy-statement from.C2 term C2.origin from as-path
C2.origin

[edit]
lab@R8# set policy-options policy-statement from.C2 term C2.origin then accept

[edit]
lab@R8# set policy-options policy-statement from.C2 term reject.rest then reject
```

Configuring export routing policies rejecting the more specific C2 routes and advertise only a summary route to the internal routers and peer U1.

```
[edit]
lab@R8# set policy-options policy-statement to.internal term C2.aggregate from protocol
aggregate

[edit]
lab@R8# set policy-options policy-statement to.internal term C2.aggregate from route-
filter 101.15.128.0/20 exact

[edit]
lab@R8# set policy-options policy-statement to.internal term C2.aggregate then accept

[edit]
lab@R8# set policy-options policy-statement to.internal term C2.specific from protocol
bgp

[edit]
lab@R8# set policy-options policy-statement to.internal term C2.specific from route-
filter 101.15.128.0/20 longer

[edit]
lab@R8# set policy-options policy-statement to.internal term C2.specific then reject
```

The U1 export policy rejects the C1 routes matched by community.

```
[edit]
lab@R8# set policy-options policy-statement to.U1 term C1.reject from protocol bgp

[edit]
lab@R8# set policy-options policy-statement to.U1 term C1.reject from community
C1.routes

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R8# set policy-options policy-statement to.U1 term C1.reject then reject

[edit]
lab@R8# set policy-options policy-statement to.U1 term C2.aggregate from protocol
aggregate

[edit]
lab@R8# set policy-options policy-statement to.U1 term C2.aggregate from route-filter
101.15.128.0/20 exact

[edit]
lab@R8# set policy-options policy-statement to.U1 term C2.aggregate then accept

[edit]
lab@R8# set policy-options policy-statement to.U1 term C2.reject from protocol bgp

[edit]
lab@R8# set policy-options policy-statement to.U1 term C2.reject from as-path C2.origin

[edit]
lab@R8# set policy-options policy-statement to.U1 term C2.reject then reject

[edit]
lab@R8# set policy-options community C1.routes members 65120:100

[edit]
lab@R8# set policy-options as-path C2.origin ".* 65123"

[edit]
lab@R8# set routing-options aggregate route 101.15.128.0/20

[edit]
lab@R8# set protocols bgp group C2 import from.C2

[edit]
lab@R8# set protocols bgp group U1 export to.U1

[edit]
lab@R8# commit
commit complete
```

b) Verify the configuration.

After the routing policies are applied to router R8, the task requirements are accomplished.

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.20.200.2 101.10.32.0/24
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
inet.0: 559 destinations, 576 routes (554 active, 0 holddown, 5 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 101.10.32.0/24          Self                                    65120 I
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.20.25.10 101.15.128.0/20

inet.0: 559 destinations, 576 routes (554 active, 0 holddown, 5 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 101.15.128.0/20         Self                            200     65123 I
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 144.44.0.1 101.15.128.0/20

inet.0: 559 destinations, 576 routes (554 active, 0 holddown, 5 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 101.15.128.0/20         Self                                    65123 I
```

```
[edit]
lab@R8# run show route receive-protocol bgp 172.20.200.2

inet.0: 559 destinations, 576 routes (554 active, 0 holddown, 5 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 101.15.128.0/24         172.20.200.2                            65123 I
* 101.15.129.0/24         172.20.200.2                            65123 I
* 101.15.130.0/24         172.20.200.2                            65123 I
* 101.15.131.0/24         172.20.200.2                            65123 I
* 101.15.132.0/24         172.20.200.2                            65123 I
* 101.15.133.0/24         172.20.200.2                            65123 I
* 101.15.134.0/24         172.20.200.2                            65123 I
* 101.15.135.0/24         172.20.200.2                            65123 I
* 101.15.136.0/24         172.20.200.2                            65123 I
* 101.15.137.0/24         172.20.200.2                            65123 I
* 101.15.138.0/24         172.20.200.2                            65123 I
* 101.15.139.0/24         172.20.200.2                            65123 I
* 101.15.140.0/24         172.20.200.2                            65123 I
* 101.15.141.0/24         172.20.200.2                            65123 I
* 101.15.142.0/24         172.20.200.2                            65123 I
* 101.15.143.0/24         172.20.200.2                            65123 I

inet6.0: 149 destinations, 166 routes (149 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:10:15:132::/64     ::ffff:172.20.200.2                     65123 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 2001:10:15:133::/64      ::ffff:172.20.200.2                65123 I
* 2001:10:15:134::/64      ::ffff:172.20.200.2                65123 I
* 2001:10:15:135::/64      ::ffff:172.20.200.2                65123 I
* 2001:10:15:136::/64      ::ffff:172.20.200.2                65123 I
```

17) The P1 routes received by R3 should not be advertised to any other peer. P1 should receive all RIP routes received from DC-1. Make sure that if DC-1 advertises more routes they are advertised automatically to P1.

a) R7 and R8

Setting a tag on routes received from DC-1 and matching on the tagged routes on R7. This will ensure that if more routes are advertised by DC-1, they will be selected too.

R7

```
[edit]
lab@R7# set policy-options policy-statement from.rip term accept.rest then tag 100

[edit]
lab@R7# commit
commit complete
```

R8

```
[edit]
lab@R8# set policy-options policy-statement from.rip term accept.rest then tag 100

[edit]
lab@R8# commit
commit complete
```

b) R3

Configure routing policy to advertise the DC-1 routes matched by this tag and reject the rest of the routes.

```
[edit]
lab@R3# set policy-options policy-statement to.P1 term rip.routes from protocol isis

[edit]
lab@R3# set policy-options policy-statement to.P1 term rip.routes from tag 100

[edit]
lab@R3# set policy-options policy-statement to.P1 term rip.routes then accept
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set policy-options policy-statement to.P1 term reject.rest then reject
```

Configuring import routing policy to set the well-known BGP no-export community.

```
[edit]
lab@R3# set policy-options policy-statement from.P1 from family inet

[edit]
lab@R3# set policy-options policy-statement from.P1 from protocol bgp

[edit]
lab@R3# set policy-options policy-statement from.P1 then community add no-export
```

Although it is well-known, the no-export community should be defined under the **policy-options** hierarchy as any other community configured in Junos.

```
[edit]
lab@R3# set policy-options community no-export members no-export

[edit]
lab@R3# set protocols bgp group P1 import from.P1

[edit]
lab@R3# set protocols bgp group P1 export to.P1

[edit]
lab@R3# commit
commit complete
```

c) Verify the configuration.

The well-known no-export community is applied to the P1 routes advertised from router 3 to the route-reflector.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.20.25.10 105.24.64.0/24 extensive

inet.0: 537 destinations, 1473 routes (537 active, 0 holddown, 0 hidden)
Restart Complete
* 105.24.64.0/24 (1 entry, 1 announced)
 BGP group internal type Internal
     Nexthop: Self
     Flags: Nexthop Change
     Localpref: 100
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
      AS path: [65001] 11599906 I
      Communities: no-export
```

The RIP tagged routes are correctly matched and advertised to peer P1.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.30.0.2

inet.0: 537 destinations, 1473 routes (537 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.15.0.0/24            Self                 22                 I
* 10.15.1.0/24            Self                 22                 I
* 10.15.2.0/24            Self                 22                 I
* 10.15.3.0/24            Self                 22                 I
* 10.15.4.0/24            Self                 22                 I
* 10.15.5.0/24            Self                 22                 I
* 10.15.6.0/24            Self                 22                 I
* 10.15.7.0/24            Self                 22                 I
* 10.15.8.0/24            Self                 22                 I
* 10.15.9.0/24            Self                 22                 I
* 10.15.10.0/24           Self                 22                 I
* 10.15.11.0/24           Self                 22                 I
* 10.15.12.0/24           Self                 22                 I
* 10.15.13.0/24           Self                 22                 I
* 10.15.14.0/24           Self                 22                 I
* 10.15.15.0/24           Self                 22                 I
```

18) Do not accept IPv4 routes that are part of RFC1918 and any default routes from any of the Upstream peers. Prefixes with an as-path attribute longer than 10 should also be rejected.

a) R1

Configuring routing policy to reject the RFC1918 addresses from all Upstream peers.

```
[edit]
lab@R1# set policy-options policy-statement from.U2 term reject.rfc1918 from route-
filter 10.0.0.0/8 orlonger reject

[edit]
lab@R1# set policy-options policy-statement from.U2 term reject.rfc1918 from route-
filter 172.16.0.0/12 orlonger reject

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set policy-options policy-statement from.U2 term reject.rfc1918 from route-
filter 192.168.0.0/16 orlonger reject

[edit]
lab@R1# set policy-options policy-statement from.U2 term reject.rfc1918 then reject
```

Configuring routing policy rejecting any default route.

```
[edit]
lab@R1# set policy-options policy-statement from.U2 term reject.any.default from route-
filter 0.0.0.0/0 through 0.0.0.0/32

[edit]
lab@R1# set policy-options policy-statement from.U2 term reject.any.default then reject
```

Configuring routing policy rejecting as-paths longer than 10 ASs.

```
[edit]
lab@R1# set policy-options policy-statement from.U2 term as-path.long from protocol bgp

[edit]
lab@R1# set policy-options policy-statement from.U2 term as-path.long from as-path as-
path.long

[edit]
lab@R1# set policy-options policy-statement from.U2 term as-path.long then reject
```

Configuring regular expression matching routes with as-path longer than 10 ASs.

```
[edit]
lab@R1# set policy-options as-path as-path.long ".{11,}"

[edit]
lab@R1# set protocols bgp group U2 import from.U2

[edit]
lab@R1# commit
commit complete
```

b) R5

Similar configuration is applied to router R5.

```
[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set policy-options policy-statement from.U3-v4 term reject.rfc1918 from route-
filter 10.0.0.0/8 orlonger reject

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term reject.rfc1918 from route-
filter 172.16.0.0/12 orlonger reject

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term reject.rfc1918 from route-
filter 192.168.0.0/16 orlonger reject

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term reject.rfc1918 then reject

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term reject.any.default from
route-filter 0.0.0.0/0 through 0.0.0.0/32

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term reject.any.default then
reject

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term as-path.long from protocol
bgp

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term as-path.long from as-path
as-path.long

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term as-path.long then reject

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term damping from protocol bgp

[edit]
lab@R5# set policy-options as-path as-path.long ".{11,}"

[edit]
lab@R5# set protocols bgp group U3-v4 import from.U3-v4

[edit]
lab@R5# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

c) R8

Similar configuration is applied to router R8.

```
[edit]
lab@R8# set policy-options policy-statement from.U1 term reject.rfc1918 from route-
filter 10.0.0.0/8 orlonger reject

[edit]
lab@R8# set policy-options policy-statement from.U1 term reject.rfc1918 from route-
filter 172.16.0.0/12 orlonger reject

[edit]
lab@R8# set policy-options policy-statement from.U1 term reject.rfc1918 from route-
filter 192.168.0.0/16 orlonger reject

[edit]
lab@R8# set policy-options policy-statement from.U1 term reject.rfc1918 then reject

[edit]
lab@R8# set policy-options policy-statement from.U1 term reject.any.default from route-
filter 0.0.0.0/0 through 0.0.0.0/32

[edit]
lab@R8# set policy-options policy-statement from.U1 term reject.any.default then reject

[edit]
lab@R8# set policy-options policy-statement from.U1 term as-path.long from protocol bgp

[edit]
lab@R8# set policy-options policy-statement from.U1 term as-path.long from as-path as-
path.long

[edit]
lab@R8# set policy-options policy-statement from.U1 term as-path.long then reject

[edit]
lab@R8# set policy-options policy-statement from.U1 term damping from protocol bgp

[edit]
lab@R8# set policy-options as-path as-path.long ".{11,}"

[edit]
lab@R8# set protocols bgp group U1 import from.U1

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

d) Verify the configuration.

After the configuration is committed, router R8 rejects a number of routes that do not comply with the routing policy.

```
[edit]
lab@R8# run show route hidden

inet.0: 556 destinations, 557 routes (551 active, 0 holddown, 5 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

0.0.0.0/4               [BGP ] 00:14:03, localpref 100
                          AS path: 16151 61671 I, validation-state: unverified
                        > to 144.44.0.1 via ge-0/0/5.511
0.0.0.0/5               [BGP ] 00:14:03, localpref 100
                          AS path: 16151 61671 27075 I, validation-state: unverified
                        > to 144.44.0.1 via ge-0/0/5.511
10.0.0.0/10             [BGP ] 00:14:03, localpref 100
                          AS path: 16151 33112 60177 54491 19593 62822 59959 22793 22889 I,
validation-state: unverified
                        > to 144.44.0.1 via ge-0/0/5.511
10.140.184.64/26        [BGP ] 00:14:03, localpref 100
                          AS path: 16151 61671 I, validation-state: unverified
                        > to 144.44.0.1 via ge-0/0/5.511
195.95.164.96/27        [BGP ] 00:14:03, localpref 100
                          AS path: 16151 33003 36621 41116 19115 26184 25505 42693 63075
6532 172 7632 8762 I, validation-state: unverified
                        > to 144.44.0.1 via ge-0/0/5.511
```

19) All peerings with Upstream peers must enforce more aggressive measures against flapping prefixes with a mask length equal or greater than 20. These routes must be suppressed after two flap occurrences and stay maximum suppressed for twice the default time. Prefixes with a mask length of 10 or less must not be subject to damping.

BGP dampening aims to reduce the spread of routing instability and excessive number of update messages in the network. It marks the routes ineligible as active for a certain time. The route is suppressed and not propagated further to the peers.

Junos provides a flexible way of tuning the damping criteria using damping profiles. When damping is enabled, the default parameters are applied. BGP sets an arbitrary value to all routes of 0 initially. Every time the route is withdrawn from the routing table, the value is increase with 1000 (non-user configurable value). The following 4 parameters define the damping profile:

© 2020 Juniper Networks, Inc. All rights reserved.

- Suppress – the threshold above that the route is no longer consider eligible and is suppress from advertisement. The default value is 3000.
- Reuse – the threshold below that the route can be used again. The default value is 750.
- Half-life – number of minutes after when the arbitrary value is decreased by half. The default value is 15 minutes.
- Max-suppress – maximum time the route can stay suppressed. The default value is 60 minutes.

a) R1

You need to define 2 new damping profiles to complete this task. One profile for disabling the damping completely and another for suppressing the routes more aggressively.

Enabling the damping functionality for the upstream neighbor.

```
[edit]
lab@R1# set protocols bgp group U2 damping
```

To configure a more aggressive damping profile you can lower the suppress threshold and increase the max-suppress time. Such profile will suppress the routes earlier and will keep the routes suppressed twice as long.

```
[edit]
lab@R1# set policy-options damping aggressive suppress 2000

[edit]
lab@R1# set policy-options damping aggressive max-suppress 120
```

To exclude certain routes from damping in Junos, again a profile is used with the `disable` option.

```
[edit]
lab@R1# set policy-options damping none disable
```

Damping profile is used in a routing policy as action to a route filter.

```
[edit]
lab@R1# set policy-options policy-statement from.U2 term damping from route-filter
0.0.0.0/0 prefix-length-range /20-/32 damping aggressive

[edit]
lab@R1# set policy-options policy-statement from.U2 term damping from route-filter
0.0.0.0/0 prefix-length-range /0-/10 damping none
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

b) R5

Similar configuration is applied to router R5.

```
[edit]
lab@R5# set protocols bgp group U3-v4 damping

[edit]
lab@R5# set policy-options damping aggressive suppress 2000

[edit]
lab@R5# set policy-options damping aggressive max-suppress 120


[edit]
lab@R5# set policy-options damping none disable

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term damping from route-filter
0.0.0.0/0 prefix-length-range /20-/32 damping aggressive

[edit]
lab@R5# set policy-options policy-statement from.U3-v4 term damping from route-filter
0.0.0.0/0 prefix-length-range /0-/10 damping none

[edit]
lab@R5# commit
commit complete
```

c) R8

Similar configuration is applied to router R8.

```
[edit]
lab@R8# set protocols bgp group U1 damping

[edit]
lab@R8# set policy-options damping aggressive suppress 2000

[edit]
lab@R8# set policy-options damping aggressive max-suppress 120
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set policy-options damping none disable

[edit]
lab@R8# set policy-options policy-statement from.U1 term damping from route-filter
0.0.0.0/0 prefix-length-range /20-/32 damping aggressive

[edit]
lab@R8# set policy-options policy-statement from.U1 term damping from route-filter
0.0.0.0/0 prefix-length-range /0-/10 damping none

[edit]
lab@R8# commit
commit complete
```

20) No hidden (except filtered by applied policies) or unresolved routes should be present in AS23.89. No static routes are allowed if not explicitly permitted by a specific task. Ensure optimal routing in the network.

a) Verify the received routes on R1.

Earlier on, router R8 was configured to advertise a higher local-preference for all routes received from U1. That was done to complete a task requiring router R8 to attract all the traffic for all Upstream peers. However, if you carefully examine the bgp summary table on router R1, you will find it prefers the routes advertised by the local U2 peer. You may also notice the number of routes received from router R6 is less than expected.

```
[edit]
lab@R1# run show bgp summary
Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     949        489          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
144.44.0.101          16151        769       1425       0       0     2:56:46
405/460/405/0         0/0/0/0
172.20.25.2           65002        493       4279       0       0     3:42:24 0/0/0/0
0/0/0/0
172.20.25.5           65002       6376       4272       0       0     3:39:16
67/472/472/0          0/0/0/0
172.20.25.6           65002       4629       4270       0       0     3:38:37
17/17/17/0            0/0/0/0
```


© 2020 Juniper Networks, Inc. All rights reserved.

Following the path of the routes back to router R8, reveals that router R4 is receiving the Upstream routes from the route-reflector router. However, for some reason it does not advertise them to router R6.

b) Verifying the routes received on R8 from U1.

If you examine the routes advertised by peer U1 to router R8 you will see that the routes are assigned with community no-export-subconfed. This is another well-known community preventing the exporting routes from one sub-AS to another.

```
[edit]
lab@R8# run show route receive-protocol bgp 144.44.0.1 extensive

inet.0: 556 destinations, 573 routes (551 active, 0 holddown, 5 hidden)
Restart Complete
* 1.64.0.0/10 (1 entry, 1 announced)
     Accepted
     Nexthop: 144.44.0.1
     AS path: 16151 61671 I
     Communities: no-export-subconfed

* 1.84.160.0/20 (1 entry, 1 announced)
     Accepted
     Nexthop: 144.44.0.1
     AS path: 16151 33112 I
     Communities: no-export-subconfed
...
```

c) R8

Applying import policy to U1 neighbor on router R8 to delete the no-export-subconfed community resolves the problem.

```
[edit]
lab@R8# set policy-options policy-statement delete.no-export-subconfed term 1 from
protocol bgp

[edit]
lab@R8# set policy-options policy-statement delete.no-export-subconfed term 1 then
community delete no-export-subconfed

[edit]
lab@R8# set policy-options community no-export-subconfed members no-export-subconfed

[edit]
lab@R8# set protocols bgp group U1 import delete.no-export-subconfed
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# commit
commit complete
```

After the change is committed, router R1 prefers the Upstream routes advertised from router R5 due to the higher local-preference.

d) Verify the configuration.

```
lab@R1# run show bgp summary
Groups: 2 Peers: 4 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       949        489          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
144.44.0.101          16151        777       1450       0       0     3:00:07
0/460/405/0           0/0/0/0
172.20.25.2           65002        500       5295       0       0     3:45:45 0/0/0/0
0/0/0/0
172.20.25.5           65002       8220       5287       0       0     3:42:37
315/315/315/0         0/0/0/0
172.20.25.6           65002       6382       5285       0       0     3:41:58
174/174/174/0         0/0/0/0
```

e) Verifying the BGP next-hop on R3.

If you look carefully at the routing table on R3 and more precisely the next-hop of routes reflected by the route-reflector you may notice that it is set to the route-reflector loopback address. By default, the route-reflector should not change the next-hop attribute. It only reflects the routes avoiding steering the data traffic to itself. Attracting the traffic through the route-reflector would lead to suboptimal routing.

```
[edit]
lab@R3# run show route extensive

Restart Complete
1.64.0.0/10 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.64.0.0/10 -> {indirect(1048582)}
Page 0 idx 0, (group sub-AS65002 type External) Type 1 val 0xdda9d80 (adv_entry)
   Advertised metrics:
     Flags: Nexthop Change
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
     Nexthop: Self
     Localpref: 200
     AS path: [65001] 16151 61671 I
     Communities:
Path 1.64.0.0 from 172.20.25.10 Vector len 4.  Val: 0
        *BGP    Preference: 170/-201
                Next hop type: Indirect, Next hop index: 0
                Address: 0xda47c50
                Next-hop reference count: 912
                Source: 172.20.25.10
                Next hop type: Router, Next hop index: 618
                Next hop: 172.20.20.34 via ge-0/0/4.39, selected
                Session Id: 0x1d6
                Protocol next hop: 172.20.25.10
                Indirect next hop: 0xc2a0300 1048582 INH Session ID: 0x1de
                State: <Active Int Ext>
                Local AS: 65001 Peer AS: 65001
                Age: 2:14       Metric2: 10
                Validation State: unverified
                ORR Generation-ID: 0
                Task: BGP_65001.172.20.25.10
                Announcement bits (3): 0-KRT 2-BGP_RT_Background 3-Resolve tree 1
                AS path: 16151 61671 I  (Originator)
                Cluster list:  1.1.1.1
                Originator ID: 172.20.25.8
                Accepted
                Localpref: 200
                Router ID: 172.20.25.10
                Indirect next hops: 1
                        Protocol next hop: 172.20.25.10 Metric: 10
                        Indirect next hop: 0xc2a0300 1048582 INH Session ID: 0x1de
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 172.20.20.34 via ge-0/0/4.39
                                Session Id: 0x1d6
                                172.20.25.10/32 Originating RIB: inet.0
                                  Metric: 10    Node path count: 1
                                  Forwarding nexthops: 1
                                        Nexthop: 172.20.20.34 via ge-0/0/4.39
                                        Session Id: 1d6
...
```

f) Fixing the BGP next hop.

The reason why the route-reflector is changing the next-hop attribute is the group applied with initial configuration displayed below. The next-hop self routing policy is configured deliberately under the neighbors R3 and R4.

© 2020 Juniper Networks, Inc. All rights reserved.

```
groups {
    initial {
        protocols {
            bgp {
                group <*> {
                    neighbor 172.20.25.4 {
                        export nhs;
                    }
                    neighbor 172.20.25.3 {
                        export nhs;
                    }
                }
            }
    }
}
```

Deleting the group applied by the initial configuration restores the default next-hop on routers R3 and R4.

```
[edit]
lab@route-reflector# delete protocols bgp apply-groups initial

[edit]
lab@route-reflector# commit
commit complete
```

g) Verify the configuration.

```
[edit]
lab@R3# run show route extensive
inet.0: 532 destinations, 548 routes (532 active, 0 holddown, 0 hidden)
Restart Complete
1.64.0.0/10 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.64.0.0/10 -> {indirect(1048576)}
Page 0 idx 0, (group sub-AS65002 type External) Type 1 val 0xda7ad54 (adv_entry)
   Advertised metrics:
     Flags: Nexthop Change
     Nexthop: Self
     Localpref: 200
     AS path: [65001] 16151 61671 I
     Communities:
Path 1.64.0.0 from 172.20.25.10 Vector len 4.  Val: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
*BGP    Preference: 170/-201
        Next hop type: Indirect, Next hop index: 0
        Address: 0xda5b210
        Next-hop reference count: 912
        Source: 172.20.25.10
        Next hop type: Router, Next hop index: 609
        Next hop: 172.20.20.2 via ae0.0, selected
        Session Id: 0x1d5
        Protocol next hop: 172.20.25.8
        Indirect next hop: 0xc29be00 1048576 INH Session ID: 0x1e7
        State: <Active Int Ext>
        Local AS: 65001 Peer AS: 65001
        Age: 1:32       Metric2: 20
        Validation State: unverified
        ORR Generation-ID: 0
        Task: BGP_65001.172.20.25.10
        Announcement bits (3): 0-KRT 2-BGP_RT_Background 3-Resolve tree 1
        AS path: 16151 61671 I  (Originator)
        Cluster list:  1.1.1.1
        Originator ID: 172.20.25.8
        Accepted
        Localpref: 200
        Router ID: 172.20.25.10
        Indirect next hops: 1
                Protocol next hop: 172.20.25.8 Metric: 20
                Indirect next hop: 0xc29be00 1048576 INH Session ID: 0x1e7
                Indirect path forwarding next hops: 1
                        Next hop type: Router
                        Next hop: 172.20.20.2 via ae0.0
                        Session Id: 0x1d5
                        172.20.25.8/32 Originating RIB: inet.0
                          Metric: 20    Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 172.20.20.2 via ae0.0
                                Session Id: 1d5
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 4: MPLS Configuration

The following part is focused on configuring MPLS to provide a building block for L3VPN and L2VPN overlay services. It is highly recommended to read the whole part before starting to solve the tasks.

1) Configure RSVP on all routers except R2 and R8. The aggregated interface between R3 and R4 must not allow more than 1Gbps to be utilized by the LSPs.

a) All devices except R2 and R8

Enabling RSVP protocol is straightforward, simply configuring the interfaces under the `protocols rsvp` hierarchy. In addition, you must enable the `family mpls` for the interface you have add the interface under the `protocol mpls` hierarchy also.

Configuring router R1 interface under RSVP and MPLS protocol.

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.15

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.16

[edit]
lab@R1# set protocols mpls interface ge-0/0/4.16

[edit]
lab@R1# set protocols mpls interface ge-0/0/4.15

[edit]
lab@R1# commit
commit complete
```

Configuring router R3 interface under RSVP and MPLS protocol.

```
[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.35

[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.37
```


© 2020 Juniper Networks, Inc. All rights reserved.

RSVP subscription for interface ae0.0 is configured for 50% of the bandwidth.

```
[edit]
lab@R3# set protocols rsvp interface ae0.0 subscription 50

[edit]
lab@R3# set protocols mpls interface ae0.0

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.35

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.37

[edit]
lab@R3# commit
commit complete
```

Configuring router R4 interface under RSVP and MPLS protocol.

```
[edit]
lab@R4# set protocols rsvp interface ae0.0 subscription 50

[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.45

[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.46

[edit]
lab@R4# set protocols rsvp interface ge-0/0/4.47

[edit]
lab@R4# set protocols mpls interface ae0.0

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.45

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.46

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.47

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.48
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# commit
commit complete
```

Configuring router R5 interface under RSVP and MPLS protocol.

```
[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.15

[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.35

[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.45

[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.56

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.15

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.35

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.45

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.56

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.52

[edit]
lab@R5# commit
commit complete
```

Configuring router R6 interface under RSVP and MPLS protocol.

```
[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.16

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.46

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.56
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# set protocols mpls interface ge-0/0/4.16

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.46

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.56

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.62

[edit]
lab@R6# commit
commit complete
```

Configuring router R7 interface under RSVP and MPLS protocol.

```
[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.37

[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.47

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.37

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.47

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.78

[edit]
lab@R7# commit
commit complete
```

b) Verify the configuration.

The RSVP interface table displays the interface enabled for RSVP protocol. The extensive view of the table is displayed below.

```
[edit]
lab@R3# run show rsvp interface extensive
RSVP interface: 4 active
ae0.0 Index 338, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, LinkProtection
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  HelloInterval 9(second)
  Address 172.20.20.1
  ActiveResv 5, PreemptionCnt 0, Update threshold 10%
  Subscription 50%,
  bc0 = ct0, StaticBW 2Gbps
  ct0: StaticBW 2Gbps, AvailableBW 1000Mbps
    MaxAvailableBW 1000Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off

fxp0.0 Index 4, State Dis/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 10.10.1.3
  Protection: Off

ge-0/0/4.35 Index 332, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.20.20.13
  ActiveResv 8, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 1000Mbps
    MaxAvailableBW 1000Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off

ge-0/0/4.37 Index 333, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 172.20.20.5
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 1000Mbps
    MaxAvailableBW 1000Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off
...
```

```
[edit]
lab@R4# run show rsvp interface extensive
RSVP interface: 4 active
ae0.0 Index 340, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.20.20.2
  ActiveResv 6, PreemptionCnt 0, Update threshold 10%
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Subscription 100%,
  bc0 = ct0, StaticBW 2Gbps
  ct0: StaticBW 2Gbps, AvailableBW 2Gbps
    MaxAvailableBW 2Gbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off

fxp0.0 Index 4, State Dis/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 10.10.1.4
  Protection: Off

ge-0/0/4.45 Index 332, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.20.20.17
  ActiveResv 5, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 1000Mbps
    MaxAvailableBW 1000Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off

ge-0/0/4.46 Index 333, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.20.20.21
  ActiveResv 5, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 1000Mbps
    MaxAvailableBW 1000Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off

ge-0/0/4.47 Index 334, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.20.20.9
  ActiveResv 6, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 1000Mbps
    MaxAvailableBW 1000Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

RSVP neighbor table is displayed on routers R1 and R5. However, the neighbor sessions will be established when an LSP is signaled.

```
[edit]
lab@R1# run show rsvp neighbor
RSVP neighbor: 0 learned
```

```
[edit]
lab@R5# run show rsvp neighbor
RSVP neighbor: 3 learned
Address              Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
172.20.20.65            0  1/0     1:47:39        9   981/981  2140
172.20.20.17            0  1/0     1:51:25        9   811/811  2144
172.20.20.13            5  1/0     1:34:45        9   731/731  2976
```

2) Ensure that the RSVP messages between R1, R5, and R6 are authenticated with password key *jncie-sp*.

a) R1

RSVP authentication is configured on the interfaces required by the Task.

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.16 authentication-key
"$9$6fseCpBW87Vb2cyMX7VsY5QF3Cu"

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.15 authentication-key
"$9$vX/8xdDjqf5F24UHqfQzcyle87"
```

b) R5

```
[edit]
lab@R5# set protocols rsvp interface ge-0/0/4.15 authentication-key "$9$u8rlBRS-
ds4oGvWNbs4JZ69AtBE"
```

c) R6

```
[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.16 authentication-key
"$9$yx/eMXaJDkqfVwoGDkmP1RhceW"
```


© 2020 Juniper Networks, Inc. All rights reserved.

d) Verify the configuration.

Displaying again the extensive RSVP interface table, shows authentication enabled.

```
[edit]
lab@R1# run show rsvp interface extensive
RSVP interface: 3 active
fxp0.0 Index 4, State Dis/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 10.10.1.1
  Protection: Off

ge-0/0/4.15 Index 337, State Ena/Up
  Authentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.20.20.65
  ActiveResv 4, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 1000Mbps
    MaxAvailableBW 1000Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off

ge-0/0/4.16 Index 338, State Ena/Up
  Authentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 172.20.20.69
  ActiveResv 2, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 1000Mbps
    MaxAvailableBW 1000Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: Off
...
```

3) Configure the appropriate colors of the links according the diagram.

Administrative groups or link coloring are manually assigned to the interfaces and used for constraining LSP's path throughout the network. The name of the groups and the values must be consistent on all devices from the MPLS domain. Then the information is exchanged from IS-IS and stored in the TED table.

© 2020 Juniper Networks, Inc. All rights reserved.

a) R3

First, you must define the groups with the values assigned to each of it. Then to assign the one or more groups to the MPLS interfaces.

```
[edit]
lab@R3# set protocols mpls admin-groups GOLD 1

[edit]
lab@R3# set protocols mpls admin-groups SILVER 2

[edit]
lab@R3# set protocols mpls admin-groups BRONZE 3

[edit]
lab@R3# set protocols mpls interface ae0.0 admin-group SILVER

[edit]
lab@R3# set protocols mpls interface ae0.0 admin-group BRONZE

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.35 admin-group SILVER

[edit]
lab@R3# set protocols mpls interface ge-0/0/4.37 admin-group BRONZE

[edit]
lab@R3# commit
commit complete
```

b) R4

```
[edit]
lab@R4# set protocols mpls admin-groups GOLD 1

[edit]
lab@R4# set protocols mpls admin-groups SILVER 2

[edit]
lab@R4# set protocols mpls admin-groups BRONZE 3

[edit]
lab@R4# set protocols mpls interface ae0.0 admin-group SILVER

[edit]
lab@R4# set protocols mpls interface ae0.0 admin-group BRONZE
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set protocols mpls interface ge-0/0/4.45 admin-group GOLD

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.46 admin-group BRONZE

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.47 admin-group GOLD

[edit]
lab@R4# commit
commit complete
```

c) R5

```
[edit]
lab@R5# set protocols mpls admin-groups GOLD 1

[edit]
lab@R5# set protocols mpls admin-groups SILVER 2

[edit]
lab@R5# set protocols mpls admin-groups BRONZE 3

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.35 admin-group SILVER

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.45 admin-group GOLD

[edit]
lab@R5# commit
commit complete
```

d) R6

```
[edit]
lab@R6# set protocols mpls admin-groups GOLD 1

[edit]
lab@R6# set protocols mpls admin-groups SILVER 2

[edit]
lab@R6# set protocols mpls admin-groups BRONZE 3

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.46 admin-group BRONZE
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# commit
commit complete
```

e) R7

```
[edit]
lab@R7# set protocols mpls admin-groups GOLD 1

[edit]
lab@R7# set protocols mpls admin-groups SILVER 2

[edit]
lab@R7# set protocols mpls admin-groups BRONZE 3

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.37 admin-group BRONZE

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.47 admin-group GOLD

[edit]
lab@R7# commit
commit complete
```

f) Verify the configuration.

The assigned administrative group and the state of the MPLS enabled interfaces are checked with the commands below.

```
[edit]
lab@R3# run show mpls interface
Interface           State       Administrative groups (x: extended)
ge-0/0/4.35         Up          SILVER
ge-0/0/4.37         Up          BRONZE
ae0.0               Up          BRONZE
                                SILVER
```

```
[edit]
lab@R4# run show mpls interface
Interface           State       Administrative groups (x: extended)
ge-0/0/4.45         Up          GOLD
ge-0/0/4.46         Up          BRONZE
ge-0/0/4.47         Up          GOLD
ge-0/0/4.48         Up          <none>
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
ae0.0                  Up            BRONZE
                                     SILVER
```

```
[edit]
lab@R5# run show mpls interface
Interface              State          Administrative groups (x: extended)
ge-0/0/4.15            Up            <none>
ge-0/0/4.35            Up             SILVER
ge-0/0/4.45            Up             GOLD
ge-0/0/4.52            Up            <none>
ge-0/0/4.56            Up            <none>
```

```
[edit]
lab@R6# run show mpls interface
Interface              State          Administrative groups (x: extended)
ge-0/0/4.16            Up            <none>
ge-0/0/4.46            Up             BRONZE
ge-0/0/4.56            Up            <none>
ge-0/0/4.62            Up            <none>
```

```
[edit]
lab@R7# run show mpls interface
Interface              State          Administrative groups (x: extended)
ge-0/0/4.37            Up             BRONZE
ge-0/0/4.47            Up             GOLD
ge-0/0/4.78            Up            <none>
```

Unlike OSPF, IS-IS exchanges traffic-engineering information automatically without explicit configuration. The information is stored in the traffic-engineering database. It is then used by the constrained shortest path first algorithm (cspf) used for building paths through the MPLS network bypassing the standard routing decision.

```
[edit]
lab@R5# run show ted database
TED database: 18 ISIS nodes 8 INET nodes
ID                             Type Age(s) LnkIn LnkOut Protocol
R2.00(172.20.25.2)             Rtr     526     2      0 IS-IS(1)
R2.02                          Net     542     0      2 IS-IS(1)
    To: R5.00(172.20.25.5), Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: R2.00(172.20.25.2), Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                             Type Age(s) LnkIn LnkOut Protocol
R2.03                          Net     535     0      2 IS-IS(1)
    To: R2.00(172.20.25.2), Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    To: R6.00(172.20.25.6), Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
R3.00(172.20.25.3)            Rtr      92     4      3 IS-IS(2)
    To: R5.00(172.20.25.5), Local: 172.20.20.13, Remote: 172.20.20.14
      Local interface index: 333, Remote interface index: 334
    To: R7.02, Local: 172.20.20.5, Remote: 0.0.0.0
      Local interface index: 334, Remote interface index: 0
    To: R4.05, Local: 172.20.20.1, Remote: 0.0.0.0
      Local interface index: 336, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
R3.02                         Net     291     0      2 IS-IS(2)
    To: R3.00(172.20.25.3), Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
    To: route-reflector.00(172.20.25.10), Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
R4.00(172.20.25.4)            Rtr      78     6      4 IS-IS(2)
    To: R5.00(172.20.25.5), Local: 172.20.20.17, Remote: 172.20.20.18
      Local interface index: 333, Remote interface index: 0
    To: R4.04, Local: 172.20.20.9, Remote: 0.0.0.0
      Local interface index: 335, Remote interface index: 0
    To: R6.00(172.20.25.6), Local: 172.20.20.21, Remote: 172.20.20.22
      Local interface index: 334, Remote interface index: 0
    To: R4.05, Local: 172.20.20.2, Remote: 0.0.0.0
      Local interface index: 338, Remote interface index: 0
ID                            Type Age(s) LnkIn LnkOut Protocol
...
```

4) Configure the following LSPs using the criteria below:

a) Configure bidirectional LSPs between R1 and R3, named *r1-to-r3* and *r3-to-r1*. Make sure the shortest path between the endpoints is used. Configure automatic protection of the link between routers R3 and R5.

b) R1

First, you must configure the destination egress PE address, a mandatory setting for point-to-point LSPs. Configure the name of the LSP as it is given in the task.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r3 to 172.20.25.3
```

The constrained shortest path first algorithm should be disabled for this LSP, because the TED database on router R1 is not populated since it is running only OSPF.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r3 no-cspf
```

Although the cspf is not enabled on the ingress PE, you can still signal fast-reroute protection. This will ensure protection on the link between routers R3 and R5 as required.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r3 fast-reroute
```

Configuring primary path for the LSP, forcing the LSP path to be established through router R5.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r3 primary r5-r3

[edit]
lab@R1# set protocols mpls path r5-r3 172.20.25.5 loose

[edit]
lab@R1# set protocols mpls path r5-r3 172.20.25.3 loose

[edit]
lab@R1# commit
commit complete
```

c) R3

Configuring the LSP in the other direction from router R3 to R1.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r1 to 172.20.25.1

[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r1 no-cspf

[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r1 fast-reroute
```

Again, the primary path is configured to force the LSP through router R5.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r1 primary r5-r1

[edit]
lab@R3# set protocols mpls path r5-r1 172.20.25.5 loose

[edit]
lab@R3# set protocols mpls path r5-r1 172.20.25.1 loose

[edit]
lab@R3# commit
commit complete
```

d) Verify the configuration.

After the configuration is committed, verify the LSPs correct operation. Extensive view of the LSP shows the last few log messages of the LSP operation.

```
[edit]
lab@R1# run show mpls lsp extensive
Ingress LSP: 4 sessions

172.20.25.3
  From: 172.20.25.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r3
  ActivePath: r5-r3 (primary)
  FastReroute desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   r5-r3              State: Up
    Priorities: 7 0
    OptimizeTimer: 7200
    SmartOptimizeTimer: 180
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
         172.20.20.66(flag=1) 172.20.20.13
   29 Dec  4 17:25:43.903 Record Route:  172.20.20.66(flag=1) 172.20.20.13
   28 Dec  4 17:25:39.989 Selected as active path
   27 Dec  4 17:25:39.988 Record Route:  172.20.20.66 172.20.20.13
   26 Dec  4 17:25:39.988 Up
...
```

```
[edit]
lab@R3# run show mpls lsp extensive
Ingress LSP: 9 sessions
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
172.20.25.1
  From: 172.20.25.3, State: Up, ActiveRoute: 0, LSPname: r3-to-r1
  ActivePath: r5-r1 (primary)
  FastReroute desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   r5-r1              State: Up
    Priorities: 7 0
    OptimizeTimer: 7200
    SmartOptimizeTimer: 180
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.20.20.14(flag=1) 172.20.20.65
   39 Dec  4 16:52:01.991 Fast-reroute Detour Up
   38 Dec  4 16:52:01.339 Fast-reroute Detour Down
   37 Dec  4 16:51:31.359 Record Route:  172.20.20.14(flag=1) 172.20.20.65
...
```

The extensive view of the transit LSPs on router R5 displays the detour LSPs created by fast-reroute functionality. Detours are calculated using the cspf algorithm on routers R5 and R3.

```
[edit]
lab@R5# run show mpls lsp transit extensive

172.20.25.3
  From: 172.20.25.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r1-to-r3, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: 300135, Recovery label sent: 3
  Resv style: 1 FF, Label in: 300135, Label out: 3
  Time left:  130, Since: R1 Dec  4 15:00:20 2016
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 53854 protocol 0
  FastReroute desired
  PATH rcvfrom: 172.20.20.65 (ge-0/0/4.15) 139 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 172.20.20.13 (ge-0/0/4.35) 112 pkts
  RESV rcvfrom: 172.20.20.13 (ge-0/0/4.35) 165 pkts, Entropy label: Yes
  Explct route: 172.20.20.13 172.20.25.3
  Record route: 172.20.20.65 <self> 172.20.20.13
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Detour adspec: received MTU 1500 sent MTU 1500
    Path MTU: received 1500
    Detour PATH sentto: 172.20.20.17 (ge-0/0/4.45) 110 pkts
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Detour RESV rcvfrom: 172.20.20.17 (ge-0/0/4.45) 133 pkts, Entropy label: Yes
    Detour Explct route: 172.20.20.17 172.20.20.1
    Detour Record route: 172.20.20.65 <self> 172.20.20.17 172.20.20.1
    Detour Label out: 300308

172.20.25.1
  From: 172.20.25.3, LSPstate: Up, ActiveRoute: 0
  LSPname: r3-to-r1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: 300130, Recovery label sent: 3
  Resv style: 1 FF, Label in: 300130, Label out: 3
  Time left:  131, Since: R1 Dec  4 12:49:12 2016
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 16981 protocol 0
  FastReroute desired
  PATH rcvfrom: 172.20.20.13 (ge-0/0/4.35) 374 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 172.20.20.65 (ge-0/0/4.15) 293 pkts
  RESV rcvfrom: 172.20.20.65 (ge-0/0/4.15) 369 pkts, Entropy label: Yes
  Explct route: 172.20.20.65 172.20.25.1
  Record route: 172.20.20.13 <self> 172.20.20.65
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Detour adspec: received MTU 1500 sent MTU 1500
    Path MTU: received 1500
    Detour PATH sentto: 172.20.20.17 (ge-0/0/4.45) 296 pkts
    Detour RESV rcvfrom: 172.20.20.17 (ge-0/0/4.45) 270 pkts, Entropy label: Yes
    Detour Explct route: 172.20.20.17 172.20.20.22 172.20.20.69
    Detour Record route: 172.20.20.13 <self> 172.20.20.17 172.20.20.22 172.20.20.69
    Detour Label out: 300302
...
```

b) Configure bi-directional LSPs between R1 and R7, named *r1-to-r7* and *r7-to-r1*. Use two diverse paths, one that goes through R5 and the other through R6. By default, the path through R6 must be used as active. Make sure both paths are signaled and ready for use at any the time. The LSPs should not revert to the path that has failed. The two paths must be able to share the same bandwidth.

a) R1

The unidirectional pair of LSPs between the routers R1 and R7 should be protected by secondary paths active all the time.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r7 to 172.20.25.7
```


© 2020 Juniper Networks, Inc. All rights reserved.

Again, CSPF should be disabled since the router R1 is not part of the same TED domain and does not have the traffic-engineering information.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r7 no-cspf
```

Adaptive option will allow the two active paths to share the same link through the network without double bandwidth reservation.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r7 adaptive
```

Configuring both paths as secondary will accomplish the requirement preventing the paths to revert to the previous path. First configure the path that must be used by default as active.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r7 secondary r6-r4-r7 standby

[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r7 secondary r6-r4-r7 select
manual

[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r7 secondary r5-r3-r7 standby
```

Configuring the two diverse paths through the network.

```
[edit]
lab@R1# set protocols mpls path r6-r4-r7 172.20.25.6 loose

[edit]
lab@R1# set protocols mpls path r6-r4-r7 172.20.25.4 loose

[edit]
lab@R1# set protocols mpls path r5-r3-r7 172.20.25.5 loose

[edit]
lab@R1# set protocols mpls path r5-r3-r7 172.20.25.3 loose

[edit]
lab@R1# commit
commit complete
```

b) R7

© 2020 Juniper Networks, Inc. All rights reserved.

Configuring the LSP in the other direction from router R7 to R1.

```
[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r1 to 172.20.25.1

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r1 no-cspf

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r1 secondary r3-r5-r1 standby

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r1 secondary r3-r5-r1 select
manual

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r1 secondary r4-r6-r1 standby

[edit]
lab@R7# set protocols mpls path r3-r5-r1 172.20.25.3 loose

[edit]
lab@R7# set protocols mpls path r3-r5-r1 172.20.25.5 loose

[edit]
lab@R7# set protocols mpls path r4-r6-r1 172.20.25.4 loose

[edit]
lab@R7# set protocols mpls path r4-r6-r1 172.20.25.6 loose

[edit]
lab@R7# commit
commit complete
```

c) Verify the configuration.

After the configuration is applied the LSPs are established following the diverse paths.

```
[edit]
lab@R1# run show mpls lsp extensive
Ingress LSP: 4 sessions
...
172.20.25.7
  From: 172.20.25.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r7
  ActivePath: r6-r4-r7 (secondary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Standby   r5-r3-r7          State: Up
    Priorities: 7 0
    OptimizeTimer: 7200
    SmartOptimizeTimer: 180
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.20.20.66 172.20.20.13 172.20.20.2 172.20.20.10
   671 Dec  4 19:00:08.799 Record Route:  172.20.20.66 172.20.20.13 172.20.20.2
172.20.20.10
   670 Dec  4 19:00:08.798 Up
   669 Dec  4 19:00:08.798 Stats related identifier changed
   668 Dec  4 19:00:04.382 ResvTear received
   667 Dec  4 19:00:04.381 172.20.20.65: Down
   666 Dec  4 19:00:04.376 172.20.20.13: No Route toward dest
   665 Dec  4 18:53:14.959 Record Route:  172.20.20.66 172.20.20.13 172.20.20.2
 *Standby   r6-r4-r7          State: Up
    Priorities: 7 0
    OptimizeTimer: 7200
    SmartOptimizeTimer: 180
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.20.20.70 172.20.20.21 172.20.20.10
   25 Dec  4 17:25:49.626 Selected as active path: due to 'select manual'
   24 Dec  4 17:25:47.618 Record Route:  172.20.20.70 172.20.20.21 172.20.20.10
   23 Dec  4 17:25:47.618 Up
   22 Dec  4 17:25:47.618 Stats related identifier changed
    1 Dec  4 12:42:55.882 Originate Call
  Created: R1 Dec  4 12:42:55 2016
Total 4 displayed, Up 4, Down 0
...
```

c) Configure LSP from R5 to R3, named r5-to-r3, with secondary path built over links with colors GOLD and BRONZE, but avoiding links with color SILVER. The LSP in the opposite direction, named *r3-to-r5*, should have a backup path with not more than two hops.

a) R3

Configuring the LSP from routers R3 to R5 with two diverse paths. The primary path is not restricted and follows the best path calculated by the IGP.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r5 to 172.20.25.5

[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r5 ldp-tunneling
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r5 primary pri-to-r5
```

The secondary path is restricted in terms of hops. It must be not cross more than 2 hops.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r5 secondary sec-to-r5 hop-limit 2
```

Configuring two empty paths to define the 2 diverse paths.

```
[edit]
lab@R3# set protocols mpls path sec-to-r5

[edit]
lab@R3# set protocols mpls path pri-to-r5

[edit]
lab@R3# commit
commit complete
```

b) R5

The secondary path of the LSP from router R5 and R3 should be built over links with specific colors only.

```
[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r3 to 172.20.25.3

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r3 primary pri-to-r3

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r3 secondary sec-to-r3 admin-group
include-any GOLD

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r3 secondary sec-to-r3 admin-group
include-any BRONZE
```

The secondary path must avoid links with color SILVER.

```
[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r3 secondary sec-to-r3 admin-group
exclude SILVER
```

Configuring two empty paths to define the 2 diverse paths.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set protocols mpls path pri-to-r3

[edit]
lab@R5# set protocols mpls path sec-to-r3

[edit]
lab@R5# commit
commit complete
```

c) Verify the configuration.

After committing the configuration, the LSPs are successfully created honoring the restrictions.

```
[edit]
lab@R3# run show mpls lsp extensive
Ingress LSP: 9 sessions
...
172.20.25.5
  From: 172.20.25.3, State: Up, ActiveRoute: 0, LSPname: r3-to-r5
  ActivePath: pri-to-r5 (primary)
  LSPtype: Static Configured, Ultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   pri-to-r5          State: Up
    Priorities: 7 0
    OptimizeTimer: 7200
    SmartOptimizeTimer: 180
    Reoptimization in 1693 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 10)
 172.20.20.14 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
         172.20.20.14
    7 Dec  4 17:33:24.403 CSPF: computation result ignored, new path no benefit[4
times]
    6 Dec  4 12:49:00.282 Selected as active path: due to 'standby'
    5 Dec  4 12:49:00.281 Record Route:  172.20.20.14
    4 Dec  4 12:49:00.281 Up
    3 Dec  4 12:49:00.276 Originate Call
    2 Dec  4 12:49:00.276 CSPF: computation result accepted  172.20.20.14
    1 Dec  4 12:48:30.666 CSPF failed: no route toward 172.20.25.5[19 times]
  Secondary sec-to-r5          State: Dn
    Priorities: 7 0
    Hoplimit: 2
    OptimizeTimer: 7200
    SmartOptimizeTimer: 180
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
        No computed ERO.
    6 Dec  4 12:49:30.144 Clear Call
    5 Dec  4 12:49:00.281 Record Route:  172.20.20.14
    4 Dec  4 12:49:00.281 Up
    3 Dec  4 12:49:00.276 Originate Call
    2 Dec  4 12:49:00.276 CSPF: computation result accepted  172.20.20.14
    1 Dec  4 12:48:30.666 CSPF failed: no route toward 172.20.25.5[18 times]
  Created: R1 Dec  4 12:29:30 2016
...
```

```
[edit]
lab@R5# run show mpls lsp extensive
Ingress LSP: 5 sessions

172.20.25.3
  From: 172.20.25.5, State: Up, ActiveRoute: 0, LSPname: r5-to-r3
  ActivePath: pri-to-r3 (primary)
  LSPtype: Static Configured, Ultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   pri-to-r3         State: Up
    Priorities: 7 0
    OptimizeTimer: 7200
    SmartOptimizeTimer: 180
    Reoptimization in 1475 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 10)
 172.20.20.13 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.20.20.13
   18 Dec  4 14:44:25.050 CSPF: computation result accepted  172.20.20.13
   17 Dec  4 14:44:25.050 CSPF: Reroute due to re-optimization
   16 Dec  4 12:50:24.920 Selected as active path: due to 'primary'
   15 Dec  4 12:49:24.150 Deselected as active
   14 Dec  4 12:49:23.533 CSPF: computation result accepted  172.20.20.13
   13 Dec  4 12:49:23.532 CSPF: link down/deleted:
172.20.20.14(172.20.25.5:333)(R5.00/172.20.25.5)-
>0.0.0.0(172.20.25.3:332)(R3.00/172.20.25.3)
   12 Dec  4 12:49:21.520 CSPF: computation result accepted  172.20.20.13
   11 Dec  4 12:49:21.520 CSPF: Reroute due to re-optimization
   10 Dec  4 12:46:57.673 Selected as active path: due to 'primary'
    9 Dec  4 12:46:13.750 Deselected as active
    8 Dec  4 12:45:56.888 CSPF: computation result accepted  172.20.20.13
    7 Dec  4 12:45:56.888 CSPF: link down/deleted:
172.20.20.14(172.20.25.5:333)(R5.00/172.20.25.5)-
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
>172.20.20.13(172.20.25.3:332)(R3.00/172.20.25.3)
    6 Dec  4 11:43:14.949 Selected as active path
    5 Dec  4 11:43:14.948 Record Route:  172.20.20.13
    4 Dec  4 11:43:14.948 Up
    3 Dec  4 11:43:14.943 Originate Call
    2 Dec  4 11:43:14.943 CSPF: computation result accepted  172.20.20.13
    1 Dec  4 11:38:23.863 CSPF failed: no route toward 172.20.25.3[8 times]
  Secondary sec-to-r3          State: Dn
    Priorities: 7 0
    OptimizeTimer: 7200
    SmartOptimizeTimer: 180
          Include Any: BRONZE GOLD    Exclude: SILVER
        No computed ERO.
   20 Dec  4 12:51:52.078 Clear Call
   19 Dec  4 12:50:24.920 Deselected as active: due to 'primary'
   18 Dec  4 12:49:24.150 Selected as active path
   17 Dec  4 12:49:23.557 Record Route:  172.20.20.17 172.20.20.10 172.20.20.5
   16 Dec  4 12:49:23.556 Up
   15 Dec  4 12:49:23.539 Originate Call
   14 Dec  4 12:49:23.539 CSPF: computation result accepted  172.20.20.17 172.20.20.10
172.20.20.5
   13 Dec  4 12:48:25.473 Clear Call
   12 Dec  4 12:46:57.673 Deselected as active: due to 'primary'
   11 Dec  4 12:46:13.750 Selected as active path
   10 Dec  4 12:45:56.991 Record Route:  172.20.20.17 172.20.20.10 172.20.20.5
    9 Dec  4 12:45:56.990 Up
    8 Dec  4 12:45:56.895 Originate Call
    7 Dec  4 12:45:56.895 CSPF: computation result accepted  172.20.20.17 172.20.20.10
172.20.20.5
    6 Dec  4 11:43:44.393 Clear Call
    5 Dec  4 11:43:14.963 Record Route:  172.20.20.17 172.20.20.10 172.20.20.5
    4 Dec  4 11:43:14.963 Up
    3 Dec  4 11:43:14.943 Originate Call
    2 Dec  4 11:43:14.943 CSPF: computation result accepted  172.20.20.17 172.20.20.10
172.20.20.5
    1 Dec  4 11:38:23.863 CSPF failed: no route toward 172.20.25.3[7 times]
...
```

d) Configure LSPs from router R3 to router R7 and in the opposite direction from R7 to R3, named *r3-to-r7* and *r3-to-r7*. Configure link-protection to automatically protect the LSPs, but make sure R5 and R6 are not used as transit nodes.

a) R3

The next pair of LSPs are between routers R3 and R7 and should be protected by automatic link-protection mechanism. The protected path must avoid the link between routers R5 and R6. To solve this you can restrict the protected path in terms of hops. Such

© 2020 Juniper Networks, Inc. All rights reserved.

restriction cannot be used with fast-reroute, therefore you have to use bypass LSP leveraged by link-protection.

The bypass LSP is not allowed to stretch more than 2 hops away from the point of protection.

```
[edit]
lab@R3# set protocols rsvp interface ge-0/0/4.37 link-protection hop-limit 2

[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r7 to 172.20.25.7
```

Link protection is configured for the LSP.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r7 node-link-protection

[edit]
lab@R3# commit
commit complete
```

b) R7

The LSP in the opposite direction is configured following the same criteria.

```
[edit]
lab@R7# set protocols rsvp interface ge-0/0/4.37 link-protection hop-limit 2

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r3 to 172.20.25.3

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r3 node-link-protection

[edit]
lab@R7# commit
commit complete
```

c) Verify the configuration.

The LSPs in both directions between routers R3 and R7 are created successfully.

```
[edit]
lab@R3# run show mpls lsp extensive
...
172.20.25.7
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 From: 172.20.25.3, State: Up, ActiveRoute: 0, LSPname: r3-to-r7
 ActivePath:  (primary)
 Node/Link protection desired
 LSPtype: Static Configured, Penultimate hop popping
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                    State: Up
   Priorities: 7 0
   OptimizeTimer: 7200
   SmartOptimizeTimer: 180
   Reoptimization in 4176 second(s).
   Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
172.20.20.14 S 172.20.20.17 S 172.20.20.10 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
         172.20.25.5(flag=0x20) 172.20.20.14(Label=300291) 172.20.25.4(flag=0x20)
172.20.20.17(Label=300554) 172.20.25.7(flag=0x20) 172.20.20.10(Label=3)
   231 Dec  4 18:22:04.700 Make-before-break: Cleaned up old instance: Hold dead expiry
   230 Dec  4 18:20:53.230 Make-before-break: Switched to new instance
   229 Dec  4 18:20:53.229 Record Route:  172.20.25.5(flag=0x20)
172.20.20.14(Label=300291) 172.20.25.4(flag=0x20) 172.20.20.17(Label=300554)
172.20.25.7(flag=0x20) 172.20.20.10(Label=3)
   228 Dec  4 18:20:53.229 Up
   227 Dec  4 18:20:53.023 Pending path switchover, skip CSPF run
   226 Dec  4 18:20:53.023 CSPF: link down/deleted: 0.0.0.0(0.0.0.0:0)(R4.02/0.0.0.0)-
>0.0.0.0(172.20.25.4:0)(R4.00/172.20.25.4)
   225 Dec  4 18:20:52.992 Pending path switchover, skip CSPF run
   224 Dec  4 18:20:52.992 CSPF: link down/deleted:
172.20.20.1(172.20.25.3:338)(R3.00/172.20.25.3)->0.0.0.0(0.0.0.0:0)(R4.02/0.0.0.0)
   223 Dec  4 18:20:52.965 Originate make-before-break call
   222 Dec  4 18:20:52.965 CSPF: computation result accepted  172.20.20.14 172.20.20.17
172.20.20.10
   221 Dec  4 18:20:52.943 172.20.20.1: Tunnel local repaired
   220 Dec  4 18:20:52.943 172.20.20.1: Down
   219 Dec  4 18:18:39.113 Make-before-break: Cleaned up old instance: Hold dead expiry
   218 Dec  4 18:17:40.635 Link-protection Up
   217 Dec  4 18:17:38.812 Make-before-break: Switched to new instance
  Created: R1 Dec  4 12:29:30 2016
...
```

```
[edit]
lab@R7# run show mpls lsp extensive
...
172.20.25.3
  From: 172.20.25.7, State: Up, ActiveRoute: 0, LSPname: r7-to-r3
  ActivePath:  (primary)
  Node/Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary                    State: Up
    Priorities: 7 0
    OptimizeTimer: 7200
    SmartOptimizeTimer: 180
    Reoptimization in 2380 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
 172.20.20.9 S 172.20.20.18 S 172.20.20.13 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
          172.20.25.4(flag=0x20) 172.20.20.9(Label=300313) 172.20.25.5(flag=0x20)
172.20.20.18(Label=300142) 172.20.25.3(flag=0x20) 172.20.20.13(Label=3)
   113 Dec  4 14:26:14.484 Link-protection Up
   112 Dec  4 14:26:05.451 Link-protection Down
   111 Dec  4 14:25:54.820 Make-before-break: Cleaned up old instance: Hold dead expiry
   110 Dec  4 14:24:50.683 Link-protection Up
   109 Dec  4 14:24:48.749 Make-before-break: Switched to new instance
   108 Dec  4 14:24:47.747 Record Route:  172.20.25.4(flag=0x20)
172.20.20.9(Label=300313) 172.20.25.5(flag=0x20) 172.20.20.18(Label=300142)
172.20.25.3(flag=0x20) 172.20.20.13(Label=3)
   107 Dec  4 14:24:47.747 Up
   106 Dec  4 14:24:47.622 Originate make-before-break call
   105 Dec  4 14:24:47.622 CSPF: computation result accepted  172.20.20.9 172.20.20.18
172.20.20.13
   104 Dec  4 14:24:47.622 CSPF: link down/deleted: 0.0.0.0(0.0.0.0:0)(R4.02/0.0.0.0)-
>0.0.0.0(172.20.25.3:0)(R3.00/172.20.25.3)
   103 Dec  4 14:24:47.578 Link-protection Down
   102 Dec  4 14:24:37.898 Link-protection Up
   101 Dec  4 14:24:27.961 Selected as active path
   100 Dec  4 14:24:27.960 Record Route:  172.20.25.4(flag=0x20)
172.20.20.9(Label=300309) 172.20.25.3(flag=0x20) 172.20.20.1(Label=3)
  Created: R1 Dec  4 11:09:21 2016
...
```

The bypass RSVP sessions are created and ready for protecting the LSP between routers R3 and R7.

```
[edit]
lab@R3# run show rsvp session
Ingress RSVP: 10 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.20.25.4     172.20.25.3     Up       0  1 SE       -   300287 Bypass->172.20.20.2
```

```
[edit]
lab@R7# run show rsvp session
Ingress RSVP: 10 sessions
To              From            State   Rt Style Labelin Labelout LSPname
172.20.25.4     172.20.25.7     Up       0  1 SE       -   300500 Bypass->172.20.20.9
```


© 2020 Juniper Networks, Inc. All rights reserved.

5) Configure all RSVP enable routers to periodically optimize the path that the LSPs take through the network. The period should be set to two hours.

a) All devices with RSVP protocol enabled.

By default, Junos does not recalculate the LSP path after it is initially built. If topology changes occur making the path suboptimal, configuring optimize-timer forces the LSP path to be optimized periodically. The optimize-timer is useful only for LSPs configured that make use of the CSPF algorithm.

Repeat this on the R1, R3, R4, R5, R6, and R7 devices.

```
[edit]
lab@R1# set protocols mpls optimize-timer 7200

[edit]
lab@R1# commit
commit complete
```

b) Verify the configuration.

Optimize-timer change is visible in the extensive view of the LSP.

```
[edit]
lab@R1# run show mpls lsp extensive
Ingress LSP: 4 sessions

172.20.25.3
  From: 172.20.25.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r3
  ActivePath: r5-r3 (primary)
  FastReroute desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   r5-r3              State: Up
    Priorities: 7 0
    OptimizeTimer: 7200
...
```

6) Routers R8 and R2 support only LDP protocol for label distribution. Connect the LDP enabled routers with rest of the routers, except router R1, in such way that they can provide MPLS VPN service at a later stage.

© 2020 Juniper Networks, Inc. All rights reserved.

The task description speculates that routers R8 and R2 support only LDP protocol for label distribution. In order to enable services deployed over MPLS between the two routers through the RSVP enabled part of the network, you can configure routers R4, R7, R5, and R6 to transport the LDP traffic in RSVP signaled LSPs. You must configure an LDP targeted session between the border routers of the RSVP enabled network. Direct LDP session are then extending the LDP between routers R2 and R8.

a) R2

Enabling LDP on the interfaces on R2 towards R5 and R6. In addition to configuring `family mpls` on the interface, already done by apply-group, you must configure the interfaces under the `protocol mpls`.

```
[edit]
lab@R2# set protocols ldp interface ge-0/0/4.52

[edit]
lab@R2# set protocols ldp interface ge-0/0/4.62

[edit]
lab@R2# set protocols mpls interface ge-0/0/4.52

[edit]
lab@R2# set protocols mpls interface ge-0/0/4.62

[edit]
lab@R2# commit
commit complete
```

b) R8

```
[edit]
lab@R8# set protocols ldp interface ge-0/0/4.48

[edit]
lab@R8# set protocols ldp interface ge-0/0/4.78

[edit]
lab@R8# set protocols mpls interface ge-0/0/4.48

[edit]
lab@R8# set protocols mpls interface ge-0/0/4.78

[edit]
lab@R8# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

c) R5

```
[edit]
lab@R5# set protocols ldp interface ge-0/0/4.52

[edit]
lab@R5# set protocols ldp interface ge-0/0/4.56

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.56

[edit]
lab@R5# set protocols mpls interface ge-0/0/4.52
```

Targeted LDP sessions are established between the loopback IP addresses and are tunneled into the LSP signaled over RSVP. Therefore, the loopback interface is enabled under the LDP protocol hierarchy on all routers.

```
[edit]
lab@R5# set protocols ldp interface lo0.0

[edit]
lab@R5# commit
commit complete
```

d) R6

```
[edit]
lab@R6# set protocols ldp interface ge-0/0/4.56

[edit]
lab@R6# set protocols ldp interface ge-0/0/4.62

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.56

[edit]
lab@R6# set protocols mpls interface ge-0/0/4.62

[edit]
lab@R6# set protocols ldp interface lo0.0

[edit]
lab@R6# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

e) R4

```
[edit]
lab@R4# set protocols ldp interface ge-0/0/4.48

[edit]
lab@R4# set protocols mpls interface ge-0/0/4.48

[edit]
lab@R4# set protocols ldp interface lo0.0

[edit]
lab@R4# commit
commit complete
```

f) R7

```
[edit]
lab@R7# set protocols ldp interface ge-0/0/4.78

[edit]
lab@R7# set protocols mpls interface ge-0/0/4.78

[edit]
lab@R7# set protocols ldp interface lo0.0

[edit]
lab@R7# commit
commit complete
```

g) R3

R3 needs to have LDP tunneling over RSVP in order to provide MPLS connectivity for the L2VPN and L3VPN having sites on the LDP only routers R2 and R8.

```
[edit]
lab@R3# set protocols ldp interface lo0.0

[edit]
lab@R3# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

h) Verify the configuration.

Before LDP tunneling is configured only the direct LDP session are established on R2 and R8.

```
[edit]
lab@R2# run show ldp neighbor
Address              Interface          Label space ID          Hold time
172.20.20.78         ge-0/0/4.52        172.20.25.5:0             10
172.20.20.82         ge-0/0/4.62        172.20.25.6:0             10
```

```
[edit]
lab@R8# run show ldp neighbor
Address              Interface          Label space ID          Hold time
172.20.20.25         ge-0/0/4.48        172.20.25.4:0             13
172.20.20.29         ge-0/0/4.78        172.20.25.7:0             12
```

7) Configure all the needed LSPs to accomplish redundant communication.

The task requires configuring full-mesh of LDP tunneled LSPs.

a) R3

Because some of the required LSPs were already configured earlier, you only need to enable LDP tunneling over them. Simply add ldp-tunneling to achieve this under the specific LSPs.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r5 ldp-tunneling

[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r7 ldp-tunneling
```

The rest of the LSP should be configure as new.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r4 to 172.20.25.4

[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r4 ldp-tunneling

[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r4 no-cspf

[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r6 to 172.20.25.6
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r6 ldp-tunneling

[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r6 no-cspf

[edit]
lab@R3# commit
commit complete
```

b) R4

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r7 to 172.20.25.7

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r7 ldp-tunneling

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r7 no-cspf

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r3 to 172.20.25.3

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r3 ldp-tunneling

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r3 no-cspf

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r6 to 172.20.25.6

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r6 ldp-tunneling

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r6 no-cspf

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r5 to 172.20.25.5

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r5 ldp-tunneling

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r5 no-cspf
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# commit
commit complete
```

c) R7

```
[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r3 ldp-tunneling

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r4 to 172.20.25.4

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r4 ldp-tunneling

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r4 no-cspf

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r5 to 172.20.25.5

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r5 ldp-tunneling

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r5 no-cspf

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r6 to 172.20.25.6

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r6 ldp-tunneling

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r6 no-cspf

[edit]
lab@R7# commit
commit complete
```

d) R5

```
[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r3 ldp-tunneling

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set protocols mpls label-switched-path r5-to-r7 to 172.20.25.7

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r7 ldp-tunneling

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r7 no-cspf

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r4 to 172.20.25.4

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r4 ldp-tunneling

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r4 no-cspf

[edit]
lab@R5# commit
commit complete
```

e) R6

```
[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r4 to 172.20.25.4

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r4 ldp-tunneling

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r4 no-cspf

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r3 to 172.20.25.3

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r3 ldp-tunneling

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r3 no-cspf

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r7 to 172.20.25.7

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r7 ldp-tunneling

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r7 no-cspf
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R6# commit
commit complete
```

- Verify the configuration.

After committing the configuration on all routers, labels are exchanged between routers R2 and R8 using LDP.

```
[edit]
lab@R2# run show ldp database
Input label database, 172.20.25.2:0--172.20.25.5:0
Labels received: 5
  Label     Prefix
 299778      172.20.25.2/32
 299865      172.20.25.4/32
      3      172.20.25.5/32
 300128      172.20.25.7/32
 300133      172.20.25.8/32

Output label database, 172.20.25.2:0--172.20.25.5:0
Labels advertised: 7
  Label     Prefix
      3      172.20.25.2/32
 299783      172.20.25.3/32
 299785      172.20.25.4/32
 299777      172.20.25.5/32
 299789      172.20.25.6/32
 299787      172.20.25.7/32
 299788      172.20.25.8/32

Input label database, 172.20.25.2:0--172.20.25.6:0
Labels received: 7
  Label     Prefix
 299777      172.20.25.2/32
 299798      172.20.25.3/32
 299802      172.20.25.4/32
 299776      172.20.25.5/32
      3      172.20.25.6/32
 299807      172.20.25.7/32
 299812      172.20.25.8/32

Output label database, 172.20.25.2:0--172.20.25.6:0
Labels advertised: 7
  Label     Prefix
      3      172.20.25.2/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 299783      172.20.25.3/32
 299785      172.20.25.4/32
 299777      172.20.25.5/32
 299789      172.20.25.6/32
 299787      172.20.25.7/32
 299788      172.20.25.8/32
```

```
[edit]
lab@R2# run show route table inet.3

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.3/32     *[LDP/9] 00:52:17, metric 1
                    > to 172.20.20.82 via ge-0/0/4.62, Push 299798
172.20.25.4/32     *[LDP/9] 00:52:17, metric 1
                    > to 172.20.20.82 via ge-0/0/4.62, Push 299802
172.20.25.5/32     *[LDP/9] 00:45:47, metric 1
                    > to 172.20.20.82 via ge-0/0/4.62, Push 299776
172.20.25.6/32     *[LDP/9] 00:51:03, metric 1
                    > to 172.20.20.82 via ge-0/0/4.62
172.20.25.7/32     *[LDP/9] 00:52:17, metric 1
                    > to 172.20.20.82 via ge-0/0/4.62, Push 299807
172.20.25.8/32     *[LDP/9] 00:52:17, metric 1
                    > to 172.20.20.82 via ge-0/0/4.62, Push 299812
```

```
[edit]
lab@R8# run show ldp database
Input label database, 172.20.25.8:0--172.20.25.4:0
Labels received: 7
  Label     Prefix
 299889      172.20.25.2/32
 299890      172.20.25.3/32
      3      172.20.25.4/32
 299876      172.20.25.5/32
 299888      172.20.25.6/32
 300299      172.20.25.7/32
 300305      172.20.25.8/32

Output label database, 172.20.25.8:0--172.20.25.4:0
Labels advertised: 7
  Label     Prefix
 299791      172.20.25.2/32
 299792      172.20.25.3/32
 299788      172.20.25.4/32
 299789      172.20.25.5/32
 299790      172.20.25.6/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
 299787      172.20.25.7/32
      3      172.20.25.8/32

Input label database, 172.20.25.8:0--172.20.25.7:0
Labels received: 7
  Label     Prefix
 299803      172.20.25.2/32
 299804      172.20.25.3/32
 299801      172.20.25.4/32
 299806      172.20.25.5/32
 299802      172.20.25.6/32
      3      172.20.25.7/32
 299805      172.20.25.8/32

Output label database, 172.20.25.8:0--172.20.25.7:0
Labels advertised: 7
  Label     Prefix
 299791      172.20.25.2/32
 299792      172.20.25.3/32
 299788      172.20.25.4/32
 299789      172.20.25.5/32
 299790      172.20.25.6/32
 299787      172.20.25.7/32
      3      172.20.25.8/32
```

```
[edit]
lab@R8# run show route table inet.3

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.2/32     *[LDP/9] 01:22:26, metric 1
                    > to 172.20.20.25 via ge-0/0/4.48, Push 299889
172.20.25.3/32     *[LDP/9] 01:22:26, metric 1
                    > to 172.20.20.25 via ge-0/0/4.48, Push 299890
172.20.25.4/32     *[LDP/9] 01:22:26, metric 1
                    > to 172.20.20.25 via ge-0/0/4.48
172.20.25.5/32     *[LDP/9] 01:22:26, metric 1
                    > to 172.20.20.25 via ge-0/0/4.48, Push 299876
172.20.25.6/32     *[LDP/9] 01:22:26, metric 1
                    > to 172.20.20.25 via ge-0/0/4.48, Push 299888
172.20.25.7/32     *[LDP/9] 01:22:32, metric 1
                    > to 172.20.20.29 via ge-0/0/4.78
```

8) Ensure that all the MPLS enabled devices do not reveal nodes/hops inside the core.

© 2020 Juniper Networks, Inc. All rights reserved.

There are two ways to decrement the TTL of the packet traversing the MPLS network in Junos. Knowing the difference between both ways is the key solving the task.

- no-decrement-ttl – this method applies only to RSVP signaled LSP. Configured on the ingress PE globally or per LSP, a proprietary RSVP object is signaled and the whole LSP is single hop in the IP packet. The egress PE is shown in the traceroute output using this method.
- no-propagate-ttl – this method applies to all, RSVP or LDP signaled. It works in multivendor environments and can be configured only globally. None of the MPLS nodes is shown in the traceroute output.

a) All devices

Because the task requires hiding the MPLS network on all routes, no-propagate-ttl will accomplish this for RSVP and LDP signaled LSPs.

```
[edit]
lab@R1# set protocols mpls no-propagate-ttl

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 5: IPv6 Configuration

The next part focuses on enabling IPV6 traffic. It is highly recommended to read the whole part before starting to solve the tasks.

1) Ensure that Customers C1 and C2 can communicate using IPv6 over your network. Configure BGP to exchange the Customer routes using a single session for IPv4 and IPv6.
2) The local summary IPv6 route should be also advertised to the C1 and C2. However, neither P1 nor U3 should be able to communicate with the Customers using IPv6.

This task requires enabling native IPv6 communication through the network. IBGP sessions are carrying both protocols already in sub-AS 65001 and on routers R5 and R6 in sub-AS 65002. To complete the task, you must configure routing policies controlling the IPv6 traffic to and from customers C1 and C2. In same time, the IPv6 communication between peers P1 and U3 should be restricted.

a) R8

The export routing policy to.C2 is amended to match and advertise the C1 routes on router R8.

```
[edit]
lab@R8# set policy-options policy-statement to.C2 term C1.routes from protocol bgp

[edit]
lab@R8# set policy-options policy-statement to.C2 term C1.routes from rib inet6.0
```

An as-path regular expression is used to match the routes advertised by peer C1 in rib inet6.0.

```
[edit]
lab@R8# set policy-options policy-statement to.C2 term C1.routes from as-path C1.origin

[edit]
lab@R8# set policy-options policy-statement to.C2 term C1.routes then accept

[edit]
lab@R8# set policy-options as-path C1.origin ".* 65120"
```

Local IPv6 summary routes is advertised to peer C2 also.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set policy-options policy-statement to.C2 term local.aggregate from protocol
bgp

[edit]
lab@R8# set policy-options policy-statement to.C2 term local.aggregate from route-
filter 2001:fedc::/64 exact

[edit]
lab@R8# set policy-options policy-statement to.C2 term local.aggregate then accept
```

The rest of the routes are rejected and not sent to peer C2.

```
[edit]
lab@R8# set policy-options policy-statement to.C2 term reject.rest-v6 from rib inet6.0

[edit]
lab@R8# set policy-options policy-statement to.C2 term reject.rest-v6 then reject

[edit]
lab@R8# set protocols bgp group C2 export to.C2

[edit]
lab@R8# commit
commit complete
```

b) R6

Similar configuration is applied to router R6 related to the C2 routes advertised to peer C1.

```
[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term C2.routes from protocol bgp

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term C2.routes from rib inet6.0

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term C2.routes from as-path
C2.origin

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term C2.routes then accept

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term local.aggregate from protocol
bgp

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R6# set policy-options policy-statement to.C1-v6 term local.aggregate from route-
filter 2001:fedc::/64 exact

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term local.aggregate then accept

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term reject.rest-v6 from rib
inet6.0

[edit]
lab@R6# set policy-options policy-statement to.C1-v6 term reject.rest-v6 then reject

[edit]
lab@R6# set policy-options as-path C2.origin ".* 65123"

[edit]
lab@R6# commit
commit complete
```

c) Verify the configuration.

After committing the changes to the export routing policies on routers R6 and R8, C1 and C2 routes are exchanged through the network.

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.20.200.2 2001:101:10:32::/64

inet6.0: 149 destinations, 166 routes (149 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:101:10:32::/64     Self                                    65120 I
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.20.200.2 2001:fedc::/64

inet6.0: 149 destinations, 166 routes (149 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  2001:fedc::/64          Self                                    I
```

```
[edit]
lab@R8# run show route advertising-protocol bgp 144.44.0.1 2001:101:10:32::/64
```


© 2020 Juniper Networks, Inc. All rights reserved.

Note: On R6 you must get the actual peer address from `show bgp summary` because this peering is using the link local address for this session.

```
[edit]
lab@R6# run show route advertising-protocol bgp fe80::206:a01:9a0e:908
2001:10:15:132::/64

inet6.0: 60 destinations, 85 routes (60 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
* 2001:10:15:132::/64     Self                                    65123 I
```

```
[edit]
lab@R6# run show route advertising-protocol bgp fe80::206:a01:9a0e:908 2001:fedc::/64

inet6.0: 60 destinations, 85 routes (60 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                  Nexthop              MED     Lclpref    AS path
  2001:fedc::/64          Self                                    I
```

3) To protect you customers, tunnel the IPv6 traffic between peers P1 and U3 in an LSP.
4) You are not allowed to add additional BGP families besides the ones enabled so far. Ensure that the transport label is not stripped off from the label stack, one hop before the egress PE.

A common way to connect IPv6 islands through IPv4 network is using 6PE method. It leverages the BGP protocol already running in the core of the network. However, the task restricts you from adding new families to the BGP protocol, which means you cannot use this approach.

The IS-IS protocol provides a mechanism to shortcut the next hop of the routes through RSVP LSP egress points. When this feature is enabled, the LSP appears like a directly connected logical interface towards the remote network used by CSPF. This effectively connects IPv6 island through the IS-IS enabled networks.

The data path works in similar way to the 6PE method. The ingress PE pushes an additional explicit null IPv6 label as first label in the stack and on top also a transport label is added.

a) R3

Enabling additional IS-IS topology for IPv6 earlier is useful for IS-IS IPv6 shortcuts also. This is required to ensure the IPv6 routes are following the route of the LSP between the

© 2020 Juniper Networks, Inc. All rights reserved.

ingress and egress PEs. The IPv6 routes must be resolved via the same interface where the LSP is initiated. Enabling a diverse IPv6 topology will provide easier control over IPv6 routes, without affecting IPv4 routes.

Enabling IS-IS IPv6 shortcuts feature on router R3.

```
[edit]
lab@R3# set protocols isis traffic-engineering family inet6 shortcuts
```

Ultimate hop popping is not required for the IS-IS IPv6 operation, but it is required by the task in this case. It will keep the transport label on the egress PE. Because the LSP between routers R3 and R5 used to carry IPv6 traffic is only a single hop, keeping the transport label will help to better observe transport operations in our lab.

```
[edit]
lab@R3# set protocols mpls label-switched-path r3-to-r5 ultimate-hop-popping
```

Because peer P1 uses a single BGP session you must replace the next-hop advertised by R3 with IPv6 direct address.

```
[edit]
lab@R3# set policy-options policy-statement to.P1 term U3-v6 from protocol bgp

[edit]
lab@R3# set policy-options policy-statement to.P1 term U3-v6 from rib inet6.0

[edit]
lab@R3# set policy-options policy-statement to.P1 term U3-v6 from as-path U3.adjacent

[edit]
lab@R3# set policy-options policy-statement to.P1 term U3-v6 then next-hop
2001:172:30::1

[edit]
lab@R3# set policy-options policy-statement to.P1 term U3-v6 then accept

[edit]
lab@R3# set policy-options policy-statement to.P1 term reject.rest then reject
```

To successfully match the U3 advertised routes only, you have to include the sub-AS 65002 added to the as-path.

```
[edit]
lab@R3# set policy-options as-path U3.adjacent "65002 74\.39 .*"

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R3# commit
commit complete
```

b) R5

Router R5 requires only the configuration of IS-IS IPv6 shortcuts and enabling ultimate hop popping.

```
[edit]
lab@R5# set protocols isis traffic-engineering family inet6 shortcuts

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r3 ultimate-hop-popping

[edit]
lab@R3# commit
commit complete
```

c) Verify the configuration.

After the IS-IS IPv6 shortcuts are enabled, the U3 advertised IPv6 prefixes are routed through the LSP between routers R3 and R5. In addition to the transport label, a label 2 is pushed at the bottom of the label stack.

```
[edit]
lab@R3# run show route 2001:240::/32 extensive

inet6.0: 148 destinations, 148 routes (148 active, 0 holddown, 0 hidden)
Restart Complete
2001:240::/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 2001:240::/32 -> {indirect(1048580)}
Page 0 idx 0, (group P1 type External) Type 1 val 0xbbe3b24 (adv_entry)
   Advertised metrics:
     Nexthop: Self
     AS path: [65001] 4849703 33112 60177 60637 30709 57763 I
     Communities:
Page 0 idx 1, (group sub-AS65002 type External) Type 1 val 0xbbe39f0 (adv_entry)
   Advertised metrics:
     Nexthop: 2001:fedc::5
     Localpref: 100
     AS path: [65001] (65002) 4849703 33112 60177 60637 30709 57763 I
     Communities:
    Advertise: 00000002
Page 0 idx 2, (group internal type Internal) Type 1 val 0xbbe8894 (adv_entry)
   Advertised metrics:
     Nexthop: 2001:fedc::5
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
     Localpref: 100
     AS path: [65001] (65002) 4849703 33112 60177 60637 30709 57763 I
     Communities:
Path 2001:240:: from 172.20.25.5 Vector len 4.  Val: 0 1 2
        *BGP    Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb2dc7a0
                Next-hop reference count: 176
                Source: 172.20.25.5
                Next hop type: Router, Next hop index: 1048581
                Next hop: 172.20.20.14 via ge-0/0/4.35 weight 0x1, selected
                Label-switched-path r3-to-r5
                Label operation: Push 2, Push 299783(top)
                Label TTL action: no-prop-ttl, no-prop-ttl(top)
                Load balance label: Label 2: None; Label 299783: None;
                Label element ptr: 0xaf16ae0
                Label parent element ptr: 0xaf150a0
                Label element references: 1
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x140
                Statistics ID Group: Kernel ID = 1, Stats IDs = { 536871118 }

                Protocol next hop: 2001:fedc::5
                Indirect next hop: 0xac9c2c0 1048580 INH Session ID: 0x15a
                State: <Active Int Ext>
                Local AS: 65001 Peer AS: 65002
                Age: 5:24:35    Metric2: 10
                Validation State: unverified
                ORR Generation-ID: 0
                Task: BGP_65002.172.20.25.5
                Announcement bits (3): 0-KRT 4-BGP_RT_Background 5-Resolve tree 5
                AS path: (65002) 4849703 33112 60177 60637 30709 57763 I
                Accepted
                Localpref: 100
                Router ID: 172.20.25.5
                Indirect next hops: 1
                        Protocol next hop: 2001:fedc::5 Metric: 10
                        Indirect next hop: 0xac9c2c0 1048580 INH Session ID: 0x15a
                        Indirect path forwarding next hops: 2
                                Next hop type: Router
                                Next hop: 172.20.20.14 via ge-0/0/4.35 weight 0x1
                                Session Id: 0x140
                                Statistics ID Group: Kernel ID = 1, Stats IDs = {
536871118 }
                                Next hop: 172.20.20.2 via ae0.0 weight 0x1
                                Session Id: 0x147
                        2001:fedc::5/128 Originating RIB: inet6.3
                          Metric: 10                        Node path count: 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                        Forwarding nexthops: 2
                              Nexthop: 172.20.20.14 via ge-0/0/4.35
...
```

In the opposite direction, the P1 prefixes are resolved on router R5 through LSP to R3.

```
[edit]
lab@R5# run show route table inet6.0 2001:f0f4:f691:d::/64 extensive

inet6.0: 146 destinations, 171 routes (146 active, 0 holddown, 0 hidden)
Restart Complete
2001:f0f4:f691:d::/64 (2 entries, 1 announced)
TSI:
KRT in-kernel 2001:f0f4:f691:d::/64 -> {indirect(1048576)}
Page 0 idx 0, (group U3-v6 type External) Type 1 val 0xb3ad55c (adv_entry)
   Advertised metrics:
     Nexthop: Self
     AS path: [65002] 11599906 I
     Communities:
Page 0 idx 1, (group sub-AS65001 type External) Type 1 val 0xb3aea94 (adv_entry)
   Advertised metrics:
     Nexthop: 2001:fedc::3
     Localpref: 100
     AS path: [65002] (65001) 11599906 I
     Communities:
    Advertise: 00000002
Path 2001:f0f4:f691:d:: from 172.20.25.3 Vector len 4.  Val: 0 1
        *BGP    Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb395fa0
                Next-hop reference count: 49
                Source: 172.20.25.3
                Next hop type: Router, Next hop index: 710
                Next hop: 172.20.20.13 via ge-0/0/4.35, selected
                Label-switched-path r5-to-r3
                Label operation: Push 2, Push 299799(top)
                Label TTL action: no-prop-ttl, no-prop-ttl(top)
                Load balance label: Label 2: None; Label 299799: None;
                Label element ptr: 0xaf18520
                Label parent element ptr: 0xaf16ba0
                Label element references: 2
                Label element child references: 0
                Label element lsp id: 0
                Session Id: 0x142
                Statistics ID Group: Kernel ID = 3, Stats IDs = { 536870980 }
                Protocol next hop: 2001:fedc::3
                Indirect next hop: 0xac9d940 1048576 INH Session ID: 0x158
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
State: <Active Int Ext>
Local AS: 65002 Peer AS: 65001
Age: 4:22:38    Metric2: 10
Validation State: unverified
ORR Generation-ID: 0
Task: BGP_65001.172.20.25.3
Announcement bits (3): 0-KRT 5-BGP_RT_Background 6-Resolve tree 5
AS path: (65001) 11599906 I
Accepted
Localpref: 100
Router ID: 172.20.25.3
Indirect next hops: 1
        Protocol next hop: 2001:fedc::3 Metric: 10
        Indirect next hop: 0xac9d940 1048576 INH Session ID: 0x158
        Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 172.20.20.13 via ge-0/0/4.35
                Session Id: 0x142
                Statistics ID Group: Kernel ID = 3, Stats IDs = {
536870980 }
        2001:fedc::3/128 Originating RIB: inet6.3
          Metric: 10                          Node path count: 1
          Forwarding nexthops: 1
                Nexthop: 172.20.20.13 via ge-0/0/4.35
...
```


© 2020 Juniper Networks, Inc. All rights reserved.

Task 6: MPLS VPNs Configuration

The last part is focused on the configuration of VPN services over the MPLS network. This includes L3VPNs, L2VPNs, and VPLS. It is highly recommended to read the whole part before starting to solve the tasks.

1) Configure a L3VPN connecting the sites: VPNA-CE1, VPNA-CE2, VPNA-CE3 and VPNA-CE4. Sites CE1 and CE2 are using OSPF protocol for PE-CE routing protocol, whereas the rest of the sites use EBGP with AS 64888. Use your local AS number for creating the route-distinguisher and route-target communities. Make sure you check all prefixes from each site and share them between the sites.
2) Site CE1 and CE2 have a backdoor connection. Make sure that the two sites prefer to send traffic through the L3VPN connection. The backdoor link has a configured metric of 50.

    a) All devices including the route-reflector.

You must enable BGP to support vpnv4 routes on all routers. Those are the routes extended by the route distinguisher and attached vrf-target communities.

```
[edit]
lab@R1# set protocols bgp group internal family inet-vpn unicast

[edit]
lab@R1# set protocols bgp group internal family inet unicast
```

Note: On R3, R4, R5, and R6 you will also need to apply these NLRI families to the sub-AS BGP groups.

In the first Part of the lab, a firewall filter protecting the RE was configured. An automatic prefix-list is used to support adding new BGP neighbors without the need of amending the configuration. To make the firewall filter accept BGP neighbors under routing-instances, you must configure a new prefix-list to account this too.

```
[edit]
lab@R1# set policy-options prefix-list bgp-neighbors-RI apply-path "routing-instances
<*> protocols bgp group <*> neighbor <*>"

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bgp from source-prefix-list
bgp-neighbors-RI
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

b) R3

Configure the interface connected to the VPNA-CE1 router.

```
[edit]
lab@R3# set interfaces ge-0/0/5 unit 900 description "=== connection to VPNA-CE1 ==="

[edit]
lab@R3# set interfaces ge-0/0/5 unit 900 vlan-id 900

[edit]
lab@R3# set interfaces ge-0/0/5 unit 900 family inet address 10.10.0.1/30
```

The VPNA-CE1 and VPNA-CE2 sites have a backdoor link in OSPF Area 0. The task speculates that the link must be used only when the connection to the MPLS network is not operational. However, because the MPLS core network appears as different area to the OSPF running on the CE site, the routes advertised from the MPLS network appear as external or inter-area routes. OSPF algorithm would always prefer intra-area routes to the inter-area and external routes. Using a sham-link, a feature similar to virtual-link, can be used to create a connection between two PEs. This will create an OSPF intra-area link between the two CEs through the MPLS core network making it comparable to the backdoor link.

The sham-link requires configuration of an additional unique loopback interface added to the L3VPN routing-instance. The actual virtual connection is built using that interface on both PEs. Because Junos allows only a single loopback interface per routing-instance before you commit you must add the second loopback interface to the routing-instance.

```
[edit]
lab@R3# set interfaces lo0 unit 1 family inet address 172.20.25.33/32
```

Junos requires the use of routing-instances as containers and for separation of the VPN configuration. Several different types of routing-instances are available, each is used for a specific purpose. L3VPNs are configured using routing-instance type vrf. The vrf instances require the configuration of a route-distinguisher and route-target policies applied to the routes.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set routing-instances VPNA instance-type vrf
```

Next, adding the CE interface to the VPNA routing-instance together with loopback interface needed for sham-link configuration.

```
[edit]
lab@R3# set routing-instances VPNA interface ge-0/0/5.900

[edit]
lab@R3# set routing-instances VPNA interface lo0.1
```

The route-distinguisher configured is based on the local AS number. Because it is a 4byte AS number, you must use the format below. A letter ‘L’ must be appended to the 4byte AS number.

```
[edit]
lab@R3# set routing-instances VPNA route-distinguisher 1507417L:100
```

vrf-import and vrf-export policy are configured for controlling the routes advertised and received by VPNA. A route-target community is assigned to the routes exported to the MP-BGP protocol. Only routes matching the VPNA target community are accepted by the import vrf-policy and inserted to the vrf routing table.

```
[edit]
lab@R3# set routing-instances VPNA vrf-import VPNA-target-import

[edit]
lab@R3# set routing-instances VPNA vrf-export VPNA-target-export
```

The vrf-table-label is used to program Junos using single vrf label for all routes part of the same VPN.

```
[edit]
lab@R3# set routing-instances VPNA vrf-table-label
```

Configure the OSPF protocol under the VPNA vrf routing-instance for connectivity with the VPNA-CE1 site. An export policy is needed to advertise the BGP routes to the OSPF process.

```
[edit]
lab@R3# set routing-instances VPNA protocols ospf export to.VPNA-CE1
```


© 2020 Juniper Networks, Inc. All rights reserved.

Configure the sham-link between the two PEs. The local address is the dedicated loopback interface address. The remote address is configured under Area 0 and points to the remote PE dedicated address.

```
[edit]
lab@R3# set routing-instances VPNA protocols ospf sham-link local 172.20.25.33

[edit]
lab@R3# set routing-instances VPNA protocols ospf area 0.0.0.0 sham-link-remote
172.20.25.77 metric 1

[edit]
lab@R3# set routing-instances VPNA protocols ospf area 0.0.0.0 interface ge-0/0/5.900
interface-type p2p
```

vrf-export is configured to advertise the OSPF and direct routes to the MP-BGP to the MPLS network.

```
[edit]
lab@R3# set policy-options policy-statement VPNA-target-export term 1 from protocol
ospf
```

The VPNA route-target community is attached to the routes.

```
[edit]
lab@R3# set policy-options policy-statement VPNA-target-export term 1 then community
add VPNA-target

[edit]
lab@R3# set policy-options policy-statement VPNA-target-export term 1 then accept

[edit]
lab@R3# set policy-options policy-statement VPNA-target-export term 2 from protocol
direct

[edit]
lab@R3# set policy-options policy-statement VPNA-target-export term 2 then community
add VPNA-target

[edit]
lab@R3# set policy-options policy-statement VPNA-target-export term 2 then accept
```

The vrf-import policy accepts the routes with VPNA route-target received by MP-BGP from the other PEs.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# set policy-options policy-statement VPNA-target-import term 1 from protocol bgp

[edit]
lab@R3# set policy-options policy-statement VPNA-target-import term 1 from community
VPNA-target

[edit]
lab@R3# set policy-options policy-statement VPNA-target-import term 1 then accept

[edit]
lab@R3# set policy-options policy-statement to.VPNA-CE1 term 1 from protocol bgp

[edit]
lab@R3# set policy-options policy-statement to.VPNA-CE1 term 1 then accept
```

Last, the VPNA route-target community is defined. Similar to the route-distinguisher when 4byte AS number is used, additional ‘L’ is used for creating the extended community.

```
[edit]
lab@R3# set policy-options community VPNA-target members target:1507417L:100

[edit]
lab@R3# commit
commit complete
```

c) R7

A new interface is configured to connect with the CE site.

```
[edit]
lab@R7# set interfaces ge-0/0/5 unit 901 description "=== connection to VPNA-CE2 ==="

[edit]
lab@R7# set interfaces ge-0/0/5 unit 901 vlan-id 901

[edit]
lab@R7# set interfaces ge-0/0/5 unit 901 family inet address 10.10.0.5/30
```

The sham-link requires configuration of an additional unique loopback interface added to the L3VPN routing-instance.

```
[edit]
lab@R7# set interfaces lo0 unit 1 family inet address 172.20.25.77/32
```


© 2020 Juniper Networks, Inc. All rights reserved.

The VPNA vrf routing-instance on router R7 is configured.

```
[edit]
lab@R7# set routing-instances VPNA instance-type vrf
```

Next, add the CE interface to the VPNA routing-instance together with the loopback interface needed for sham-link configuration.

```
[edit]
lab@R7# set routing-instances VPNA interface ge-0/0/5.901

[edit]
lab@R7# set routing-instances VPNA interface lo0.1

[edit]
lab@R7# set routing-instances VPNA route-distinguisher 1507417L:100

[edit]
lab@R7# set routing-instances VPNA vrf-import VPNA-target-import

[edit]
lab@R7# set routing-instances VPNA vrf-export VPNA-target-export

[edit]
lab@R7# set routing-instances VPNA vrf-table-label

[edit]
lab@R7# set routing-instances VPNA routing-options multicast ssm-groups
239.255.255.3/32
```

Configure the OSPF protocol for VPNA for connectivity with the VPNA-CE2 site. The export policy is needed to advertise the BGP routes to the OSPF process.

```
[edit]
lab@R7# set policy-options policy-statement to.VPNA-CE2 term 1 from protocol bgp

[edit]
lab@R7# set policy-options policy-statement to.VPNA-CE2 term 1 then accept

[edit]
lab@R7# set routing-instances VPNA protocols ospf export to.VPNA-CE2
```

Configure the sham-link between the two PEs:

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set routing-instances VPNA protocols ospf sham-link local 172.20.25.77

[edit]
lab@R7# set routing-instances VPNA protocols ospf area 0.0.0.0 sham-link-remote
172.20.25.33 metric 1

[edit]
lab@R7# set routing-instances VPNA protocols ospf area 0.0.0.0 interface ge-0/0/5.901
interface-type p2p

[edit]
lab@R7# set policy-options policy-statement VPNA-target-export term 1 from protocol
ospf

[edit]
lab@R7# set policy-options policy-statement VPNA-target-export term 1 then community
add VPNA-target

[edit]
lab@R7# set policy-options policy-statement VPNA-target-export term 1 then accept

[edit]
lab@R7# set policy-options policy-statement VPNA-target-export term 2 from protocol
direct

[edit]
lab@R7# set policy-options policy-statement VPNA-target-export term 2 then community
add VPNA-target

[edit]
lab@R7# set policy-options policy-statement VPNA-target-export term 2 then accept

[edit]
lab@R7# set policy-options policy-statement VPNA-target-import term 1 from protocol bgp

[edit]
lab@R7# set policy-options policy-statement VPNA-target-import term 1 from community
VPNA-target

[edit]
lab@R7# set policy-options policy-statement VPNA-target-import term 1 then accept

[edit]
lab@R7# set policy-options community VPNA-target members target:1507417L:100

[edit]
lab@R7# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

d) Verify the configuration

You should notice that the sham link is not present.

```
[edit]
lab@R3# run show ospf neighbor instance VPNA
Address          Interface              State     ID               Pri  Dead
10.10.0.2        ge-0/0/5.900           Full      10.10.100.1      128    35
```

Review the routes you are sending and receiving from the route reflector for the VPN.

```
[edit]
lab@R3# run show route advertising-protocol bgp 172.20.25.10 table VPNA.inet.0

VPNA.inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.10.0.0/30            Self                         100        I
* 10.10.0.4/30            Self                 52      100        I
* 10.10.0.96/30           Self                 51      100        I
* 10.10.100.1/32          Self                 1       100        I
* 10.10.100.2/32          Self                 51      100        I
* 10.100.0.0/24           Self                 1       100        I
* 10.100.1.0/24           Self                 1       100        I
* 10.100.2.0/24           Self                 1       100        I
* 10.100.3.0/24           Self                 1       100        I
* 10.100.4.0/24           Self                 1       100        I
* 10.100.5.0/24           Self                 1       100        I
* 10.100.6.0/24           Self                 1       100        I
* 10.100.7.0/24           Self                 1       100        I
* 10.100.8.0/24           Self                 51      100        I
* 10.100.9.0/24           Self                 51      100        I
* 10.100.10.0/24          Self                 51      100        I
* 10.100.11.0/24          Self                 51      100        I
* 10.100.12.0/24          Self                 51      100        I
* 10.100.13.0/24          Self                 51      100        I
* 10.100.14.0/24          Self                 51      100        I
* 10.100.15.0/24          Self                 51      100        I
* 172.20.25.33/32         Self                         100        I
```

```
[edit]
lab@R3# run show route receive-protocol bgp 172.20.25.10 table VPNA.inet.0

VPNA.inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
```

Review the `bgp.l3vpn.0` table on the route reflector to determine why no routes are being sent to the VPN PE devices.

© 2020 Juniper Networks, Inc. All rights reserved.

e) Route-reflector

```
[edit]
lab@route-reflector# run show route table bgp.l3vpn.0

bgp.l3vpn.0: 23 destinations, 44 routes (0 active, 0 holddown, 44 hidden)
```

Review the hidden routes in `bgp.l3vpn.0` routing table.

```
[edit]
lab@route-reflector# run show route table bgp.l3vpn.0 hidden extensive

bgp.l3vpn.0: 23 destinations, 44 routes (0 active, 0 holddown, 44 hidden)
Restart Complete
1507417L:100:10.10.0.0/30 (2 entries, 0 announced)
         BGP    Preference: 170/-101
                Route Distinguisher: 1507417L:100
                Next hop type: Unusable, Next hop index: 0
                Address: 0xbe74e50
                Next-hop reference count: 44
                State: <Hidden Int Ext ProtectionPath ProtectionCand>
                Local AS: 65001 Peer AS: 65001
                Age: 8:47       Metric: 52
                Validation State: unverified
                Task: BGP_65001.172.20.25.7
                AS path: I
                Communities: target:1507417L:100 rte-type:0.0.0.0:1:0
                Accepted
                VPN Label: 16
                Localpref: 100
                Router ID: 172.20.25.7
                Indirect next hops: 1
                        Protocol next hop: 172.20.25.7
                        Label operation: Push 16
                        Label TTL action: prop-ttl
                        Load balance label: Label 16: None;
                        Indirect next hop: 0x0 - INH Session ID: 0x0
      BGP    Preference: 170/-101
             Route Distinguisher: 1507417L:100
             Next hop type: Unusable, Next hop index: 0
             Address: 0xbe74e50
             Next-hop reference count: 44
             State: <Hidden Int Ext ProtectionPath ProtectionCand>
             Local AS: 65001 Peer AS: 65001
             Age: 8:48
             Validation State: unverified
             Task: BGP_65001.172.20.25.3
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                AS path: I
                Communities: target:1507417L:100
                Accepted
                VPN Label: 16
                Localpref: 100
                Router ID: 172.20.25.3
                Indirect next hops: 1
                        Protocol next hop: 172.20.25.3
                        Label operation: Push 16
                        Label TTL action: prop-ttl
                        Load balance label: Label 16: None;
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
...
```

Junos provides several ways to populate the `inet.3` routing table when MPLS is not enabled on the route-reflector. However, the task narrows the choices to a particular one, which is adding a static default route in the `inet.3` routing table.

Use static routing to resolve any hidden routes in RIB `inet.3` on the route-reflector.

f) Route-reflector

```
[edit]
lab@route-reflector# set routing-options rib inet.3 static route 0.0.0.0/0 reject

[edit]
lab@route-reflector# set routing-options rib inet6.3 static route 0:0:0:0:0:ffff::/96
reject

[edit]
lab@route-reflector# commit
commit complete
```

g) Verify the changes have fixed the routing problems.

The OSPF neighbor with CE1 is in Full state. The sham-link is also established successfully.

```
[edit]
lab@R3# run show ospf neighbor instance VPNA
Address          Interface              State     ID               Pri  Dead
10.10.0.2        ge-0/0/5.900           Full      10.10.100.1      128    27
172.20.25.77     shamlink.0             Full      172.20.25.77       0    32
```


© 2020 Juniper Networks, Inc. All rights reserved.

OSPF routes are received from site CE1 by router R3, whereas the routes from the remote site CE2 are received through the MPLS network. If a sham-link was not used, then the backdoor link would be preferred, and the routes advertised by CE2 would be reachable through CE1.

```
[edit]
lab@R3# run show route table VPNA.inet.0

VPNA.inet.0: 25 destinations, 35 routes (25 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/30       *[Direct/0] 01:17:47
                    > via ge-0/0/5.900
10.10.0.1/32       *[Local/0] 01:17:47
                      Local via ge-0/0/5.900
10.10.0.4/30       *[BGP/170] 00:01:19, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                    > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                      to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
10.10.0.96/30      *[OSPF/10] 01:17:38, metric 51
                    > to 10.10.0.2 via ge-0/0/5.900
10.10.100.1/32     *[OSPF/10] 01:17:38, metric 1
                    > to 10.10.0.2 via ge-0/0/5.900
10.10.100.2/32     *[BGP/170] 00:01:19, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                    > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                      to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
10.100.0.0/24      *[OSPF/10] 01:17:38, metric 1
                    > to 10.10.0.2 via ge-0/0/5.900
10.100.1.0/24      *[OSPF/10] 01:17:38, metric 1
                    > to 10.10.0.2 via ge-0/0/5.900
10.100.2.0/24      *[OSPF/10] 01:17:38, metric 1
                    > to 10.10.0.2 via ge-0/0/5.900
10.100.3.0/24      *[OSPF/10] 01:17:38, metric 1
                    > to 10.10.0.2 via ge-0/0/5.900
10.100.4.0/24      *[OSPF/10] 01:17:38, metric 1
                    > to 10.10.0.2 via ge-0/0/5.900
10.100.5.0/24      *[OSPF/10] 01:17:38, metric 1
                    > to 10.10.0.2 via ge-0/0/5.900
10.100.6.0/24      *[OSPF/10] 01:17:38, metric 1
                    > to 10.10.0.2 via ge-0/0/5.900
10.100.7.0/24      *[OSPF/10] 01:17:38, metric 1
                    > to 10.10.0.2 via ge-0/0/5.900
10.100.8.0/24      *[BGP/170] 00:01:19, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                    > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                      to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
10.100.9.0/24      *[BGP/170] 00:01:19, MED 1, localpref 100, from 172.20.25.10
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    AS path: I, validation-state: unverified
                  > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                    to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
10.100.10.0/24     *[BGP/170] 00:01:19, MED 1, localpref 100, from 172.20.25.10
                    AS path: I, validation-state: unverified
                  > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                    to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
10.100.11.0/24     *[BGP/170] 00:01:19, MED 1, localpref 100, from 172.20.25.10
                    AS path: I, validation-state: unverified
                  > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                    to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
10.100.12.0/24     *[BGP/170] 00:01:19, MED 1, localpref 100, from 172.20.25.10
                    AS path: I, validation-state: unverified
                  > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                    to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
10.100.13.0/24     *[BGP/170] 00:01:19, MED 1, localpref 100, from 172.20.25.10
                    AS path: I, validation-state: unverified
                  > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                    to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
10.100.14.0/24     *[BGP/170] 00:01:19, MED 1, localpref 100, from 172.20.25.10
                    AS path: I, validation-state: unverified
                  > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                    to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
10.100.15.0/24     *[BGP/170] 00:01:19, MED 1, localpref 100, from 172.20.25.10
                    AS path: I, validation-state: unverified
                  > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                    to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
172.20.25.33/32    *[Direct/0] 01:17:47
                  > via lo0.1
172.20.25.77/32    *[BGP/170] 00:01:19, localpref 100, from 172.20.25.10
                    AS path: I, validation-state: unverified
                    to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
                  > to 172.20.20.14 via ge-0/0/4.35, label-switched-path r3-to-r7
224.0.0.5/32       *[OSPF/10] 01:17:48, metric 1
                    MultiRecv
```

On router R7 the OSPF neighbors are in the Full state.

```
[edit]
lab@R7# run show ospf neighbor instance VPNA
Address          Interface              State     ID               Pri  Dead
10.10.0.6        ge-0/0/5.901           Full      10.10.100.2      128    31
172.20.25.33     shamlink.0             Full      172.20.25.33       0    39
```

VPNA routing table on router R7 contains the OSPF routes advertised by site CE2. The CE1 routes are reachable through router R3 across the MPLS network.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# run show route table VPNA

VPNA.inet.0: 25 destinations, 36 routes (25 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/30        *[BGP/170] 00:03:00, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                     > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.10.0.4/30        *[Direct/0] 01:15:01
                     > via ge-0/0/5.901
10.10.0.5/32        *[Local/0] 01:15:01
                      Local via ge-0/0/5.901
10.10.0.96/30       *[OSPF/10] 01:14:47, metric 51
                     > to 10.10.0.6 via ge-0/0/5.901
                     [BGP/170] 00:03:00, MED 51, localpref 100, from 172.20.25.10
                       AS path: I, validation-state: unverified
                     > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                       to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.10.100.1/32      *[BGP/170] 00:03:00, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                     > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.10.100.2/32      *[OSPF/10] 01:14:47, metric 1
                     > to 10.10.0.6 via ge-0/0/5.901
10.100.0.0/24       *[BGP/170] 00:03:00, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                     > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.100.1.0/24       *[BGP/170] 00:03:00, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                     > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.100.2.0/24       *[BGP/170] 00:03:00, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                     > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.100.3.0/24       *[BGP/170] 00:03:00, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                     > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.100.4.0/24       *[BGP/170] 00:03:00, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                     > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.100.5.0/24       *[BGP/170] 00:03:00, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                  > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.100.6.0/24       *[BGP/170] 00:03:00, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                    > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.100.7.0/24       *[BGP/170] 00:03:00, MED 1, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                    > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
10.100.8.0/24       *[OSPF/10] 01:14:47, metric 1
                    > to 10.10.0.6 via ge-0/0/5.901
10.100.9.0/24       *[OSPF/10] 01:14:47, metric 1
                    > to 10.10.0.6 via ge-0/0/5.901
10.100.10.0/24      *[OSPF/10] 01:14:47, metric 1
                    > to 10.10.0.6 via ge-0/0/5.901
10.100.11.0/24      *[OSPF/10] 01:14:47, metric 1
                    > to 10.10.0.6 via ge-0/0/5.901
10.100.12.0/24      *[OSPF/10] 01:14:47, metric 1
                    > to 10.10.0.6 via ge-0/0/5.901
10.100.13.0/24      *[OSPF/10] 01:14:47, metric 1
                    > to 10.10.0.6 via ge-0/0/5.901
10.100.14.0/24      *[OSPF/10] 01:14:47, metric 1
                    > to 10.10.0.6 via ge-0/0/5.901
10.100.15.0/24      *[OSPF/10] 01:14:47, metric 1
                    > to 10.10.0.6 via ge-0/0/5.901
172.20.25.33/32     *[BGP/170] 00:03:00, localpref 100, from 172.20.25.10
                      AS path: I, validation-state: unverified
                    > to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
                      to 172.20.20.9 via ge-0/0/4.47, label-switched-path r7-to-r3
172.20.25.77/32     *[Direct/0] 01:15:01
                    > via lo0.1
224.0.0.5/32        *[OSPF/10] 01:15:02, metric 1
                      MultiRecv

VPNA.inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

fe80::250:5603:85a9:44ce/128
                    *[Local/0] 01:14:50
                       Local via ge-0/0/5.901
fe80::250:560f:fca9:7bb1/128
                    *[Direct/0] 01:15:02
                     > via lo0.1
ff02::2/128         *[INET6/0] 01:15:02
                       MultiRecv
```


© 2020 Juniper Networks, Inc. All rights reserved.

3) Site CE3 is multihomed to both routers R1 and R8. Make sure the path through R8 is the preferred for sending and receiving traffic to CE3. You must use already applied policy to the route-reflector on R8 to complete the task.

Sites VPNA-CE3 is connected to the MPLS network on routers R1 and R8 using EBGP. You must make sure that R8 is the preferred PE for sending and receiving traffic to the site. To achieve this, you must use the existing policy on R8 applied to increase the local-preference and make R8 attractive for sending traffic to site CE3. In the direction from site CE3 you can increase the metric on the routes advertised from the MPLS network. Because the lowest metric is preferred and the default MED is 0, you must increase the MED on the routes sent by router R1 to CE3.

a) R8

Configuring the interface connected to the VPNA-CE3 router.

```
[edit]
lab@R8# set interfaces ge-0/0/5 unit 911 description "=== connection to VPNA-CE3 ==="

[edit]
lab@R8# set interfaces ge-0/0/5 unit 911 vlan-id 911

[edit]
lab@R8# set interfaces ge-0/0/5 unit 911 family inet address 10.10.0.9/30
```

Configuring the vrf routing-instance and adding the CE3 interface to it.

```
[edit]
lab@R8# set routing-instances VPNA instance-type vrf

[edit]
lab@R8# set routing-instances VPNA interface ge-0/0/5.911
```

The route-distinguisher and the vrf policies are applied to the routing-instance.

```
[edit]
lab@R8# set routing-instances VPNA route-distinguisher 1507417L:100

[edit]
lab@R8# set routing-instances VPNA vrf-import VPNA-target-import

[edit]
lab@R8# set routing-instances VPNA vrf-export VPNA-target-export
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set routing-instances VPNA vrf-table-label
```

VPNA-CE3 is using EBGP session to connect to routers R8 and R1.

```
[edit]
lab@R8# set routing-instances VPNA protocols bgp group VPNA-CE3 peer-as 64888

[edit]
lab@R8# set routing-instances VPNA protocols bgp group VPNA-CE3 neighbor 10.10.0.10
```

Next, the vrf-import and vrf-export policies are defined to advertise and receive VPNA routes

```
[edit]
lab@R8# set policy-options policy-statement VPNA-target-export term 1 from protocol bgp

[edit]
lab@R8# set policy-options policy-statement VPNA-target-export term 1 then community
add VPNA-target

[edit]
lab@R8# set policy-options policy-statement VPNA-target-export term 1 then accept

[edit]
lab@R8# set policy-options policy-statement VPNA-target-import term 1 from protocol bgp

[edit]
lab@R8# set policy-options policy-statement VPNA-target-import term 1 from community
VPNA-target

[edit]
lab@R8# set policy-options policy-statement VPNA-target-import term 1 then accept
```

The VPNA route-target community is defined under the **policy-options** hierarchy.

```
[edit]
lab@R8# set policy-options community VPNA-target members target:1507417L:100
```

By default, Junos does not apply the BGP group or neighbor policies to the vrf routes advertised to the other PEs. The vpn-apply-export knob changes the default behavior and applies both vrf export and BGP group export policies to the vrf routes.

```
[edit]
lab@R8# set protocols bgp vpn-apply-export
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# commit
commit complete
```

b) R1

Similar configuration is applied to the router R1 connecting the second interface of VPNA-CE3.

```
[edit]
lab@R1# set interfaces ge-0/0/5 unit 912 description "=== connection to VPNA-CE3 ==="

[edit]
lab@R1# set interfaces ge-0/0/5 unit 912 vlan-id 912

[edit]
lab@R1# set interfaces ge-0/0/5 unit 912 family inet address 10.10.0.13/30

[edit]
lab@R1# set routing-instances VPNA instance-type vrf

[edit]
lab@R1# set routing-instances VPNA interface ge-0/0/5.912

[edit]
lab@R1# set routing-instances VPNA route-distinguisher 1507417L:100

[edit]
lab@R1# set routing-instances VPNA vrf-import VPNA-target-import

[edit]
lab@R1# set routing-instances VPNA vrf-export VPNA-target-export

[edit]
lab@R1# set routing-instances VPNA vrf-table-label
```

Increasing the metric or MED attribute will accomplish the requirement of VPNA-CE3 to prefer sending traffic through router R8.

```
[edit]
lab@R1# set routing-instances VPNA protocols bgp group VPNA-CE3 metric-out 100

[edit]
lab@R1# set routing-instances VPNA protocols bgp group VPNA-CE3 peer-as 64888
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set routing-instances VPNA protocols bgp group VPNA-CE3 neighbor 10.10.0.14

[edit]
lab@R1# set policy-options policy-statement VPNA-target-export term 1 from protocol bgp

[edit]
lab@R1# set policy-options policy-statement VPNA-target-export term 1 then community
add VPNA-target

[edit]
lab@R1# set policy-options policy-statement VPNA-target-export term 1 then accept

[edit]
lab@R1# set policy-options policy-statement VPNA-target-export term 2 from protocol
direct

[edit]
lab@R1# set policy-options policy-statement VPNA-target-export term 2 then community
add VPNA-target

[edit]
lab@R1# set policy-options policy-statement VPNA-target-export term 2 then accept

[edit]
lab@R1# set policy-options policy-statement VPNA-target-import term 1 from protocol bgp

[edit]
lab@R1# set policy-options policy-statement VPNA-target-import term 1 from community
VPNA-target

[edit]
lab@R1# set policy-options policy-statement VPNA-target-import term 1 then accept

[edit]
lab@R1# set policy-options community VPNA-target members target:1507417L:100
```

Create a bidirectional MPLS LSP between R1 and R5 and disable CSPF calculations.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r5 to 172.20.25.5 no-cspf

[edit]
lab@R1# commit
commit complete
```

c) R5

© 2020 Juniper Networks, Inc. All rights reserved.

Configure return LSP to R1.

```
[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r1 to 172.20.25.1 no-cspf

[edit]
lab@R5# commit
commit complete
```

d) Verify the configuration.

Router R8 receives number of routes advertised by VPNA-CE3.

```
[edit]
lab@R8# run show route receive-protocol bgp 10.10.0.10

inet.0: 557 destinations, 574 routes (552 active, 0 holddown, 5 hidden)
Restart Complete

inet.3: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete

VPNA.inet.0: 33 destinations, 33 routes (33 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.100.16.0/24          10.10.0.10                              64888 I
* 10.100.17.0/24          10.10.0.10                              64888 I
* 10.100.18.0/24          10.10.0.10                              64888 I
* 10.100.19.0/24          10.10.0.10                              64888 I
* 10.100.20.0/24          10.10.0.10                              64888 I
* 10.100.21.0/24          10.10.0.10                              64888 I
* 10.100.22.0/24          10.10.0.10                              64888 I
* 10.100.23.0/24          10.10.0.10                              64888 I
...
```

The output below confirms that VPNA vrf routes are advertised to the route-reflector with local-preference of 200.

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.20.25.10 10.100.16.0/24

VPNA.inet.0: 33 destinations, 33 routes (33 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.100.16.0/24          Self                            200        64888 I
```

Router R1 is receiving the same routes from site VPNA-CE3. The CE1 and CE2 VPNA routes are advertised by router R1 to VPNA-CE3 with an increased metric of 100.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# run show route receive-protocol bgp 10.10.0.14

inet.0: 473 destinations, 473 routes (418 active, 0 holddown, 55 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

VPNA.inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.100.16.0/24          10.10.0.14                               64888 I
* 10.100.17.0/24          10.10.0.14                               64888 I
* 10.100.18.0/24          10.10.0.14                               64888 I
* 10.100.19.0/24          10.10.0.14                               64888 I
* 10.100.20.0/24          10.10.0.14                               64888 I
* 10.100.21.0/24          10.10.0.14                               64888 I
* 10.100.22.0/24          10.10.0.14                               64888 I
* 10.100.23.0/24          10.10.0.14                               64888 I
...
```

```
[edit]
lab@R1# run show route advertising-protocol bgp 10.10.0.14

VPNA.inet.0: 33 destinations, 49 routes (33 active, 0 holddown, 8 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.10.0.0/30            Self                 100                I
* 10.10.0.4/30            Self                 100                I
* 10.10.0.96/30           Self                 100                I
* 10.10.100.1/32          Self                 100                I
* 10.10.100.2/32          Self                 100                I
* 10.100.0.0/24           Self                 100                I
* 10.100.1.0/24           Self                 100                I
* 10.100.2.0/24           Self                 100                I
* 10.100.3.0/24           Self                 100                I
* 10.100.4.0/24           Self                 100                I
* 10.100.5.0/24           Self                 100                I
* 10.100.6.0/24           Self                 100                I
* 10.100.7.0/24           Self                 100                I
* 10.100.8.0/24           Self                 100                I
* 10.100.9.0/24           Self                 100                I
* 10.100.10.0/24          Self                 100                I
* 10.100.11.0/24          Self                 100                I
* 10.100.12.0/24          Self                 100                I
* 10.100.13.0/24          Self                 100                I
* 10.100.14.0/24          Self                 100                I
* 10.100.15.0/24          Self                 100                I
* 172.20.25.33/32         Self                 100                I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 172.20.25.77/32          Self                    100                 I
```

4) Site C4 is configured to peer with AS 23.890 by mistake. You must connect the site to the MPLS network. However, you are not allowed to use *local-as* command to accomplish that. Ensure that the mistaken AS number is not added to the AS path attribute on the routes send to the other sites.

The task requires connecting site CE4 to the MPLS network using EBGP, similar to site CE3. However, by mistake the CE router is configured to peer with a wrong AS number. The task restricts the use of `local-as` knob, which could make router R5 pretend it is part of the AS 23.890. Junos allows configuring the routing-instance to have a AS number different from the global AS configured. However, the global AS number is still checked for AS loop detection for all routing-instances.

a) R5

Configuring the interface connecting the VPNA-CE4 router.

```
[edit]
lab@R5# set interfaces ge-0/0/5 unit 921 description "=== connection to VPNA-CE4 ==="

[edit]
lab@R5# set interfaces ge-0/0/5 unit 921 vlan-id 921

[edit]
lab@R5# set interfaces ge-0/0/5 unit 921 family inet address 10.10.0.17/30
```

The interface is then added to a new vrf routing-instance.

```
[edit]
lab@R5# set routing-instances VPNA instance-type vrf

[edit]
lab@R5# set routing-instances VPNA interface ge-0/0/5.921
```

VPNA route-distinguisher is attached to the routing-instance. vrf-import and vrf-export policies are configured to assign and match the VPNA route-target community.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set routing-instances VPNA route-distinguisher 1507417L:100

[edit]
lab@R5# set routing-instances VPNA vrf-import VPNA-target-import

[edit]
lab@R5# set routing-instances VPNA vrf-export VPNA-target-export

[edit]
lab@R5# set routing-instances VPNA vrf-table-label
```

The routing-instance is configured with the AS number used by site CE4 by mistake. However, the task requires that the wrong AS number is not be present in the as-path when the routes are advertised to rest of the sites. Basically, only site CE3 is speaking BGP, so you must ensure that it does not receive the AS 23.890 from the MPLS network.

The purpose of the Independent-domain feature is to connect L3VPN sites with IBGP sessions transparently through the MPLS network. However, It can be used with EBGP sessions in a similar way to tunnel BGP attributes between sites CE3 and CE4. This way the MPLS network AS number including the wrong AS CE4 peers with, are hidden from the CE sites.

```
[edit]
lab@R5# set routing-instances VPNA routing-options autonomous-system 23.890
independent-domain
```

The BGP session is configured similarly to the session with CE3 on routers R1 and R8.

```
[edit]
lab@R5# set routing-instances VPNA protocols bgp group VPNA-CE4 peer-as 64888

[edit]
lab@R5# set routing-instances VPNA protocols bgp group VPNA-CE4 neighbor 10.10.0.18
```

The vrf-import and vrf-export policies are also similar to policies configured on routers R1 and R8.

```
[edit]
lab@R5# set policy-options policy-statement VPNA-target-export term 1 from protocol bgp

[edit]
lab@R5# set policy-options policy-statement VPNA-target-export term 1 then community
add VPNA-target
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement VPNA-target-export term 1 then accept

[edit]
lab@R5# set policy-options policy-statement VPNA-target-export term 2 from protocol
direct

[edit]
lab@R5# set policy-options policy-statement VPNA-target-export term 2 then community
add VPNA-target

[edit]
lab@R5# set policy-options policy-statement VPNA-target-export term 2 then accept

[edit]
lab@R5# set policy-options policy-statement VPNA-target-import term 1 from protocol bgp

[edit]
lab@R5# set policy-options policy-statement VPNA-target-import term 1 from community
VPNA-target

[edit]
lab@R5# set policy-options policy-statement VPNA-target-import term 1 then accept

[edit]
lab@R5# set policy-options community VPNA-target members target:1507417L:100

[edit]
lab@R5# commit
commit complete
```

b) Verify the configuration.

The EBGP session between router R5 and VPNA-CE4 router is successfully established. The use of local AS configuration for the VPNA routing-instance, resolves the issue with the wrong configuration on the CE router.

```
[edit]
lab@R5# run show bgp summary
Groups: 5 Peers: 8 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    1899        489          0          0          0          0
bgp.l3vpn.0
                     102         32          0          0          0          0
bgp.l2vpn.0
                      14          6          0          0          0          0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
bgp.mvpn.0
                       7          3          0          0          0          0
inet6.0
                     132        109          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.10.0.18            64888        210        271       0       0     1:52:06 Establ
  VPNA.inet.0: 8/8/8/0
...
```

Routes are received from the VPNA-CE4 site.

```
[edit]
lab@R5# run show route receive-protocol bgp 10.10.0.18

inet.0: 546 destinations, 1955 routes (541 active, 0 holddown, 5 hidden)
Restart Complete

inet.3: 39 destinations, 46 routes (5 active, 0 holddown, 39 hidden)
Restart Complete

VPNA.inet.0: 45 destinations, 115 routes (45 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.100.24.0/24          10.10.0.18                              64888 I
* 10.100.25.0/24          10.10.0.18                              64888 I
* 10.100.26.0/24          10.10.0.18                              64888 I
* 10.100.27.0/24          10.10.0.18                              64888 I
* 10.100.28.0/24          10.10.0.18                              64888 I
* 10.100.29.0/24          10.10.0.18                              64888 I
* 10.100.30.0/24          10.10.0.18                              64888 I
* 10.100.31.0/24          10.10.0.18                              64888 I
...
```

With the use of the independent-domain feature on router R5, R8 advertises the CE4 routes to CE3 site without as-path.

```
[edit]
lab@R8# run show route advertising-protocol bgp 10.10.0.10

VPNA.inet.0: 43 destinations, 43 routes (43 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.10.0.0/30                Self                                       I
* 10.10.0.4/30                Self                                       I
* 10.10.0.16/30               Self                                       I
* 10.10.0.96/30               Self                                       I
* 10.10.100.1/32              Self                                       I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 10.10.100.2/32          Self                                  I
* 10.100.0.0/24           Self                                  I
* 10.100.1.0/24           Self                                  I
* 10.100.2.0/24           Self                                  I
* 10.100.3.0/24           Self                                  I
* 10.100.4.0/24           Self                                  I
* 10.100.5.0/24           Self                                  I
* 10.100.6.0/24           Self                                  I
* 10.100.7.0/24           Self                                  I
* 10.100.8.0/24           Self                                  I
* 10.100.9.0/24           Self                                  I
* 10.100.10.0/24          Self                                  I
* 10.100.11.0/24          Self                                  I
* 10.100.12.0/24          Self                                  I
* 10.100.13.0/24          Self                                  I
* 10.100.14.0/24          Self                                  I
* 10.100.15.0/24          Self                                  I
* 10.100.24.0/24          Self                                  I
* 10.100.25.0/24          Self                                  I
* 10.100.26.0/24          Self                                  I
* 10.100.27.0/24          Self                                  I
* 10.100.28.0/24          Self                                  I
* 10.100.29.0/24          Self                                  I
* 10.100.30.0/24          Self                                  I
* 10.100.31.0/24          Self                                  I
* 172.20.25.33/32         Self                                  I
* 172.20.25.77/32         Self                                  I
```

For comparison, the output below is taken before independent-domain knob was added to router R5.

```
[edit]
lab@R8# run show route advertising-protocol bgp 10.10.0.10

VPNA.inet.0: 43 destinations, 43 routes (43 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.10.0.16/30           Self                                    I
* 10.10.0.96/30           Self                                    I
* 10.10.100.1/32          Self                                    I
* 10.10.100.2/32          Self                                    I
* 10.100.0.0/24           Self                                    I
* 10.100.1.0/24           Self                                    I
* 10.100.2.0/24           Self                                    I
* 10.100.3.0/24           Self                                    I
* 10.100.4.0/24           Self                                    I
* 10.100.5.0/24           Self                                    I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 10.100.6.0/24          Self                                  I
* 10.100.7.0/24          Self                                  I
* 10.100.8.0/24          Self                                  I
* 10.100.9.0/24          Self                                  I
* 10.100.10.0/24         Self                                  I
* 10.100.11.0/24         Self                                  I
* 10.100.12.0/24         Self                                  I
* 10.100.13.0/24         Self                                  I
* 10.100.14.0/24         Self                                  I
* 10.100.15.0/24         Self                                  I
* 10.100.24.0/24         Self                                  1508218 64888 I
* 10.100.25.0/24         Self                                  1508218 64888 I
* 10.100.26.0/24         Self                                  1508218 64888 I
* 10.100.27.0/24         Self                                  1508218 64888 I
* 10.100.28.0/24         Self                                  1508218 64888 I
* 10.100.29.0/24         Self                                  1508218 64888 I
* 10.100.30.0/24         Self                                  1508218 64888 I
* 10.100.31.0/24         Self                                  1508218 64888 I
* 172.20.25.33/32        Self                                  I
* 172.20.25.77/32        Self                                  I
```

5) Make sure the EBGP sites, CE3 and CE4, exchange routes through the core successfully without using the *as-override* command. The two sites are configured to tolerate AS loops.

By default, BGP honors the split-horizon rule and does not advertise routes learned by one EBGP peer back to any EBGP peers that are part from the same AS. On way to disable that behavior is using the `as-override` option. It will overwrite the originating AS with the MPLS core AS and the routes will be accepted by the other sites participating in the same VPN. However, it is not allowed to use this approach to solve the task as its explicitly stated that sites using EBGP, CE3 and CE4, are tolerating AS loops. This means that they could accept routes with their own AS number in the as-path attribute. Other way to disable the suppression of routes send to the same AS, is using the `advertise-peer-as` option.

a) R1

Due to the use of the independent-domain feature, the routes from CE4 are advertised with an empty as-path by router R1 and R8, hence the advertise-peer-as is not required. It is configured there as well for consistency.

```
[edit]
lab@R1# set routing-instances VPNA protocols bgp group VPNA-CE3 advertise-peer-as
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

b) R5

```
[edit]
lab@R5# set routing-instances VPNA protocols bgp group VPNA-CE4 advertise-peer-as

[edit]
lab@R5# commit
commit complete
```

c) R8

```
[edit]
lab@R8# set routing-instances VPNA protocols bgp group VPNA-CE3 advertise-peer-as

[edit]
lab@R8# commit
commit complete
```

d) Verify the configuration

Router R5 is advertising CE3 routes to the site CE4 successfully despite the two sites in configured with the same AS.

```
[edit]
lab@R5# run show route advertising-protocol bgp 10.10.0.18

VPNA.inet.0: 42 destinations, 81 routes (42 active, 0 holddown, 8 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 10.10.0.0/30            Self                                    65002 I
* 10.10.0.4/30            Self                                    65002 I
* 10.10.0.12/30           Self                                    65002 I
* 10.10.0.96/30           Self                                    65002 I
* 10.10.100.1/32          Self                                    65002 I
* 10.10.100.2/32          Self                                    65002 I
* 10.100.0.0/24           Self                                    65002 I
* 10.100.1.0/24           Self                                    65002 I
* 10.100.2.0/24           Self                                    65002 I
* 10.100.3.0/24           Self                                    65002 I
* 10.100.4.0/24           Self                                    65002 I
* 10.100.5.0/24           Self                                    65002 I
* 10.100.6.0/24           Self                                    65002 I
* 10.100.7.0/24           Self                                    65002 I
* 10.100.8.0/24           Self                                    65002 I
* 10.100.9.0/24           Self                                    65002 I
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
* 10.100.10.0/24          Self                                65002 I
* 10.100.11.0/24          Self                                65002 I
* 10.100.12.0/24          Self                                65002 I
* 10.100.13.0/24          Self                                65002 I
* 10.100.14.0/24          Self                                65002 I
* 10.100.15.0/24          Self                                65002 I
* 10.100.16.0/24          Self                                65002 64888 I
* 10.100.17.0/24          Self                                65002 64888 I
* 10.100.18.0/24          Self                                65002 64888 I
* 10.100.19.0/24          Self                                65002 64888 I
* 10.100.20.0/24          Self                                65002 64888 I
* 10.100.21.0/24          Self                                65002 64888 I
* 10.100.22.0/24          Self                                65002 64888 I
* 10.100.23.0/24          Self                                65002 64888 I
* 172.20.25.33/32         Self                                65002 I
* 172.20.25.77/32         Self                                65002 I
```

6) Configure VPNA to support multicast transmission using NG-MVPN protocol and BGP for control protocol with RSVP multicast data plane. The sources are located on the CE1 and CE2 routers. They provide source specific multicast service for group 239.255.255.3 from source IP addresses, 10.10.0.2 and 10.10.0.6.
7) Only router R1 must be used for multicast traffic received by site CE3. Static routing is configured on CE3 to prefer R1 as a gateway to the sources.

Next, you must enable VPNA for multicast support using NG-MVPN. NG-MVPN uses a PIM free core and BGP for transporting the customer PIM routes connecting the receivers with the sources. In the MPLS network, the C-PIM routes are exchanged using MP-BGP protocol. NG-MVPN, on the other hand, provides a higher flexibility. In addition to PIM/GRE supports, a few more options for the multicast data plane, including multicast transport over LSPs signaled by RSVP and mLDP. The task explicitly requires the use of RSVP for creating p2mp LSPs for transporting the multicast through the MPLS network.

a) All devices including the route reflector.

First, the IPv4 mvpn signaling family is added to MP-BGP.

Note: On R3, R4, R5, and R6 you must also apply these NLRI families to the sub-AS BGP groups.

```
[edit]
lab@R1# set protocols bgp group internal family inet-mvpn signaling
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

b) R3

The multicast sources are connected behind sites CE1 and CE2. This makes the routers R3 and R7 source sites and root of the p2mp RSVP LSP.

```
[edit]
lab@R3# set protocols mpls label-switched-path p2mp-template template
```

```
[edit]
lab@R3# set protocols mpls label-switched-path p2mp-template no-cspf
```

```
[edit]
lab@R3# set protocols mpls label-switched-path p2mp-template p2mp
```

By default, Junos reserved multicast group addresses in the range from 232.0.0.0 through 232.255.255.255 for source specific multicast. However, according the Task description the SSM group used by the Source to stream multicast traffic is out of that range. You must add the group to the range of SSM groups.

```
[edit]
lab@R3# set routing-instances VPNA routing-options multicast ssm-groups
239.255.255.3/32
```

Configuring the p2mp RSVP LSP as an inclusive provider tunnel.

```
[edit]
lab@R3# set routing-instances VPNA provider-tunnel rsvp-te label-switched-path-template
p2mp-template
```

Enable PIM on the interface connected to the CE site.

```
[edit]
lab@R3# set routing-instances VPNA protocols pim interface ge-0/0/5.900 mode sparse
```

Configure the mvpn protocol under the routing-instance. This will create the mvpn.0 rib and enable t mvpn functionality for the routing instance.

```
[edit]
lab@R3# set routing-instances VPNA protocols mvpn
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R3# commit
commit complete
```

c) R7

The second multicast source is attached to router R7 and configured similarly as R3.

```
[edit]
lab@R7# set protocols mpls label-switched-path p2p-template template

[edit]
lab@R7# set protocols mpls label-switched-path p2p-template no-cspf

[edit]
lab@R7# set protocols mpls label-switched-path p2p-template p2mp

[edit]
lab@R7# set routing-instances VPNA routing-options multicast ssm-groups
239.255.255.3/32

[edit]
lab@R7# set routing-instances VPNA provider-tunnel rsvp-te label-switched-path-template
p2p-template

[edit]
lab@R7# set routing-instances VPNA protocols pim interface ge-0/0/5.901 mode sparse

[edit]
lab@R7# set routing-instances VPNA protocols mvpn

[edit]
lab@R7# commit
commit complete
```

d) R1

The configuration of routers R1 and R5, where the receivers are located, is much simpler. Because they are dedicated as receiver-only sites, only the mvpn protocol must be enabled, and c-PIM is configured on the interface to the CE sites.

```
[edit]
lab@R1# set routing-instances VPNA routing-options multicast ssm-groups
239.255.255.3/32


[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R1# set routing-instances VPNA protocols mvpn

[edit]
lab@R1# set routing-instances VPNA protocols pim interface ge-0/0/5.912 mode sparse

[edit]
lab@R1# commit
commit complete
```

e) R5

```
[edit]
lab@R5# set routing-instances VPNA routing-options multicast ssm-groups
239.255.255.3/32

[edit]
lab@R5# set routing-instances VPNA protocols pim interface ge-0/0/5.921 mode sparse

[edit]
lab@R5# set routing-instances VPNA protocols mvpn

[edit]
lab@R5# commit
commit complete
```

8) Manually designate the receiving sites and ensure that they cannot be the root of a provider-tunnel. Configure the routers to choose the upstream multicast PE router based on the unicast route.

f) R1

Explicitly configuring routers R1 and R5 as receiver sites under the VPNA routing-instance mvpn configuration.

```
[edit]
lab@R1# set routing-instances VPNA protocols mvpn receiver-site
```

Because the source is multihomed to routers R3 and R7, the rest of the PEs must select single forwarded to avoid duplicating the multicast traffic. By default, the PE with the numerically highest IP address is selected. To amend the default behavior, and make sure that the standard BGP algorithm is used for selecting the designated source, the knob below should be configured under the **protocols mvpn** hierarchy for the proper routing-instance.

```
[edit]
lab@R1# set routing-instances VPNA protocols mvpn unicast-umh-election
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# commit
commit complete
```

g) R5

```
[edit]
lab@R5# set routing-instances VPNA protocols mvpn receiver-site

[edit]
lab@R5# set routing-instances VPNA protocols mvpn unicast-umh-election

[edit]
lab@R5# commit
commit complete
```

h) R3

```
[edit]
lab@R3# set routing-instances VPNA protocols mvpn unicast-umh-election

[edit]
lab@R3# commit
commit complete
```

i) R7

```
[edit]
lab@R7# set routing-instances VPNA protocols mvpn unicast-umh-election

[edit]
lab@R7# commit
commit complete
```

9) Ensure that the multicast PEs use a dedicated tunnel interface for the multicast traffic and assign different route-target to the mvpn routes compared to the unicast vpn routes.

j) R1

Configure a dedicated virtual-tunnels interface to serve only the multicast traffic.

```
[edit]
lab@R1# set chassis fpc 0 pic 0 tunnel-services
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R1# set interfaces vt-0/0/0 apply-groups-except enable_mpls_iso

[edit]
lab@R1# set interfaces vt-0/0/0 unit 0 family inet

[edit]
lab@R1# set routing-instances VPNA interface vt-0/0/0.0 multicast
```

By default, the mvpn routes advertised by the PEs have the same route-target community assigned as used for the unicast routes. To assign specific router-target community for the mvpn routes only, use the commands below.

```
[edit]
lab@R1# set routing-instances VPNA protocols mvpn route-target import-target target
target:1507417L:200

[edit]
lab@R1# set routing-instances VPNA protocols mvpn route-target export-target target
target:1507417L:200

[edit]
lab@R1# commit
commit complete
```

k) R5

```
[edit]
lab@R5# set chassis fpc 0 pic 0 tunnel-services

[edit]
lab@R5# set interfaces vt-0/0/0 apply-groups-except enable_mpls_iso

[edit]
lab@R5# set interfaces vt-0/0/0 unit 0 family inet

[edit]
lab@R1# set routing-instances VPNA interface vt-0/0/0.0 multicast



[edit]
lab@R5# set routing-instances VPNA protocols mvpn route-target import-target target
target:1507417L:200

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R5# set routing-instances VPNA protocols mvpn route-target export-target target
target:1507417L:200

[edit]
lab@R5# commit
commit complete
```

l) R3

```
[edit]
lab@R3# set chassis fpc 0 pic 0 tunnel-services

[edit]
lab@R3# set interfaces vt-0/0/0 apply-groups-except enable_mpls_iso

[edit]
lab@R3# set interfaces vt-0/0/0 unit 0 family inet

[edit]
lab@R3# set routing-instances VPNA interface vt-0/0/0.0 multicast

[edit]
lab@R3# set routing-instances VPNA protocols mvpn route-target import-target target
target:1507417L:200

[edit]
lab@R3# set routing-instances VPNA protocols mvpn route-target export-target target
target:1507417L:200

[edit]
lab@R3# commit
commit complete
```

m) R7

```
[edit]
lab@R7# set chassis fpc 0 pic 0 tunnel-services

[edit]
lab@R7# set interfaces vt-0/0/0 apply-groups-except enable_mpls_iso


[edit]
lab@R7# set interfaces vt-0/0/0 unit 0 family inet

[edit]
lab@R7# set routing-instances VPNA interface vt-0/0/0.0 multicast
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# set routing-instances VPNA protocols mvpn route-target import-target target
target:1507417L:200

[edit]
lab@R7# set routing-instances VPNA protocols mvpn route-target export-target target
target:1507417L:200
```

n) Verifying the multicast configuration.

Run the operational commands to check the NG-MVPN multicast operation. All routers enabled for mvpn advertise type 1 routes used for auto discovery.

```
[edit]
lab@R1# run show route table VPNA.mvpn.0

VPNA.mvpn.0: 6 destinations, 8 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1507417L:100:172.20.25.1/240
                   *[MVPN/70] 00:20:17, metric2 1
                      Indirect
1:1507417L:100:172.20.25.3/240
                   *[BGP/170] 00:20:17, localpref 100, from 172.20.25.5
                      AS path: (65001) I, validation-state: unverified
                    > to 172.20.20.66 via ge-0/0/4.15, label-switched-path r1-to-r3
                    [BGP/170] 00:20:17, localpref 100, from 172.20.25.6
                      AS path: (65001) I, validation-state: unverified
                    > to 172.20.20.66 via ge-0/0/4.15, label-switched-path r1-to-r3
1:1507417L:100:172.20.25.5/240
                   *[BGP/170] 00:17:02, localpref 100, from 172.20.25.5
                      AS path: I, validation-state: unverified
                    > to 172.20.20.66 via ge-0/0/4.15, label-switched-path r1-to-r5
1:1507417L:100:172.20.25.7/240
                   *[BGP/170] 00:20:17, localpref 100, from 172.20.25.5
                      AS path: (65001) I, validation-state: unverified
                    > to 172.20.20.70 via ge-0/0/4.16, label-switched-path r1-to-r7
                      to 172.20.20.66 via ge-0/0/4.15, label-switched-path r1-to-r7
                    [BGP/170] 00:20:17, localpref 100, from 172.20.25.6
                      AS path: (65001) I, validation-state: unverified
                    > to 172.20.20.70 via ge-0/0/4.16, label-switched-path r1-to-r7
                      to 172.20.20.66 via ge-0/0/4.15, label-switched-path r1-to-r7
7:1507417L:100:65001:32:10.10.0.2:32:239.255.255.3/240
                   *[PIM/105] 00:01:57, metric 0
                      Multicast (IPv4) Composite
7:1507417L:100:65001:32:10.10.0.6:32:239.255.255.3/240
                   *[PIM/105] 00:01:57, metric 0
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
                    Multicast (IPv4) Composite
```

Router R3 is the p2mp RSVP LSP root node, with all other PEs as leafs.

```
[edit]
lab@R3# run show mpls lsp p2mp
Ingress LSP: 1 sessions
P2MP name: 1507417L:100:mvpn:VPNA, P2MP branch count: 3
To              From            State Rt P     ActivePath       LSPname
172.20.25.5     172.20.25.3     Up     0 *
172.20.25.5:1507417L:100:mvpn:VPNA
172.20.25.1     172.20.25.3     Up     0 *
172.20.25.1:1507417L:100:mvpn:VPNA
172.20.25.7     172.20.25.3     Up     0 *
172.20.25.7:1507417L:100:mvpn:VPNA
Total 3 displayed, Up 3, Down 0

Egress LSP: 6 sessions, 1 detours
P2MP name: 1507417L:100:mvpn:VPNA, P2MP branch count: 1
To              From            State   Rt Style Labelin Labelout LSPname
172.20.25.3     172.20.25.7     Up       0  1 SE  301028        3
172.20.25.3:1507417L:100:mvpn:VPNA
Total 1 displayed, Up 1, Down 0

Transit LSP: 5 sessions
Total 0 displayed, Up 0, Down 0
```

On the receiver side:

```
[edit]
lab@R1# run show mpls lsp p2mp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 7 sessions, 1 detours
P2MP name: 1507417L:100:mvpn:VPNA, P2MP branch count: 1
To              From            State   Rt Style Labelin Labelout LSPname
172.20.25.1     172.20.25.3     Up       0  1 SE  299800        3
172.20.25.1:1507417L:100:mvpn:VPNA
P2MP name: 1507417L:100:mvpn:VPNA, P2MP branch count: 1
To              From            State   Rt Style Labelin Labelout LSPname
172.20.25.1     172.20.25.7     Up       0  1 SE  299822        3
172.20.25.1:1507417L:100:mvpn:VPNA
Total 2 displayed, Up 2, Down 0

Transit LSP: 1 sessions
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
Total 0 displayed, Up 0, Down 0
```

The SSM join message is received via BGP from the receivers to the sources.

```
[edit]
lab@R3# run show pim join instance VPNA extensive
Instance: PIM.VPNA Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.255.255.3
    Source: 10.10.0.2
    Flags: sparse
    Upstream interface: ge-0/0/5.900
    Upstream neighbor: Direct
    Upstream state: Local Source
    Keepalive timeout:
    Uptime: 00:31:57
    Downstream neighbors:
        Interface: Pseudo-MVPN
            Uptime: 00:18:32 Time since last Join: 00:18:32
    Number of downstream interfaces: 1
    Number of downstream neighbors: 1
```

On the receiver PEs, the upstream interface to the Source is learned through MVPN.

```
[edit]
lab@R1# run show pim join instance VPNA extensive
Instance: PIM.VPNA Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.255.255.3
    Source: 10.10.0.2
    Flags: sparse,spt
    Upstream protocol: BGP
    Upstream interface: Through BGP
    Upstream neighbor: Through MVPN
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 00:21:39
    Downstream neighbors:
        Interface: ge-0/0/5.912
            10.10.0.14 State: Join Flags: S Timeout: 177
            Uptime: 00:21:39 Time since last Join: 00:00:33
    Number of downstream interfaces: 1
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
    Number of downstream neighbors: 1

Group: 239.255.255.3
    Source: 10.10.0.6
    Flags: sparse,spt
    Upstream protocol: BGP
    Upstream interface: Through BGP
    Upstream neighbor: Through MVPN
    Upstream state: Join to Source
    Keepalive timeout:
    Uptime: 00:21:39
    Downstream neighbors:
        Interface: ge-0/0/5.912
            10.10.0.14 State: Join Flags: S Timeout: 177
            Uptime: 00:21:39 Time since last Join: 00:00:33
Number of downstream interfaces: 1
Number of downstream neighbors: 1
```

10) Configure VPLS connecting the VPLS-HUB site to the other two VPLS sites. Make sure you use BGP for VPN signaling and auto discovery.
11) The two VPLS spoke sites should be restricted to communicate only with HUB site. Because they are not allowed to communicate with each other, ensure that vpls connection between them is not available.

This task requires configuring VPLS connectivity between the three VPLS sites. Although VPLS provides point-to-multipoint communications, you must emulate hub-and-spoke connectivity. The site connected to router R2 should be designated as HUB and communicate with other two sites. However, the two other sites should not communicate with each other. To prevent this, a vpls connection should not be built between VPLS-S1 and VPLS-S2. The type of VPLS used should be BGP signaled with auto-discovery.

a) All devices including the route-reflector.

Enabling new MP-BGP family to carry VPLS signaling and auto-discovery.

Note: On R3, R4, R5, and R6 you will also need to apply these NLRI families to the sub-AS BGP groups.

```
[edit]
lab@R1# set protocols bgp group internal family l2vpn signaling
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) R2

Configure a new physical interface to connect site VPLS-HUB. You must exclude the apply-group automatically which applies family mpls to the logical interfaces. Otherwise, the commit will fail due to error.

```
[edit]
lab@R2# set interfaces ge-0/0/3 apply-groups-except enable_mpls_iso

[edit]
lab@R2# set interfaces ge-0/0/3 description "=== connection to VPLS-HUB ==="

[edit]
lab@R2# set interfaces ge-0/0/3 flexible-vlan-tagging

[edit]
lab@R2# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R2# set interfaces ge-0/0/3 unit 1500 encapsulation vlan-vpls

[edit]
lab@R2# set interfaces ge-0/0/3 unit 1500 vlan-id 1500
```

VPLS configuration is defined under routing instance of type vpls. Because BGP is used for auto-discovery, you must assign a route-distinguisher and route-target

```
[edit]
lab@R2# set routing-instances VPLS instance-type vpls
```

Configuring the normalization VLAN for the VPLS routing-instance.

```
[edit]
lab@R2# set routing-instances VPLS vlan-id 1500
```

Configure R2s interface connecting to site VPNA-HUB.

```
[edit]
lab@R2# set routing-instances VPLS interface ge-0/0/3.1500
```

Configure the route-distinguisher and router-target community for the VPLS routing-instance.

```
[edit]
lab@R2# set routing-instances VPLS route-distinguisher 1507417L:1500
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R2# set routing-instances VPLS vrf-target target:1507417L:1500
```

Enabling the vpls protocol and adding the customer interface and logical tunnel interface. The no-tunnel-service eliminates the need of hardware tunnel interface.

```
[edit]
lab@R2# set routing-instances VPLS protocols vpls no-tunnel-services
```

Define the identifier of the local site and configure the encapsulation to carry VLAN tag inside the VPLS tunnel.

```
[edit]
lab@R2# set routing-instances VPLS protocols vpls site VPLS-HUB site-identifier 1

[edit]
lab@R2# set routing-instances VPLS protocols vpls site VPLS-HUB encapsulation-type
ethernet-vlan

[edit]
lab@R2# set routing-instances VPLS protocols vpls site VPLS-HUB interface ge-0/0/3.1500

[edit]
lab@R2# commit
commit complete
```

c) R7

Configure a new physical interface to connect site VPLS-S1. You must exclude the apply-group, automatically applying family mpls to the logical interfaces.

```
[edit]
lab@R7# set interfaces ge-0/0/3 apply-groups-except enable_mpls_iso

[edit]
lab@R7# set interfaces ge-0/0/3 flexible-vlan-tagging


[edit]
lab@R7# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R7# set interfaces ge-0/0/3 unit 1050 description "=== connection to VPLS-S1 ==="

[edit]
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
lab@R7# set interfaces ge-0/0/3 unit 1050 encapsulation vlan-vpls

[edit]
lab@R7# set interfaces ge-0/0/3 unit 1050 vlan-id 1050

[edit]
lab@R7# set routing-instances VPLS instance-type vpls
```

Assign R7's interface connecting to site VPNA-S1.

```
[edit]
lab@R7# set routing-instances VPLS interface ge-0/0/3.1050
```

Configure a route-distinguisher and router-target community for the VPLS routing-instance.

```
[edit]
lab@R7# set routing-instances VPLS route-distinguisher 1507417L:1500

[edit]
lab@R7# set routing-instances VPLS vrf-target target:1507417L:1500
```

Lower the site-range to 2 from the default value of 65534, to enable this site to build a VPLS connection to sites with identifier of 1 or 2. In fact, site-range of 1 would accomplish the task requirements also, but here we use 2, in case another HUB is added later for high availability.

```
[edit]
lab@R7# set routing-instances VPLS protocols vpls site-range 2

[edit]
lab@R7# set routing-instances VPLS protocols vpls no-tunnel-services
```

The local site-identifier is configured to 7.

```
[edit]
lab@R7# set routing-instances VPLS protocols vpls site VPLS-S1 site-identifier 7


[edit]
lab@R7# set routing-instances VPLS protocols vpls site VPLS-S1 encapsulation-type
ethernet-vlan

[edit]
lab@R7# set routing-instances VPLS protocols vpls site VPLS-S1 interface ge-0/0/3.1050
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R7# commit
commit complete
```

d) R8

Configure a new physical interface to connect site VPLS-S2. You must exclude the apply-group automatically applying family mpls to the logical interfaces.

```
[edit]
lab@R8# set interfaces ge-0/0/3 apply-groups-except enable_mpls_iso

[edit]
lab@R8# set interfaces ge-0/0/3 description "=== connection to VPLS-S2 ==="

[edit]
lab@R8# set interfaces ge-0/0/3 flexible-vlan-tagging

[edit]
lab@R8# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services

[edit]
lab@R8# set interfaces ge-0/0/3 unit 1500 encapsulation vlan-vpls

[edit]
lab@R8# set interfaces ge-0/0/3 unit 1500 vlan-tags outer 600

[edit]
lab@R8# set interfaces ge-0/0/3 unit 1500 vlan-tags inner 1500

[edit]
lab@R8# set routing-instances VPLS instance-type vpls
```

Assign R8's interface connecting to site VPNA-S2.

```
[edit]
lab@R8# set routing-instances VPLS interface ge-0/0/3.1500

[edit]
lab@R8# set routing-instances VPLS route-distinguisher 1507417L:1500

[edit]
lab@R8# set routing-instances VPLS vrf-target target:1507417L:1500
```

Lower the site-range to 2 from the default value of 65534. This way this site could build VPLS connection to sites with identifier of 1 or 2.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R8# set routing-instances VPLS protocols vpls site-range 2

[edit]
lab@R8# set routing-instances VPLS protocols vpls no-tunnel-services
```

The local site-identifier is configured to 8. However, the HUB PE site-range is left by default. Which means it would be able to establish a VPLS session to VPLS-S2. However, this will prevent building a VPLS connection between VPLS-S1 and VPLS-S2, because they are out of range.

```
[edit]
lab@R8# set routing-instances VPLS protocols vpls site VPLS-S2 site-identifier 8

[edit]
lab@R8# set routing-instances VPLS protocols vpls site VPLS-S2 encapsulation-type
ethernet-vlan

[edit]
lab@R8# set routing-instances VPLS protocols vpls site VPLS-S2 interface ge-0/0/3.1500

[edit]
lab@R8# commit
commit complete
```

12) According to the diagram, the two VPLS spoke sites are configured with different VLANs than the HUB. You must normalize the VLAN used in the VPN to VLAN 1500. Do not make changes under the interface level to achieve the requirement.

Before you verify reachability between the HUB and spokes sites, you must perform VLAN normalization. The task restricts you from using VLAN map under the R7's interface. Therefore, routing-instance VLAN normalization is used to swap the VLAN numbers.

e) R7

```
[edit]
lab@R7# set routing-instances VPLS vlan-id 1500

[edit]
lab@R7# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

f) R8

Site VPLS-S2 is configured with stacked VLANs. Using a VLAN map applied to the interface removes the outer or service tag from the input frames and put it back on the output frames.

```
[edit]
lab@R8# set interfaces ge-0/0/3 unit 1500 input-vlan-map pop

[edit]
lab@R8# set interfaces ge-0/0/3 unit 1500 output-vlan-map push

[edit]
lab@R8# commit
commit complete
```

g) Verify the configuration.

Verify the interface status for applied VLAN normalization.

```
[edit]
lab@R7# run show interfaces ge-0/0/3.1050
  Logical interface ge-0/0/3.1050 (Index 347) (SNMP ifIndex 548)
    Description: === connection to VPLS-S1 ===
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1050 ] In(swap .1500) Out(swap .1050)
Encapsulation: VLAN-VPLS
    Input packets : 0
    Output packets: 0
    Protocol vpls, MTU: 1522
      Flags: Is-Primary
```

```
[edit]
lab@R8# run show interfaces ge-0/0/3.1500
  Logical interface ge-0/0/3.1500 (Index 345) (SNMP ifIndex 537)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.600 0x8100.1500 ] In(pop) Out(push
0x8100.600)  Encapsulation: VLAN-VPLS
    Input packets : 0
    Output packets: 0
    Protocol vpls, MTU: 1522
      Flags: Is-Primary
```


© 2020 Juniper Networks, Inc. All rights reserved.

13) Limit the mac-address table to not more than 20 hosts on the Spoke sites. Drop traffic if this limit is exceeded. Disable the mac learning function on the HUB site.

h) R7 and R8

Limit the MAC address table on interfaces connected to VPLS-S1 and VPLS-S2. If the limit is reached, the packets are dropped.

```
[edit]
lab@R7# set routing-instances VPLS protocols vpls site VPLS-S1 interface ge-0/0/3.1050
interface-mac-limit 20

[edit]
lab@R7# set routing-instances VPLS protocols vpls site VPLS-S1 interface ge-0/0/3.1050
interface-mac-limit packet-action drop

[edit]
lab@R7# commit
commit complete
```

i) R2

The mac learning functionality on the interface connected to VPLS-HUB site.

```
[edit]
lab@R2# set routing-instances VPLS protocols vpls site VPLS-HUB interface ge-0/0/3.1500
no-mac-learning

[edit]
lab@R2# commit
commit complete
```

j) Verifying the VPLS configuration.

The router R2 has connection to both Spoke PEs.

```
[edit]
lab@R2# run show vpls connections
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
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
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: VPLS
Edge protection: Not-Primary
  Local site: VPLS-HUB (1)
    connection-site           Type  St     Time last up          # Up trans
    7                         rmt   Up     Dec  4 22:22:50 2016           1
      Remote PE: 172.20.25.7, Negotiated control-word: No
      Incoming label: 262151, Outgoing label: 262145
      Local interface: lsi.1048576, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls VPLS local site 1 remote site 7
      Flow Label Transmit: No, Flow Label Receive: No
    8                         rmt   Up     Dec  4 22:27:05 2016           1
      Remote PE: 172.20.25.8, Negotiated control-word: No
      Incoming label: 262152, Outgoing label: 262169
      Local interface: lsi.1048577, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls VPLS local site 1 remote site 8
      Flow Label Transmit: No, Flow Label Receive: No
```

The Spoke PEs have built connection only with HUB PE, but not between each other.

```
[edit]
lab@R7# run show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
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: VPLS
Edge protection: Not-Primary
  Local site: VPLS-S1 (7)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Dec  4 17:59:05 2016           1
      Remote PE: 172.20.25.2, Negotiated control-word: No
      Incoming label: 262145, Outgoing label: 262151
      Local interface: lsi.1049088, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls VPLS local site 7 remote site 1
      Flow Label Transmit: No, Flow Label Receive: No
    8                         rmt   CN
```

```
[edit]
lab@R8# run show vpls connections
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
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
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: VPLS
Edge protection: Not-Primary
  Local site: VPLS-S2 (8)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Dec  4 18:33:28 2016           1
      Remote PE: 172.20.25.2, Negotiated control-word: No
      Incoming label: 262169, Outgoing label: 262152
      Local interface: lsi.1049088, Status: Up, Encapsulation: VPLS
        Description: Intf - vpls VPLS local site 8 remote site 1
      Flow Label Transmit: No, Flow Label Receive: No
    7                         rmt   CN
```

The ping tests are successful from the Spoke sites to the Hub. However, the ping between the Spokes is not possible.

```
[edit]
lab@vr-device-1> ping routing-instance VPLS-S2 192.168.5.1
PING 192.168.5.1 (192.168.5.1): 56 data bytes
64 bytes from 192.168.5.1: icmp_seq=0 ttl=63 time=21.056 ms
64 bytes from 192.168.5.1: icmp_seq=1 ttl=64 time=2.633 ms
64 bytes from 192.168.5.1: icmp_seq=2 ttl=64 time=2.561 ms
64 bytes from 192.168.5.1: icmp_seq=3 ttl=64 time=2.629 ms
^C
--- 192.168.5.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 2.561/7.220/21.056/7.988 ms
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@vr-device-1> ping routing-instance VPLS-S1 192.168.5.1
PING 192.168.5.1 (192.168.5.1): 56 data bytes
64 bytes from 192.168.5.1: icmp_seq=0 ttl=63 time=26.416 ms
64 bytes from 192.168.5.1: icmp_seq=1 ttl=64 time=2.574 ms
^C
--- 192.168.5.1 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 2.574/14.495/26.416/11.921
ms
```

14) Configure interconnectivity between the L2VPN site and L2circuit site, using router R4 as stitching point. L2VPN site is using BGP signaled VPN, whereas the L2circuit is using LDP based VPN.
15) Dedicate router R4 to provide connectivity between both L2 VPN sites. You are not allowed to use *lt-0/0/0* interface to accomplish the task.

The last task requires configuring a successful connectivity between the L2VPN and L2circuit sites. However, they are using different type of L2VPN. Therefore, router R4 should be used to stitch the two L2VPN and successfully establish connectivity.

a) R1

Configuring a new physical interface for connecting site L2VPN. You must exclude the apply-group automatically applying family mpls to the logical interfaces.

```
[edit]
lab@R1# set interfaces ge-0/0/3 apply-groups-except enable_mpls_iso

[edit]
lab@R1# set interfaces ge-0/0/3 description "=== connection to L2VPN ==="

[edit]
lab@R1# set interfaces ge-0/0/3 flexible-vlan-tagging

[edit]
lab@R1# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services
```

The interface unit is configured with tagging and circuit cross connect encapsulation.

```
[edit]
lab@R1# set interfaces ge-0/0/3 unit 1700 encapsulation vlan-ccc

[edit]
lab@R1# set interfaces ge-0/0/3 unit 1700 vlan-id 1700
```


© 2020 Juniper Networks, Inc. All rights reserved.

Create an MPLS LSP from R1 to R4 with CSPF disabled.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r4 to 172.20.25.4 no-cspf
```

Configuring the BGP signaled L2VPN for the routing-instance.

```
[edit]
lab@R1# set routing-instances L2VPN instance-type l2vpn
```

Adding the CE interfaces to the routing-instance.

```
[edit]
lab@R1# set routing-instances L2VPN interface ge-0/0/3.1700
```

Configure the route-distinguisher and router-target community for the L2VPN routing-instance. The L2VPN is using the same BGP signaling method used for the BGP VPLS configuration in the previous task.

```
[edit]
lab@R1# set routing-instances L2VPN route-distinguisher 1507417L:1700

[edit]
lab@R1# set routing-instances L2VPN vrf-target target:1507417L:1700
```

Configure the l2vpn protocol under the routing-instance. Local site identifier is configured with value of 1.

```
[edit]
lab@R1# set routing-instances L2VPN protocols l2vpn encapsulation-type ethernet-vlan

[edit]
lab@R1# set routing-instances L2VPN protocols l2vpn site L2VPN site-identifier 1

[edit]
lab@R1# set routing-instances L2VPN protocols l2vpn site L2VPN interface ge-0/0/3.1700

[edit]
lab@R1# commit
commit complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

b) R4

This task does not allow using lt- interface to stitch the two L2VPNs. However, Junos provides another way for doing that, using the internetworking interface and protocol. The interface is configured similarly like the lt- interface, if used for the same purpose.

Configuring the two units of the iw0 interface.

```
[edit]
lab@R4# set interfaces iw0 apply-groups-except enable_mpls_iso

[edit]
lab@R4# set interfaces iw0 unit 0 encapsulation vlan-ccc

[edit]
lab@R4# set interfaces iw0 unit 0 vlan-id 1700

[edit]
lab@R4# set interfaces iw0 unit 0 peer-unit 1

[edit]
lab@R4# set interfaces iw0 unit 1 encapsulation vlan-ccc

[edit]
lab@R4# set interfaces iw0 unit 1 vlan-id 1700

[edit]
lab@R4# set interfaces iw0 unit 1 peer-unit 0
```

In addition to configuring the internetworking interface, the l2iw protocol must be enabled. The configuration below is the only one needed to enable the internetworking successfully.

```
[edit]
lab@R4# set protocols l2iw
```

Configuring the LDP L2VPN on router R4 using the internetworking interface.

```
[edit]
lab@R4# set protocols l2circuit neighbor 172.20.25.5 interface iw0.1 virtual-circuit-id
1

[edit]
lab@R4# set protocols l2circuit neighbor 172.20.25.5 interface iw0.1 ignore-mtu-
mismatch
```


© 2020 Juniper Networks, Inc. All rights reserved.

Configure a return MPLS LSP from R4 to R1 with CSPF disabled.

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r1 to 172.20.25.1 no-cspf
```

Configure the BGP L2VPN for the other unit of the internetworking interface.

```
[edit]
lab@R4# set routing-instances L2VPN instance-type l2vpn

[edit]
lab@R4# set routing-instances L2VPN interface iw0.0

[edit]
lab@R4# set routing-instances L2VPN route-distinguisher 1507417L:1700

[edit]
lab@R4# set routing-instances L2VPN vrf-target target:1507417L:1700

[edit]
lab@R4# set routing-instances L2VPN protocols l2vpn encapsulation-type ethernet-vlan

[edit]
lab@R4# set routing-instances L2VPN protocols l2vpn site L2VPN site-identifier 2

[edit]
lab@R4# set routing-instances L2VPN protocols l2vpn site L2VPN interface iw0.0

[edit]
lab@R4# commit
commit complete
```

c) R5

Configuring the L2circuit on router R5 to establish L2VPN with router R4. Configuring the interface to connect L2circuit site to the MPLS network.

```
[edit]
lab@R5# set interfaces ge-0/0/3 apply-groups-except enable_mpls_iso

[edit]
lab@R5# set interfaces ge-0/0/3 description "=== connection to L2Circuit ==="

[edit]
lab@R5# set interfaces ge-0/0/3 flexible-vlan-tagging

[edit]
lab@R5# set interfaces ge-0/0/3 encapsulation flexible-ethernet-services
```


© 2020 Juniper Networks, Inc. All rights reserved.

The same tagged encapsulation on the interface for the L2circuit VPN as with the L2VPN.

```
[edit]
lab@R5# set interfaces ge-0/0/3 unit 1700 encapsulation vlan-ccc

[edit]
lab@R5# set interfaces ge-0/0/3 unit 1700 vlan-id 1700
```

The virtual-circuit-id must match on both ends of the L2circuit connection.

```
[edit]
lab@R5# set protocols l2circuit neighbor 172.20.25.4 interface ge-0/0/3.1700 virtual-
circuit-id 1


[edit]
lab@R5# set protocols l2circuit neighbor 172.20.25.4 interface ge-0/0/3.1700 ignore-
mtu-mismatch
```

16) Configure an additional LSP between R4 and R5 avoiding links with GOLD color. Make sure the L2circuit VPN is using that specific LSP for forwarding traffic.

Junos allows for steering L2VPN traffic over a specific LSP. Use a community assigned to the L2circuit and then match on it using routing policy, which is then applied to the forwarding table

a) R4

Configuring a dedicated LSP avoiding links with GOLD color.

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r5-l2circit to 172.20.25.5

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r5-l2circit admin-group exclude
GOLD
```

Assigning community to the l2circuit configuration.

```
[edit]
lab@R4# set protocols l2circuit neighbor 172.20.25.5 interface iw0.1 community
l2circuit-map
```

Configure routing policy that matches on the community and the l2circuit protocol. The action applied by the policy is manually setting the dedicated LSP as a next hop.

© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# set policy-options policy-statement mapping-l2circuit-lsp term 1 from protocol
l2circuit

[edit]
lab@R4# set policy-options policy-statement mapping-l2circuit-lsp term 1 from community
l2circuit-map

[edit]
lab@R4# set policy-options policy-statement mapping-l2circuit-lsp term 1 then install-
nexthop lsp r4-to-r5-l2circit
```

Defying the community under the `policy-options` hierarchy.

```
[edit]
lab@R4# set policy-options community l2circuit-map members 1700:1700
```

Applying the routing-policy to the forwarding table in the export direction.

```
[edit]
lab@R4# set routing-options forwarding-table export mapping-l2circuit-lsp

[edit]
lab@R4# commit
commit complete
```

b) R5

Similar configuration is applied to the router R5 steering the traffic in the opposite direction.

```
[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r4-l2circit to 172.20.25.4

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r4-l2circit admin-group exclude
GOLD

[edit]
lab@R5# set protocols l2circuit neighbor 172.20.25.4 interface ge-0/0/3.1700 community
l2circuit-map

[edit]
lab@R5# set policy-options policy-statement mapping-l2circuit-lsp term 1 from protocol
l2circuit

[edit]
lab@R5# set policy-options policy-statement mapping-l2circuit-lsp term 1 from community
l2circuit-map
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R5# set policy-options policy-statement mapping-l2circuit-lsp term 1 then install-
nexthop lsp r5-to-r4-l2circit

[edit]
lab@R5# set policy-options community l2circuit-map members 1700:1700

[edit]
lab@R5# set routing-options forwarding-table export mapping-l2circuit-lsp
```

c) Verify the configuration.

Checking the status of the two L2VPN connection on router R4.

```
[edit]
lab@R4# run 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  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- Hot-standby Connection
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 172.20.25.5
    Interface                 Type  St     Time last up          # Up trans
    iw0.1(vc 1)               rmt   Up     Dec  5 09:46:21 2016           1
      Remote PE: 172.20.25.5, Negotiated control-word: Yes (Null)
      Incoming label: 299776, Outgoing label: 299776
      Negotiated PW status TLV: No
      Local interface: iw0.1, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
[edit]
lab@R4# run show l2vpn connections
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: L2VPN
Edge protection: Not-Primary
  Local site: L2VPN (2)
    connection-site           Type  St     Time last up          # Up trans
    1                         rmt   Up     Dec  5 09:46:21 2016           1
      Remote PE: 172.20.25.1, Negotiated control-word: Yes (Null)
      Incoming label: 800000, Outgoing label: 800001
      Local interface: iw0.0, Status: Up, Encapsulation: VLAN
      Flow Label Transmit: No, Flow Label Receive: No
```

The mpls routing table reveals that the forwarding policy installs the correct LSP for next hop.

```
[edit]
lab@R4# run show route table mpls.0

mpls.0: 48 destinations, 50 routes (48 active, 2 holddown, 0 hidden)
Restart Complete
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
+ = Active Route, - = Last Active, * = Both
...
299776              *[L2IW/6] 01:12:57, metric2 11
                     > to 172.20.20.22 via ge-0/0/4.46, label-switched-path r4-to-r1
                     [L2CKT/7] 01:12:57
                     > via iw0.1, Pop       Offset: 4
800000              *[L2IW/6] 00:06:17, metric2 10
                        to 172.20.20.18 via ge-0/0/4.45, label-switched-path r4-to-r5
                     >  to 172.20.20.1 via ae0.0, label-switched-path r4-to-r5-l2circuit
                     [L2VPN/7] 01:12:57
                     > via iw0.0, Pop       Offset: 4
iw0.0               *[L2VPN/7] 01:12:57, metric2 11
                     > to 172.20.20.22 via ge-0/0/4.46, label-switched-path r4-to-r1
iw0.1               *[L2CKT/7] 00:06:17, metric2 10
                       to 172.20.20.1 via ae0.0, label-switched-path r4-to-r5-l2circuit
```

On router R5, the L2circuit is UP and the correct LSP is used for sending the traffic to router R4.

```
[edit]
lab@R5# run 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  SP -- Static Pseudowire
LD -- local site signaled down   RS -- remote site standby
RD -- remote site signaled down  HS -- Hot-standby Connection
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 172.20.25.4
    Interface                 Type  St     Time last up          # Up trans
    ge-0/0/3.1700(vc 1)       rmt   Up     Dec  5 09:28:07 2016           1
      Remote PE: 172.20.25.4, Negotiated control-word: Yes (Null)
      Incoming label: 299776, Outgoing label: 299776
      Negotiated PW status TLV: No
      Local interface: ge-0/0/3.1700, Status: Up, Encapsulation: VLAN
```


© 2020 Juniper Networks, Inc. All rights reserved.

```
      Flow Label Transmit: No, Flow Label Receive: No
```

```
[edit]
lab@R5# run show route table mpls.0

mpls.0: 32 destinations, 32 routes (32 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
...
299776             *[L2CKT/7] 01:09:20
                    > via ge-0/0/3.1700, Pop       Offset: 4
ge-0/0/3.1700      *[L2CKT/7] 01:09:20, metric2 10
                      to 172.20.20.13 via ge-0/0/4.35, label-switched-path r5-to-r4-
l2circuit
```

Use ping to test the connectivity between the sites.

```
[edit]
lab@vr-device-1> ping routing-instance L2circuit 192.168.200.1
PING 192.168.200.1 (192.168.200.1): 56 data bytes
64 bytes from 192.168.200.1: icmp_seq=0 ttl=63 time=28.920 ms
64 bytes from 192.168.200.1: icmp_seq=1 ttl=64 time=3.505 ms
64 bytes from 192.168.200.1: icmp_seq=2 ttl=64 time=3.119 ms
64 bytes from 192.168.200.1: icmp_seq=3 ttl=64 time=3.261 ms
^C
--- 192.168.200.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.119/9.701/28.920/11.097 ms
```

You have completed Super Lab 2.

© 2020 Juniper Networks, Inc. All rights reserved.