

iNET ZERO

your JNCIE training partner

JNCIE-SP (Service Provider) Lab preparation workbook v2.0

For Juniper Networks, inc - JNCIE-SP Lab Exam

<http://www.inetzero.com> - Copyright 2018 iNET ZERO. All rights reserved
for personal non commercial use only – do not distribute

<https://t.me/learningnets>

Copyright and licensing information

This workbook, iNET ZERO's JNCIE-SP Lab Preparation workbook is developed by iNET ZERO. All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means without the prior written permission of iNET ZERO a registered company in the Netherlands. This product cannot be used by or transferred to any other person. You are not allowed to rent, lease, loan or sell iNET ZERO training products including this workbook and its configurations. You are not allowed to modify, copy, upload, email or distribute this workbook in any way. This product may only be used and printed for your own personal use and may not be used in any commercial way. Juniper (c), Juniper Networks inc, JNCIE, JNCIP, JNCIS, JNCIA, Juniper Networks Certified Internet Expert, are registered trademarks of Juniper Networks, Inc.

About iNET ZERO's content developers and authors:Maxim Frolov

Maxim lives in Russia and speaks Russian and English. He started his networking career in 1999. Throughout the years Maxim has designed and implemented several large scale networks for enterprise and service provider customers. Over the years he has developed several high quality courseware materials for industry leading networking vendors. Maxim has the following certifications: JNCIE, JNCIP-ENT, JNCIS-SEC, Nortel NNCSS. For technology Max values efficiency and pragmatic design. When Max is not at work he likes to spend time with his family. Max enjoys being outside in the nature and loves to travel and exploring the world.

Ivan Ivanov

Ivan van lives in East Europe country of Bulgaria. He has more than 10 years experience with IP technologies, working at several Internet Service Providers, big enterprise companies and International system integrators. Throughout his career, Ivan gained extensive experience designing, implementing and supporting IP networks based mostly on Juniper Networks and Cisco Systems solutions and devices. Ivan worked on various international projects, designing, securing and implementing MPLS/IP backbone for multinational mobile operators. Ivan has the following certificates: JNCIE, JNCIP-SEC and various Cisco certificates.

Jörg Buesink



Jörg lives in the Netherlands and brings more than 15 years of experience in the IT and networking industry. He worked for several large ISPs / service providers in the role of technical consultant, designer and network architect. He has extensive experience in network implementation, design and architecture. Jörg is quadruple JNCIE certified (JNCIE-DC#007, JNCIE-ENT#21, JNCIE-SP#284 and JNCIE-SEC#30) as well as triple CCIE#15032 (Routing/ Switching, Service provider and Security), Cisco CCDE#20110002, Huawei HCIE#2188 Routing and Switching certified.

General information

Rack rental service

Did you know that this workbook can be used in combination with iNETZERO's JNCIE rack rental service? Take a look on our website for more information www.inetzero.com

Target audience

This workbook is developed for experienced network engineers who are preparing for the Juniper Networks 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. iNET ZERO's JNCIE-SP preparation workbook is developed in such a way that we expect you to have theoretical knowledge about the JNCIE-SP lab exam blueprint topics (JNCIP-SP certified or working towards this certification). For example, in this workbook we will not explain what rib-groups, LSP's or Multicast VPNs are. What we will do is test if you are able to configure all these technologies based on certain requirements and understand how they interact in a typical SP environment.

How to use this workbook

We recommend that you start your JNCIE lab preparation with the workbook chapters only. Always take a note on the time spent for each chapter/ task to see if you improved once you go over the chapters again. Ensure that at least you go the workbook chapters twice before you start with the full day lab challenges. You are ready to try the Full day lab challenges if you are able to configure the chapter's tasks without the need of the chapter's answers. Each Full day lab challenge must be completed within 8 hours.

Topology diagrams

In the chapters you will find several topology diagrams in small format. In the appendix of this workbook you will find larger versions of the topology diagrams for better readability. We recommend to print the topology diagrams.

iNET ZERO support

Always feel free to ask us questions regarding the workbook or JNCIE rack rental. You can reach us at info@inetzero.com. We love to hear from you regarding your preparation progress. Your feedback regarding our products is also very appreciated!

Exam strategy

Lab preparation

To earn the prestigious JNCIE-SP credential you must first pass the following Juniper Networks written exams: JNCIA-Junos, JNCIS-SP and JNCIP-SP. Once you are JNCIP-SP certified you are eligible to take the 8 hour JNCIE-SP lab. You must have extensive hands-on experience with Junos Service Provider technologies (v)MX platforms. When preparing for the JNCIE-SP lab exam ensure you fully understand the exam topics and objectives.

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 at least to go twice through all the workbook tasks and the final lab challenges. 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 chapters tasks and super labs without looking at the solutions.

While using this workbook ensure you take study notes and record the time spend on each task. Time is of essence for the JNCIE-SP lab exam. Try to complete the final lab exam within 8 hours!

As mentioned you need extensive hands-on experience with (v)MX devices. If you do not have these devices available, you might consider using iNET ZERO's rack rental service. When using iNET ZERO's rack rental service you do not need to worry about buying your own equipment, patching cables, correct port numbering, software versions, etc as the iNET ZERO rack rental service will match the physical topologies used in this workbook.

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 exams begins, you are allowed 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 have to 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 or not it will have impact on the rest of your exam. If it does not impact any other task on the exam, move to the next task and revisit the question again later on 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. As long as 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 to 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.

JNCIE Hall of Fame

Let us know once you passed the JNCIE-SP lab exam and would like to be added to the JNCIE Hall of Fame.

www.inetzero.com/hall-of-fame

www.jnciehof.com

Workbook and configuration file updates

Together with this workbook you should have received configuration files that you will need while working through the workbook chapters. The workbook configurations .zip file contains initial and final configurations for each chapter. For example if you would like to start with Chapter 4, you must load baseline configs (baseline folder) and the final configuration files of Chapter 3 as the chapters build upon each other. The Full day lab scenarios have separate initial and final configurations. Although we test our products and have an extensive review process, it still can happen that an error is found and needs to be corrected in either this workbook or the configuration files. Also iNET ZERO might make changes or additions to the workbook or the configuration files. Check our website on a regular basis and subscribe to our social media newsfeeds for more information.

Table of Contents

General information	5
Rack rental service	5
Target audience	5
How to use this workbook	5
Topology diagrams	5
iNET ZERO support	5
Exam strategy	6
JNCIE Hall of Fame	8
Workbook and configuration file updates	8
Chapter One: General System Features	14
Task 1. Initial System Settings	15
Task 2. SNMP Configuration	18
Task 3. Firewall Filters	19
Task 4. Interface Configuration	20
Task 5. Scripting	22
Chapter Two: IGP Configuration and Troubleshooting	23
Task 1. OSPF Troubleshooting	23
Task 2. ISIS Troubleshooting	25
Task 3. IGP Rollout	28
Chapter Three: BGP and Routing Policy	32
Task 1. IBGP and Confederation	32
Task 2. EBGp Configuration	33
Task 3. Routing Policies	35
Task 4. IBGP and Route Reflection	36
Chapter Four: MPLS Configuration	38
Task 1. LDP Configuration	38
Task 2. RSVP Configuration	39
Task 3. RSVP Protection	43
Task 4. IPv6 Tunneling with 6PE	44
Chapter Five: L3VPN Configuration	45
Task 1. L3VPN Configuration	45
Task 2. Multicast in L3VPN	48
Task 3. IPv6 Tunneling with 6VPE	49
Chapter Six: L2VPN and VPLS Configuration	50
Task 1. L2VPN Configuration	50
Task 2. VPLS Configuration	52
Chapter Seven: Inter-provider VPN Configuration	54
Task 1. Inter-provider VPN Option B	54
Task 2. Inter-provider VPN Option C	55
Chapter Eight: Class of Service	56
Task 1. Forwarding Classes, Queues and Schedulers	56
Task 2. Classification, Policing and Marking	58

Chapter Nine: A Full Day Lab Challenge	59
Task 1: Initial System Configuration	60
Task 2: Building the Network	61
Task 3: IGP Configuration	64
Task 4: BGP Configuration	66
Task 5: MPLS Configuration	68
Task 6: VPN Configuration	70
Task 7: Class of Service Configuration	72
A Full Day Lab Challenge II	74
Part 1: System Features.....	75
Task 1.1 Service Configuration	75
Task 1.2: Centralized authentication management	75
Task 1.3: Local user configuration	76
Task 1.4: Active configuration archival and logging	76
Task 1.5: Advanced Interface configuration and chassis features.....	77
Task 1.6: Advanced RE Protection	77
Part 2: Troubleshooting and Configuring IGP	78
Task 2.1: Troubleshooting	78
Task 2.2: Connectivity to OSPFv3 Area.....	78
Task 2.3: RIP redistribution	78
Task 2.4: Multilevel IS-IS configuration	78
Task 2.5: Advanced IS-IS configuration.....	78
Part 3: Troubleshooting and Configuring BGP.....	80
Task 3.1: iBGP design and configuration	80
Task 3.2: eBGP peers and configuration.....	80
Task 3.3: BGP policy configuration	81
Task 3.4: BGP general requirements	81
Part 4: MPLS configuration.....	82
Task 4.1: MPLS and RSVP configuration	82
Task 4.2: MPLS and LDP configuration	82
Part 5: IPv6 configuration.....	83
Task 5.1: Native IPv6 configuration	83
Task 5.2: IPv6 tunneling configuration	83
Part 6: MPLS VPNs configuration	84
Task 6.1: VPNA configuration	84
Task 6.2: VPNA multicast configuration	84
Task 6.3: VPLS configuration	85
Task 6.4: L2VPN configuration.....	85
A Full Day Lab Challenge III	86
Part 1: System Features.....	87
Task 1.1 Service Configuration	87
Task 1.2: Centralized authentication management	88
Task 1.3: Local user configuration	88
Task 1.4: Active configuration archival and logging	88
Task 1.5: Advanced Interface configuration and chassis functions.....	88

Task 1.6: Advanced RE Protection	89
Part 2: Troubleshooting and Configuring IGP	90
Task 2.1: Troubleshooting	90
Task 2.2: RIP redistribution	90
Task 2.3: Multi-area OSPF configuration	90
Task 2.4: Advanced OSPF configuration	91
Part 3: Troubleshooting and Configuring BGP	92
Task 3.1: Internal BGP design and configuration.....	92
Task 3.2: external BGP peers and configuration	92
Task 3.3: BGP policy configuration	92
Task 3.4: BGP general requirements	93
Part 4: MPLS configuration.....	94
Task 4.1: MPLS and RSVP configuration	94
Task 4.2: MPLS and LDP configuration	94
Part 5: IPv6 tunneling and Multicast configuration.....	95
Task 5.1: IPv6 tunneling configuration	95
Task 5.2: Multicast configuration	95
Part 6: MPLS VPNs configuration	96
Task 6.1: VPNA configuration	96
Task 6.2: VPNB configuration	96
Task 6.3: VPLS configuration	96
Appendix 1: Additional Theory	98
OSPF adjacency troubleshooting.....	98
BGP adjacency troubleshooting	102
BGP IPV6 NLRI over IPV4 peering.....	105
Troubleshooting: Multicast traffic engineering using RIB-groups.....	112
Advanced firewall filtering	115
Appendix 2 : Topology diagrams	118
Task 1: Initial System Configuration	138
Task 2: SNMP Configuration.....	151
Task 3: Firewall filters.....	157
Task 4: Interface Configuration.	168
Task 5: Scripting.	178
Chapter Two solutions: IGP Configuration and Troubleshooting.....	184
Task 1: OSPF Troubleshooting	185
Task 2: IS-IS Troubleshooting	251
Task 3: IGP Rollout.....	305
Chapter Three solutions: BGP and Routing policy.....	379
Task 1: IBGP and Confederation.....	379
Task Two: EBGP Configuration.	390
Task 3: Routing Policies	430
Task 4: IBGP and Route Reflection	487
Chapter Four solutions: MPLS configuration.....	535
MPLS Overview	535

LDP Overview	536
Task 1: LDP Configuration	536
Task 2: RSVP Configuration	548
Task 3: RSVP Protection	609
Task 4: IPv6 tunneling with 6PE.....	632
Chapter Five solutions: L3VPN Configuration	642
Task 1: L3VPN configuration	642
Task 2: Multicast in L3VPNs.....	699
Task 3: IPv6 Tunneling with 6VPE.....	734
Chapter Six solutions: L2VPN and VPLS configuration	747
Task 1: L2VPN Configuration	747
Task 2: VPLS Configuration	764
Chapter Seven solutions: Inter-provider VPN Configuration	795
Task 1: Inter-provider VPN Option B	795
Task 2: Inter-provider VPN Option C	808
Chapter Eight solutions: Class of Service.....	822
Task 1: Forwarding Classes, Queues and Schedulers	823
Task 2: Classification, Policing and Marking	832
Chapter Nine solutions: Full Day Lab Challenge I	841
Task 1: Initial System Configuration	841
Task 2: Building the network.....	846
Task 3: IGP Configuration	848
Task 4: BGP Configuration	909
Task 5: MPLS configuration	966
Task 6: VPN configuration	1057
Task 7: Class of Service Configuration	1107
Chapter Ten solutions: Full Day Lab Challenge II	1123
Part 1: System Features.....	1124
Solution - Task 1.1 Service Configuration	1124
Solution - Task 1.2: Centralized authentication management	1125
Solution - Task 1.3: Local user configuration.....	1127
Solution - Task 1.4: Active configuration archival and logging	1129
Solution - Task 1.5: Advanced Interface configuration and chassis features	1130
Solution - Task 1.6: Advanced RE Protection.....	1133
Appendix - Part 2: Troubleshooting and Configuring IGP.....	1140
Solution - Task 2.1: Troubleshooting	1140
Solution - Task 2.2: Connectivity to OSPFv3 Area.....	1153
Solution - Task 2.3: RIP redistribution	1160
Solution - Task 2.4: Multilevel IS-IS configuration	1169
Solution - Task 2.5: Advanced IS-IS configuration	1173
Part 3: Troubleshooting and Configuring BGP.....	1176
Solution - Task 3.1: iBGP design and configuration	1176
Solution - Task 3.2: eBGP peers and configuration	1189
Solution - Task 3.3: BGP policy configuration	1197
Solution - Task 3.4: BGP general requirements.....	1217

Part 4: MPLS configuration.....	1223
Solution - Task 4.1: MPLS and RSVP configuration.....	1223
Part 5: IPv6 configuration.....	1248
Solution - Task 5.1: Native IPv6 configuration.....	1248
Solution - Task 5.2: IPv6 tunneling configuration.....	1251
Part 6: MPLS VPNs configuration	1255
Solution - Task 6.1: VPNA configuration.....	1255
Solution - Task 6.2: VPNA multicast configuration	1271
Solution - Task 6.3: VPLS configuration	1278
Solution - Task 6.4: L2VPN configuration	1285
Chapter eleven solutions: Full Day Lab Challenge III.....	1292
Part 1: System Features.....	1294
Solution - Task 1.1 Service Configuration.....	1294
Solution - Task 1.2: Centralized authentication management	1295
Solution - Task 1.3: Local user configuration.....	1296
Solution - Task 1.4: Active configuration archival and logging.....	1298
Solution - Task 1.5: Advanced Interface configuration and chassis functions.....	1299
Solution - Task 1.6: Advanced RE Protection.....	1302
Part 2: Troubleshooting and Configuring IGP	1311
Solution - Task 2.1: Troubleshooting	1311
Solution - Task 2.2: RIP redistribution	1330
Solution - Task 2.3: Multi-area OSPF configuration.....	1341
Solution - Task 2.4: Advanced OSPF configuration.....	1354
Part 3: Troubleshooting and Configuring BGP.....	1358
Solution - Task 3.1: Internal BGP design and configuration	1358
Solution - Task 3.2: External BGP peers and configuration	1365
Solution - Task 3.3: BGP policy configuration	1374
Solution - Task 3.4: BGP general requirements.....	1393
Part 4: MPLS configuration.....	1396
Solution - Task 4.1: MPLS and RSVP configuration.....	1396
Solution - Task 4.2: MPLS and LDP configuration	1413
Part 5: IPv6 tunneling and Multicast configuration.....	1417
Solution - Task 5.1: IPv6 tunneling configuration.....	1417
Solution - Task 5.2: Multicast configuration.....	1423
Part 6: MPLS VPNs configuration	1430
Solution - Task 6.1: VPNA configuration.....	1430
Solution - Task 6.2: VPNB configuration.....	1442
Solution - Task 6.3: VPLS configuration	1454

Chapter One: General System Features

TIP: Throughout the workbook before you begin a chapter, we recommend you to 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.




Figure 1




Figure 2

Task 1. Initial System Settings

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

NOTE: The lab uses a dedicated VR-device to emulate external systems interacting with your domain. The device is reachable at 10.10.1.9 IP address using user name "lab" and password "lab123".

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.

Load the initial/ baseline configurations on the devices and use root password *root123* on every router. Please **do not** change the root password on our devices to prevent unnecessary password recovery.

- 1) Configure the host names according to Table 1.

Table 1

Router	Router Type	Host Name
R1	vMX	Sun
R2	vMX	Sirius
R3	vMX	Canopus
R4	vMX	Arcturus
R5	vMX	A-Centauri
R6	vMX	Vega
R7	vMX	Rigel
R8	vMX	Procyon

- 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.

Table 2

Router	OoB Interface Name	OoB Interface IP Address
R1	ge-0/0/0	10.10.1.1/24
R2	ge-0/0/0	10.10.1.2/24
R3	ge-0/0/0	10.10.1.3/24
R4	ge-0/0/0	10.10.1.4/24
R5	ge-0/0/0	10.10.1.5/24
R6	ge-0/0/0	10.10.1.6/24
R7	ge-0/0/0	10.10.1.7/24
R8	ge-0/0/0	10.10.1.8/24

- 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 your 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 user name *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.

TIP: From this point on we recommend you to 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 user name.

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

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

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 users interactive commands
User “ops”	All warning level messages
All users	All critical level messages

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

Task 2. SNMP Configuration

In this task you will configure SNMP v3 for secure NMS interaction.

- 1) Configure SNMP v3 view parameters according to Table 5. Make sure that SNMP v3 provides read only access.

Table 5

Parameter	Value
USM user name	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

- 2) Configure SNMP v3 notification parameters according to Table 6.

Table 6

Parameter	Value
Target address	S1 server IP address
Target processing model	v3
Target security model	usm
Target security level	privacy
Target security name	lab
Notification OID filter	snmpTraps, jnxTraps
Notification type	trap

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.
- 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 i1 and i4 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.




Figure 3

NOTE: The interface unit numbers match the VLAN tags.

Table 7

Router	Interface	Interface Name	IP Address	IPv6 Address
R1	i1	ge-0/0/1	802.3ad	
	i2	ge-0/0/4.114	172.30.0.5/30	
	i3	ge-0/0/4.118	172.30.0.9/30	link-local
	i4	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	i1	ge-0/0/1	802.3ad	
	i2	ge-0/0/4.127	172.30.0.17/30	
	i3	ge-0/0/4.123	172.30.0.13/30	link-local
	i4	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	i1	ge-0/0/4.134	172.30.0.21/30	link-local
	i2	ge-0/0/4.136	172.30.0.25/30	
	i3	ge-0/0/4.123	172.30.0.14/30	link-local
	i4	ge-0/0/4.200	172.30.1.1/24	
	i5	ge-0/0/4.201	172.30.2.1/24	
		lo0.0	172.30.5.3/32	fd17:f0f4:f691:5::3/128
R4	i1	ge-0/0/4.134	172.30.0.22/30	link-local
	i2	ge-0/0/4.114	172.30.0.6/30	
	i3	ge-0/0/4.145	172.30.0.29/30	link-local
	i4	ge-0/0/4.200	172.30.1.2/24	
	i5	ge-0/0/4.201	172.30.2.2/24	
		lo0.0	172.30.5.4/32	fd17:f0f4:f691:5::4/128
R5	i1	ge-0/0/1	802.3ad	
	i2	ge-0/0/4.158	172.30.0.37/30	
	i3	ge-0/0/4.145	172.30.0.30/30	link-local
	i4	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	i1	ge-0/0/1	802.3ad	
	i2	ge-0/0/4.136	172.30.0.26/30	
	i3	ge-0/0/4.167	172.30.0.41/30	link-local
	i4	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	i1	ge-0/0/4.178	172.30.0.45/30	link-local
	i2	ge-0/0/4.127	172.30.0.18/30	
	i3	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	i1	ge-0/0/4.178	172.30.0.46/30	link-local
	i2	ge-0/0/4.158	172.30.0.38/30	
	i3	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

- 2) On R3 and R4 configure VRRP such as R3 is the VRRP master on i4 interface and R4 is the VRRP master on i5 interface. Use .254 Virtual Router IP address on the i4 and i5 subnets.
- 3) Make sure that R3 and R4 track their uplink interfaces i2 and i3 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.

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 also 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.

Chapter Two: IGP Configuration and Troubleshooting

This chapter contains three independent tasks: OSPF troubleshooting, ISIS troubleshooting and the new IGP rollout. NOTE: You need the final configurations you have saved in the previous chapter.

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 4

Table 8

Router	Interface	Area
R1	ae0.0	3
	i3	2
	lo0.0	2
R2	ae0.0	3
	i2	0
	i3	0
	lo0.0	0
R3	i1	4
	i2	0
	i3	0
	lo0.0	0
	i5	0

R4	i1	4
	i3	4
	lo0.0	4
R5	ae0.0	4
	i3	4
	lo0.0	4
R6	ae0.0	4
	i2	0
	i3	0
	lo0.0	0
R7	i1	1
	i2	0
	i3	0
	lo0.0	0
R8	i1	1
	i3	2
	lo0.0	2

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.
 - 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. ISIS Troubleshooting

In this task you load a broken ISIS 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.




Figure 5
Table 9

Router	Interface	Level
R1	ae0.0	2
	i3	1
	lo0.0	1
R2	ae0.0	2
	i2	2
	i3	2
	lo0.0	2
R3	i1	1
	i2	1
	i3	2
	lo0.0	1
R4	i1	1
	i3	1
	lo0.0	1
R5	ae0.0	1

	i3	1
	lo0.0	1
R6	ae0.0	1
	i2	1
	i3	2
	lo0.0	1
R7	i1	2
	i2	2
	i3	2
	lo0.0	2
R8	i1	2
	i3	1
	lo0.0	1

Table 10

Router	Area
R1	49.0001
R2	49.0002
R3	49.0002
R4	49.0002
R5	49.0002
R6	49.0002
R7	49.0002
R8	49.0001


The ISIS network must meet the following criteria:

- All ISIS adjacencies are up.
- All your routers can reach all other routers loopbacks.
- No routing loops are allowed anywhere.
- Each ISIS interface must have no more than one adjacency.
- All routers must use MD5 authentication for Hello ISIS PDU only on all ISIS 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 ISIS 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 ISIS 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.

- 4) Using operational and configuration mode commands troubleshoot the ISIS network and fix the errors.
- 5) Write a summary report on all the issues found.

Task 3. IGP Rollout

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







Figure 6

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.

Table 11

Router	Interface	Interface Name	IP Address	IPv6 Address
R4	i6	ge-0/0/4.202	172.30.0.49/30	
	i7	ge-0/0/4.203	172.30.0.53/30	link-local
R5	i5	ge-0/0/4.204	172.30.0.57/30	
	i6	ge-0/0/4.205	172.30.0.61/30	link-local



3) Figure 6. Table 12 lists the routers NET addresses.

Table 12

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

- 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 ISIS domain but the ISIS 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 ISIS domain.
- 6) Make sure that no pseudo nodes enter into the ISIS database.
- 7) Configure MD5 authentication on all ISIS enabled interfaces for all ISIS 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 ISIS neighbors can detect the adjacency loss in less than 500ms.
- 10) Make sure that all adjacencies are up and all routers can reach all other routers' IPv4 loopback addresses.
- 11) Configure RIP on R4 i6 and R5 i5 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 ISIS. Make sure that any of the R4 or R5 failure will not result in the RIP routes disappearing from the ISIS 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.
- 16) Configure OSPFv3 area 0 on R4 i7 and R5 i6 interfaces respectively. Make sure that OSPFv3 supports both IPv4 and IPv6 routing.
- 17) Advertise IPv4 and IPv6 ISIS routes to OSPFv3. Advertise IPv4 and IPv6 OSPFv3 routes to ISIS. Make sure that any of the R4 or R5 failure will not disrupt the routing between the ISIS 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.

Chapter Three: 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.




Figure 7

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

Table 13

Router	Interface	Interface Name	IP Address	IPv6 Address
R1	i5	ge-0/0/5.300	192.168.1.1/24	
R2	i5	ge-0/0/5.300	192.168.1.2/24	
R3	i6	ge-0/0/5.301	192.168.0.1/30	link-local
	i7	ge-0/0/5.302	192.168.0.5/30	
R5	i7	ge-0/0/5.303	192.168.0.9/30	IPv4 compatible/126
	i8	ge-0/0/5.304	192.168.0.13/30	IPv4 compatible/126
R6	i5	ge-0/0/5.305	192.168.0.17/30	
	i6	ge-0/0/5.306	192.168.0.21/30	
	i7	ge-0/0/5.307	192.168.0.25/30	
R7	i4	ge-0/0/5.308	192.168.0.29/30	fc09:c0:ffee::1/126
	i5	ge-0/0/5.309	192.168.0.33/30	

<http://www.inetzero.com> - Copyright 2018 iNET ZERO. All rights reserved
For personal non commercial use only – do not distribute - Steganoid=ON

<https://t.me/learningnets>

R8	i4	ge-0/0/5.310	192.168.0.37/30	fc09:c0:ffee::5/126


- 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 i6 or i7 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.
- 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.

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 it's OoB management port at 10.10.1.19 address. Feel free to modify the RR settings as needed.

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.

Table 14

Router	Interface	Interface Name	IP Address
R1	i6	ge-0/0/4.206	172.30.0.65/30
R2	i6	ge-0/0/4.207	172.30.0.69/30

- 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.
- 9) Ensure that all the IBGP session state changes are logged to syslog.
- 10) No unresolved IPv4 routes are allowed anywhere.

Chapter Four: 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 8

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 ISIS 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 333Mbps of bandwidth reservation.
- 3) Configure link administrative groups as shown in Table 15.




Figure 9
Table 15

Router	Interface	Admin. Group
R1	i2	green
	i3	red
	ae0.0	green, red
R2	i2	green
	i3	red
	ae0.0	green, red
R3	i1	green, red
	i2	green

	i3	red
R4	i1	green, red
	i2	green
	i3	red
R5	i2	green
	i3	red
	ae0.0	green, red
R6	i2	green
	i3	red
	ae0.0	green, red
R7	i1	green, red
	i2	green
	i3	red
R8	i1	green, red
	i2	green
	i3	red

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




Figure 10

Table 16

Ingress	Egress	LSP ID
Sun	Procyon	A
Sun	Vega	C
Sirius	Rigel	E
Sirius	A-Centauri	G
Canopus	Procyon	J
Canopus	Procyon	L
Canopus	Vega	Q
Arcturus	Rigel	N
Arcturus	Rigel	P
Arcturus	A-Centauri	S
A-Centauri	Sirius	H
A-Centauri	Arcturus	T
Vega	Sun	D
Vega	Canopus	R
Rigel	Sirius	F
Rigel	Arcturus	M
Rigel	Arcturus	O
Procyon	Sun	B
Procyon	Canopus	I
Procyon	Canopus	K

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 have to 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.
- 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 back 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 focus is 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.

Chapter Five: L3VPN Configuration

In this chapter tasks you implement L3VPN's. 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. L3VPN 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.

Table 17

Router	Interface	Interface Name	IP Address	IPv6 Address
R1	i7	ge-0/0/5.311	192.168.0.41/30	
	i8	ge-0/0/5.312	192.168.0.45/30	
	i9	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	i7	ge-0/0/5.314	192.168.0.53/30	
	i8	ge-0/0/5.315	192.168.0.57/30	
	i9	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	i8	ge-0/0/5.317		fc09:c0:ffee::9/126
	i9	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	i8	ge-0/0/5.319	192.168.0.73/30	
	i9	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	i9	ge-0/0/5.321	192.168.0.81/30	
		lo0.1	172.30.5.25/32	
R6	i8	ge-0/0/5.322	192.168.0.85/30	
		lo0.1	172.30.5.29/32	
R7	i6	ge-0/0/5.323	192.168.0.89/30	
		lo0.1	172.30.5.33/32	
R8	i5	ge-0/0/5.324	192.168.0.93/30	
	i6	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

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




Figure 11
Table 18

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

- 3) You may not have any MPLS LSPs on Route Reflector. A static route is allowed on the RR if needed.
- 4) Make sure that the customer C1 OSPF area 0 appears as a contiguous area without ABRs.

- 5) Customer C1 has some backdoor OSPF connections but prefers that your MPLS network would be used for traffic forwarding between the customer sites.
- 6) Make sure that your MPLS network can be used as a backup path between CE1-2 and CE1-3.
- 7) 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.
- 8) 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.
- 9) Make sure that if a route is originated in customer C2 site S1 or S2, it is never advertised back to the same site.
- 10) Make sure that PE-CE link subnets in customer C2 VPN are advertised to the customer remote VPN sites.
- 11) Make sure that all PE routers receive only the routes with those targets that they specifically request for.
- 12) 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.
- 13) 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.

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

- 14) Customer C2 must be provided with Internet access at the customer site S1, using a dedicated interface i9 at both R1 and R2 routers. All other customer sites should be able to reach the Internet via the site S1.

Task 2. Multicast in L3VPN

In this task you implement Draft-Rosen and Next Generation multicast in the L3VPNs.

NOTE: Both customers C1 and C2 use 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 RP's. 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 Draft-Rosen multicast in customer C1 VPN. Customer C1 uses auto-RP with CE1-2 and CE1-3 acting as both RP candidates and mapping agents.
- 5) Configure multicast data MDT in the customer C1 site S2 for multicast groups 239.0.0.1 and 239.0.0.2 from any source. The cutoff rate to switch over to the data MDT should be set to 30Mbps. Make sure that no more than 5 data MDTs are allowed.
- 6) 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.
- 7) Customer C2 outsources its RP to your network. Make sure that your routers R1 and R2 act as the customer anycast RPs.
- 8) 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.
- 9) 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.
- 10) Make sure that the customer C2 receiver sites join only source based multicast distribution trees.

Table 19

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

Table 20

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

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.

Chapter Six: L2VPN and VPLS Configuration

In this chapter tasks you implement L2VPN and VPLS applications in your network. The tasks include LDP and BGP signaled L2VPN and VPLS, dual-homed customers and loop prevention, L2VPN and VPLS interworking, LDP-signaled and BGP-signaled VPLS interworking and VPLS L3 interface configuration.

Task 1. L2VPN Configuration

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




Figure 12

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

Table 21

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

<http://www.inetzero.com> - Copyright 2018 iNET ZERO. All rights reserved
For personal non commercial use only – do not distribute - Steganoid=ON

<https://t.me/learningnets>

	S4	CE5-4	BGP	ge-0/0/3

Table 22

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

- 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.




Figure 13

- 1) Configure VPLS as shown in Figure 13. Table 23 specifies the VPLS details. Configure customer VLANs as shown in Table 24.

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

Table 24

Customer	VLAN
C5	600

	601
C6	700
	701

- 2) No L2 switching loops are allowed anywhere in the customers' VPLS networks. You may not use Spanning Tree protocol for loop prevention.
- 3) 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.
- 4) 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.
- 5) 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.

Chapter Seven: Inter-provider VPN Configuration

In this chapter you will 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.




Figure 14

- 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 has to 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.




Figure 15

- 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.

Chapter Eight: 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 25.

Table 25

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

- 2) Configure Schedulers with parameters shown in Table 26 and map them to the Forwarding Classes as indicated in Table 25.

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

- 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 27.

Table 27

Fill Level	Drop Probability
25	5
50	15
75	40

- 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 28.

Table 28

Fill Level	Drop Probability
25	10
50	30
75	65

- 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 29.

Table 29

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

- 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 30. 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.

Table 30

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	

- 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 30.

Chapter Nine: A Full Day Lab Challenge

In this chapter you will be presented with a complete 8 hour lab emulation scenario covering the tasks on multiple different ISP applications all together. **Error! Reference source not found.** shows the network topology used for this chapter.




Figure 16

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

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 31.

Table 31

Router	Router Type	Host Name
R1	MX	R1
R2	MX	R2
R3	MX	R3
R4	MX	R4
R5	MX	R5
R6	MX	R6
R7	MX	R7
R8	MX	R8

- 2) Configure OoB management interfaces on each device with the appropriate IP addresses. The devices and their respective IP addresses are listed in Table 32. Set the interface description.

Table 32

Device	OoB Interface Name	OoB Interface IP Address
R1	ge-0/0/0	10.10.1.1/24
R2	ge-0/0/0	10.10.1.2/24
R3	ge-0/0/0	10.10.1.3/24
R4	ge-0/0/0	10.10.1.4/24
R5	ge-0/0/0	10.10.1.5/24
R6	ge-0/0/0	10.10.1.6/24
R7	ge-0/0/0	10.10.1.7/24
R8	ge-0/0/0	10.10.1.8/24

- 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 redistributed to any dynamic routing protocol. Ensure 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 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 33.

Table 33

Username	Password	Privileges
noc	noc123	Class "operator" permissions. Additionally 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 cannot execute the "clear", "configure" or "edit" commands

- 11) Configure Syslog settings on all your devices as indicated in Table 4.

Table 34

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
User "lab"	All emergency level messages

- 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 user name *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.

- 18) Enable the scripts.

Task 2: Building the Network

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

<http://www.inetzero.com> - Copyright 2018 iNET ZERO. All rights reserved
For personal non commercial use only – do not distribute - Steganoid=ON

<https://t.me/learningnets>

- 1) Build the network by configuring interfaces as indicated in Table 7. Aggregated Ethernet interfaces are listed in Table 35.
- 2) Enable LACP continuity checking on the Aggregated Ethernet interfaces.
- 3) Set all the interfaces descriptions.

Table 35

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

NOTE: The interface unit numbers match the VLAN tags.

Table 36

Router	Interface	Interface Name	IP Address	IPv6 Address
R1	i1	ae0.0	172.30.0.1/30	link-local
	i2	ge-0/0/4.117	172.30.0.5/30	link-local
	i3	ge-0/0/4.118	172.30.0.9/30	link-local
	i4	ge-0/0/4.206	172.30.0.65/30	
	i5	ge-0/0/5.318	192.168.0.69/30	
	i6	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	i1	ae0.0	172.30.0.2/30	link-local
	i2	ge-0/0/4.126	172.30.0.17/30	link-local
	i3	ge-0/0/4.123	172.30.0.13/30	link-local
	i4	ge-0/0/4.207	172.30.0.69/30	
	i5	ge-0/0/5.303	192.168.0.9/30	IPv4 compatible/126
	i6	ge-0/0/3.601		
		lo0.0	172.30.5.2/32	fd17:f0f4:f691:5::2/128
R3	i1	ge-0/0/4.123	172.30.0.14/30	link-local
	i2	ge-0/0/4.138	172.30.0.33/30	link-local
	i3	ge-0/0/4.137	172.30.0.29/30	link-local
	i4	ge-0/0/4.135	172.30.0.85/30	link-local
	i5	ae0.0	172.30.0.81/30	link-local
	i6	ge-0/0/5.306	192.168.0.21/30	
	i7	ge-0/0/5.307	192.168.0.25/30	
	i8	ge-0/0/3.600		

		lo0.0	172.30.5.3/32	fd17:f0f4:f691:5::3/128
R4	i1	ge-0/0/4.146	172.30.0.89/30	link-local
	i2	ae0.0	172.30.0.82/30	link-local
	i3	ge-0/0/3.600		
	i4	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	i1	ge-0/0/4.135	172.30.0.86/30	link-local
	i2	ae0.0	172.30.0.93/30	link-local
	i3	ge-0/0/5.305	192.168.0.17/30	
	i4	ge-0/0/4.202	172.30.0.49/30	
		lo0.0	172.30.5.5/32	fd17:f0f4:f691:5::5/128
R6	i1	ge-0/0/4.126	172.30.0.18/30	link-local
	i2	ge-0/0/4.146	172.30.0.90/30	link-local
	i3	ae0.0	172.30.0.94/30	link-local
	i4	ge-0/0/4.167	172.30.0.45/30	link-local
	i5	ge-0/0/4.168	172.30.0.21/30	link-local
	i6	ge-0/0/4.204	172.30.0.57/30	
		lo0.0	172.30.5.6/32	fd17:f0f4:f691:5::6/128
R7	i1	ge-0/0/4.117	172.30.0.6/30	link-local
	i2	ge-0/0/4.137	172.30.0.30/30	link-local
	i3	ge-0/0/4.167	172.30.0.46/30	link-local
	i4	ge-0/0/5.311	192.168.0.41/30	
	i5	ge-0/0/5.312	192.168.0.45/30	
	i6	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	i1	ge-0/0/4.118	172.30.0.10/30	link-local
	i2	ge-0/0/4.138	172.30.0.34/30	link-local
	i3	ge-0/0/4.168	172.30.0.22/30	link-local
	i4	ge-0/0/5.308	192.168.0.29/30	fc09:c0:ffee::1/126
	i5	ge-0/0/5.302	192.168.0.5/30	
		lo0.0	172.30.5.8/32	fd17:f0f4:f691:5::8/128

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 37 specifications. Make sure that OSPF is not running on the OoB management interface and on the AS external interfaces.

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

Table 37

Router	Interface	Area
R1	i1	0
	i2	0
	i3	0
	lo0.0	0
R2	i1	0
	i2	0
	i3	0
	lo0.0	0
R3	i1	0
	i2	0
	i3	0
	i4	1
	i5	1
	lo0.0	0
R4	i1	1
	i2	1
	lo0.0	1
R5	i1	1
	i2	1
	lo0.0	1
R6	i1	0
	i2	1
	i3	1
	i4	0
	i5	0
	lo0.0	0
R7	i1	0
	i2	0
	i3	0
	lo0.0	0
R8	i1	0
	i2	0
	i3	0
	lo0.0	0

- 2) Configure OSPFv2 only on R1 and R2 as shown in Table 38. Enable OSPFv2 on Route Reflector.

Table 38

Router	Interface	Area
R1	i4	0
R2	i4	0

- 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 i4 and R6 i6 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.

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 may 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 may not 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 39.

Table 39

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	

- 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 may 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.
- 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.

- 1) Configure LDP interfaces as shown in Table 40. Enable LDP on Route Reflector.

Table 40

Router	Interface
R1	i4
R2	i4
R3	i4
	i5
R4	i1
	i2
R5	i1
	i2
R6	i2
	i3

- 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 41. Enable RSVP message aggregation.
- 6) Configure link administrative groups as shown in Table 41.

Table 41

Router	Interface	Color
R1	i1	green
	i2	blue
	i3	purple
R2	i1	green
	i2	blue
	i3	purple
R3	i1	purple
	i2	blue
	i3	green
R6	i1	blue
	i4	purple
	i5	green
R7	i1	blue
	i2	green
	i3	purple
R8	i1	purple
	i2	blue
	i3	green

- 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.
- 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 paths 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 infrastructure including dual-homed customer sites and VLAN normalization. Ensure that your configuration meets the following criteria.

- 1) Configure L3VPN as shown in Table 42.

Table 42

Customer	Site	Router	PE-CE Protocol	Protocol details
CE1	S1	CE1-1	OSPF	Area 0
	S2	CE1-2	OSPF	Area 0
CE2	S1	CE2-1	BGP	AS 64600
	S2	CE2-2	BGP	AS 64600
	S3	CE2-3	BGP	AS 64600

- 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 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 43. The customer uses VLANs 600 and 601.

Table 43

Customer	Site	Router	VPLS	CE facing interface

			signaling	
CE3	S1	CE3-1	BGP	ge-0/0/3.601
	S2	CE3-2	BGP	ge-0/0/3.600

- 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.

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 44.

Table 44

Traffic Type	Criteria	Forwarding Class
VPN regular	DSCP 0b000000	I3vpn
VPN priority	DSCP 0b101110	I3vpn-priority

- 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 "I2vpn" forwarding class.
- 3) Make sure that traffic entering PE routers from L3VPN customers and classified as I3vpn-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 45.

Table 45

Forwarding Class	Queue	Scheduler
be	0	be-sc
I3vpn	1	I3vpn-sc
I2vpn	2	I3vpn-pri-sc
I3vpn-priority	3	I2vpn-sc
nc	4	nc-sc

- 6) Configure schedulers with parameters shown in Table 46.

Table 46

Scheduler	Parameter	Value
be-sc	Priority	low
	Transmit rate	remainder
	Buffer size	remainder
	Drop profile LP any	high-drop
I3vpn-sc	Priority	medium-low
	Transmit rate	20%
	Buffer size	20%
I2vpn-sc	Priority	medium-high
	Transmit rate	20%
	Buffer size	20%
	Drop profile LP low	low-drop
	Drop profile LP high	high-drop
I3vpn-pri-sc	Priority	high
	Transmit rate	10%
	Buffer size	5 msec

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

- 7) Configure drop profiles as shown in Table 47. Have a router to automatically build a smooth graph line based on the defined data points.

Table 47

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


- 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 48. Make sure that all PE and BGP ASBR routers mark both IPv4 and MPLS packet headers' CoS fields.
- 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.

Table 48

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	


A Full Day Lab Challenge II

This lab scenario represents a complete 8 hour lab challenge, aiming to prepare potential candidates for the JNCIE-SP lab exam. It contains all major topics found in the exam blueprint. This lab is also supported on iNETZERO rack rental service. 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

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.

Part 1: System Features

Task 1.1 Service Configuration

- 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 the SSH protocol, enable telnet access. 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.

Task 1.2: Centralized authentication management

- 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.

Task 1.3: Local user configuration

- 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 week days 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!
#####
```

Task 1.4: Active configuration archival and logging

- 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*.
- 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.

Task 1.5: Advanced Interface configuration and chassis features

- 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.
- 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.

Task 1.6: Advanced RE Protection

- 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 subsequences 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 subsequences 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

Part 2: Troubleshooting and Configuring IGP

Task 2.1: Troubleshooting

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

Task 2.2: Connectivity to OSPFv3 Area

- 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 level 1 area 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.

Task 2.3: RIP redistribution

- 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* stanza 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.

Task 2.4: Multilevel IS-IS configuration

- 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.

Task 2.5: Advanced IS-IS configuration

- 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.

Part 3: Troubleshooting and Configuring BGP

Task 3.1: iBGP design and configuration

- 1) Configure the necessary connectivity to the dedicated route-reflector. It already has partial configuration, however you have to 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.

Task 3.2: eBGP peers and configuration

- 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-next-hop*.
- 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 have to 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 have to discover the eBGP AS and configure the eBGP session using md5 authentication password *jncie*.

Task 3.3: BGP policy configuration

- 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.
- 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 effect 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 then 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.

Task 3.4: BGP general requirements

- 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.

Part 4: MPLS configuration

Task 4.1: MPLS and RSVP 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.
- 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 back 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 *r7-to-r3*. 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.

Task 4.2: MPLS and LDP configuration

- 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.

Part 5: IPv6 configuration

Task 5.1: Native 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.

Task 5.2: IPv6 tunneling configuration

- 3) To protect your customers , tunnel the IPv6 traffic between peers P1 and U3 in 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.

Part 6: MPLS VPNs configuration

Task 6.1: VPNA 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 have to 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 have to 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.

Task 6.2: VPNA multicast configuration

- 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.

Task 6.3: VPLS configuration


- 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 have to 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 if this limit is exceeded. Disable the mac learning function on the HUB site.

Task 6.4: L2VPN configuration


- 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.

A Full Day Lab Challenge III

This lab scenario represents a complete 8 hour lab challenge, aiming to prepare potential candidates for the JNCIE-SP lab exam. It contains all major topics found in the exam blueprint. This lab is also supported on iNETZERO rack rental service. 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.



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



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.

Part 1: System Features

Task 1.1 Service Configuration

- 1) SSH access is preconfigured as part of the initial configuration on all routers. Limit the SSH connections to 2 at a time and not more than 2 connection attempts per minute.
- 2) Enable NETCONF protocol over SSH using the standard port – TCP 830.
- 3) Configure SNMPv2c with community *superlab* on all routers, polled from management system with address 10.10.10.1/32. Since SNMPv2c is using transmission in clear text, make sure that SNMP traffic is accepted only over the management interface fxp0.0. Ensure that the reachability to 10.10.10.1/32 is provided even if the *rpd* process is not yet running. Assume the IP address 10.10.1.200 is used as the gateway in the management segment.

Task 1.2: Centralized authentication management

- 4) Configure all devices to use TACACS server located at 10.10.10.1 for authentication. The communication with the TACACS server should be encrypted with password *jncie123*.
- 5) Ensure that local user database is used for authentication only when the TACACS server is unreachable.

Task 1.3: Local user configuration

- 6) Configure user *noc1* and *noc2* using local passwords *noc123* and *noc456* respectively. Any session of users *noc1* and *noc2* should be disconnected after 5 minutes of inactivity. The users should have full access to the device without the possibility to manipulate or list files on the local storage. Entering the shell mode should be restricted as well.
- 7) Users must be locked out of the device at least for 3 minutes if they fail to provide a correct password 3 times.
- 8) Make sure that system alarms are displayed during login of any user not part of default login classes. You have to ensure that users of any current or future custom login classes would have this requirement enabled.

Task 1.4: Active configuration archival and logging

- 9) You should ensure that all device transfer the configuration every 2 hours to SCP server located on address 10.10.10.1. Using username *lab* and password *workbook*. The configuration files must be copied to destination folder */home/archive* on the server.
- 10) Configure all routers to send all log messages related to the *rpd* process and the kernel to syslog server 10.10.10.1.
- 11) Amend the default syslog configuration in the way that all commands run by any user should be saved in a file *interactive-commands* stored locally. To collect enough data for audit trail, but in the same to conserve space, make sure that no more than 5 files each of 3MB in filesize are kept in the storage.

Task 1.5: Advanced Interface configuration and chassis functions

- 12) Configure an aggregated Ethernet interface connecting R3 and R4. Assume that configuration template requires the interface to have the name *ae3*. Ensure the configuration template continues to be used successfully without adjusting it
- 13) Configure the aggregated Ethernet interface to monitor the member links. Due to the fact, that vMX is software based device increase the monitoring interval to 30s.
- 14) Enable support for switching label packets on all logical interfaces on all routers except the RR. You must use the most efficient way configuration-wise to accomplish this. This change should not affect the management interface.

Task 1.6: Advanced RE Protection

15) You must protect the control plane of all routers, including the route-reflector router.

Configure an input firewall filter using the following criteria:

- Allow the following IPv4 protocols: SSH, SNMP, BFD, OSPF, RIP, PIM, LDP, RSVP, TACACS+, ping and traceroute. If needed, also only allow protocols required for tasks.
- Allow IPv6 protocols required by the subsequences tasks: OSPFv3, BGP, RIPng, ping and traceroute.
- Ensure that SNMP packets are accepted only from already configured hosts, as well as any added in the future without requiring a change in the firewall filter configuration.
- You must accept BGP sessions only from configured neighbors. Any additional neighbors must be automatically accepted. In addition, you have to ensure that only BGP packets destined to local interface addresses are accepted.
- Limit the traffic to the SSH, SNMP, TACACS+, ping and traceroute protocol to not more than 5Mbps.
- All other traffic must be silently discarded. Make sure that the last 400 packets are logged in the local volatile memory.

Part 2: Troubleshooting and Configuring IGP

Task 2.1: Troubleshooting

- 1) Ensure all of the OSPF neighbors connections are successfully established (Full state). Make sure the following criteria are met:
 - All neighbors that require authentication must authenticate successfully.
 - Area 0 routers should not exchange any Network LSAs.
 - Ensure that all routers in Area 0 can reach the loopback address of all other routers in the Area. Make sure that they are advertised as internal OSPF routes.

Task 2.2: RIP redistribution

- 2) Make sure that R4 and R7 are able to receive RIP and RIPng routes from DC-1. Advertise the most specific summary route representing these routes to the OSPF domain.
- 3) DC-1 must have reachability to the loopback addresses of all routers in the OSPF domain.
- 4) Avoid changing the default preference value of any routing protocol if you need to resolve a routing loops or suboptimal routing.

Task 2.3: Multi-area OSPF configuration

- 5) Configure multi-area OSPF network as described on the diagram. Ensure that the following criteria are met:
 - No Type 5 LSAs are allowed in Area 50. Ensure that only R7 is advertising a summary of the DC-1 routes as external in Area 50. The loopback addresses of R3 and R7 routers must appear as internal in Area 50 routing table.
 - To spread the traffic between both ABRs of Area 50, make sure that R5 is the preferred transit node for IPv4 traffic to the Backbone area, whereas IPv6 traffic prefers the path through R4. However, you must preserve full reachability in case one of the ABRs fail. Do not advertise anything apart from the IPv4 and IPv6 RIP prefixes summary routes and the R3 and R7 loopback addresses out of Area 50.
 - Keep the same traffic separation between R4 and R5 from the backbone area to Area 50.
 - R8 is restricted to run single OSPF process. However, it should still be able to exchange IPv4 and IPv6 with rest of the network.
 - R8 must receive only the loopback address routes of the other routers in the network plus a summary route of DC-1. Due to limited size of the RIB, make sure that the IPv6 routes appear as summarized in the R8 routing table. Including the IPv6 loopback addresses.
 - Make sure that rest of the network can reach the R8 IPv4 and IPv6 loopback addresses.

Task 2.4: Advanced OSPF configuration

- 6) Ensure all links with a bandwidth of 1Gbps have an OSPF cost of 100 and those with a bandwidth of 2Gbps have an OSPF cost of 50. Do not configure the cost manually on the interfaces.
- 7) Make sure that all non backbone OSPF routers wait at least 360 seconds after the OSPF process has started before they forward traffic.
- 8) Configure R1 to protect the link to R6 using the alternate path through the network. R1 must have that path into the forwarding table ready to be used in case the direct link to R6 fails.

Part 3: Troubleshooting and Configuring BGP

Task 3.1: Internal BGP design and configuration

- 1) Configure the necessary connectivity to the dedicated route-reflector. It already has partial configuration, moreover you have to ensure that appropriate IGP configuration is in place for providing connectivity to and from the route-reflector.
- 2) Design iBGP network with two clusters. Both of the clusters must reside on the dedicated route-reflector device named respectively – *cluster-1* and *cluster-2*. All Area 0 routers should share one the cluster and the rest of the routers the other. iBGP sessions must tolerate a single link failure. Configure only a single IPv4 iBGP session on each router to the route reflector.
- 3) All iBGP sessions must use the BFD protocol for faster failure detection.
- 4) Advertise a summary IPv4 route 10.50.250.0/28 representing the loopback addresses of all routers into iBGP on R4 and R5 only. Ensure that all routers have the summary present in the routing table learnt from BGP.

Task 3.2: external BGP peers and configuration

- 5) Configure eBGP sessions with the Upstream (U1, U2 and U3) peers using the diagram IP addresses. U3 provides requires a separate eBGP session for advertising IPv6 prefixes. Make sure you receive IPv4 and IPv6 routes from upstream U3.
- 6) Configure an eBGP session with the Partner (P1) peer using the diagram IP addresses. Make sure that you receive IPv4 and IPv6 prefixes from P1 over single IPv4 eBGP session.
- 7) Configure an eBGP session with the Customer (C1 and C2) peers using the diagram IP addresses. C1 supports two separate eBGP sessions, one for IPv4 and other for IPv6, however the connectivity IPv4 and IPv6 addresses are not known. You have to find a way to discover them.
- 8) The C2 customer is recently migrated from other provider and is currently configured to peer with AS7862. Until C2's configuration is corrected, make sure it successfully establishes an eBGP session with R2. To prevent C2 prefixes being discarded due to loop protection make sure AS7862 is not added to the AS path.
- 9) Ensure all eBGP neighbors state changes are logged to syslog.

Task 3.3: BGP policy configuration

- 10) All IPv4 prefixes received from Upstream, Partner and Customer peers must be active and reachable on all routers in your AS unless they are filtered by policy in some of the subsequent tasks.
- 11) Advertise a summary route representing all your internal loopback IPv4 address to all external eBGP peers. Make sure that you avoid amending policies where possible.
- 12) All three Upstream peers are advertising the same prefixes. Ensure that all internal routers are selecting the best path to the Upstream peer based on the MED attribute. Apart from R3 which has to prefer the link to U1 when sending traffic to AS 87.109.

- 13) Accept only prefixes originated from P1 autonomous system on R7. Traffic from P1 must not use your AS for transit to Upstreams or Customers networks. You are allowed to do changes only to P1 policies to accomplish this.
- 14) IPv4 Customers prefixes should be advertised out of your AS to the Upstream peers only. The peering agreement with U3 does not allow you to send any routes with communities to AS 9687.
- 15) The ingress traffic from AS 87.109 should prefer the link from U2 to R3. Ensure traffic going to C1 prefixes prefer the path through R5. U3 peer must be the least preferred exit point from your AS.
- 16) Advertise all IPv4 Upstream prefixes to C1 and C2 except routes with mask higher than /28 only. Make sure the MED value is preserved and sent to the Customers (C1 and C2).
- 17) Because of the peer agreement with the Customers (C1 and C2) all routes advertised by their AS must be tagged with community 65457:2021 for C1 and 65456:2022 for C2. You must use a policy with a reject action to accomplish this. Use these communities to disable the communication between C1 and C2.
- 18) Do not accept IPv4 routes that are part of RFC1918 or any default routes from any of the external peers.
- 19) Ensure that 4byte AS numbers are displayed in as-dot notation format on all routers.

Task 3.4: BGP general requirements

- 20) No hidden (except filtered by applied policies) or unresolved routes should be present in AS 5673.873.
- 21) No static routes are allowed if not explicit permitted by a particular task.
- 22) Ensure optimal routing in the network.

Part 4: MPLS configuration

Task 4.1: MPLS and RSVP configuration

- 1) Enable the RSVP protocol on all routers in the OSPF domain apart from R8 and the route-reflector. Ensure that additional information like link coloring and bandwidth availability is exchanged.
- 2) Configure the appropriate colors of the links according the diagram. Also configure the available bandwidth for the LSPs as follows:
 - GOLD = 700Mbps.
 - SILVER = 500Mbps.
 - BRONZE = 300Mbps.
- 3) Configure the LSPs using the criteria below:
 - Configure LSPs named *r1-to-r5* and *r5-to-r1* connecting R1 with R5 without using cspf. However, you need to ensure they cross the link between R1 and R6. Protect those LSPs from failure of the link between R1 and R6. You have to ensure that the protecting LSPs have at least 100Mbps available bandwidth.
 - Configure LSPs named *r2-to-r5* and *r5-to-r2* providing connectivity between R2 and R5. Make sure to restrict those LSPs to use links with maximum available bandwidth. Protect the LSPs with one-to-one backup method not shared by other LSPs. The protecting LSPs should not be restricted by any means.
 - Configure LSP between R4 and R6 named *r4-to-r6* and *r6-to-r4* having a secondary path different from the primary. Make sure that the primary path is established only over links marked with color SILVER, whereas the secondary path must be using links marked with GOLD and BRONZE. The LSPs should have at least 250Mbps available bandwidth. If the primary path is preempted for any reason, ensure that the secondary path is established beforehand. However, to preserve resources on the routers, the secondary path must not be signaled if not needed.
 - Configure LSPs from R7 to R4 and R6 with primary path following the shortest route through the network. Protect the LSPs with secondary path avoiding completely the primary path. Make sure that the secondary path is signaled and ready for use. Configure LSPs from R6 and R4 to R7 without restrictions.
- 4) Make sure that all devices use reduction of soft-state RSVP messages to increase the scalability together with reliable message delivery inside OSPF Area 0.

Task 4.2: MPLS and LDP configuration

- 5) Assume that R8 supports only the LDP protocol. Configure R5 and R6 to provide MPLS transport gateway functionality to the rest of network for R8. CEs attached to R8 should be able to exchange labeled packet with the CEs behind R1 and R2.
- 6) Make sure that LDP neighbors are not sending traps during transition from one state to other.

Part 5: IPv6 tunneling and Multicast configuration

Task 5.1: IPv6 tunneling configuration

- 1) Configure IPv6 connectivity between P1, C1 and U3 peers. Ensure IPv6 is encapsulated in labeled packets within your AS. Make sure that IPv6 routes are exchanged between all three peers. You are not allowed to configure routing-instances of any type or enable native IPv6 transport in the core of your network.
- 2) Avoid changing the resolution table or enabling MPLS on the interfaces of the route-reflector to resolve any hidden routes on the route-reflector. You are not allowed to have any routes except the Loopback addresses in inet.3 on RR device.

Task 5.2: Multicast configuration

- 3) Configure all additional interfaces to connect the Src1, Rcv1 and Rcv2 to the network. Make sure those interfaces are reachable by all routers.
- 4) Configure PIM protocol on all routers in OSPF Area 0 and Area 50. Ensure no PIM neighbor can be established on the management interface.
- 5) Configure IP address 192.168.50.50 as RP to provide support for ASM multicast receivers. However, to eliminate the possibility of RP failure dedicate R1 and R2 as RPs in the network. You are not allowed to use different protocols apart from PIM to accomplish this requirement.
- 6) Rcv1 is statically joined to group 239.255.255.2. Whereas Rcs2 is configured for source specific multicast, joined to group 239.255.255.3. Make sure multicast transmission from the source to the receivers is possible.

Part 6: MPLS VPNs configuration

Task 6.1: VPNA configuration

- 1) Configure a Layer 3 VPNA connecting the sites: VPNA-CE1, VPNA-CE2, VPNA-CE3 and VPNA-CE4. Sites CE1 and CE2 peer using BGP protocol. Whereas sites CE3 and CE4 are peering with the corresponding PEs using OSPF Area 0. Use your local AS number to define the route-distinguisher and route-target communities. Make sure you receive two prefixes from each site and share them between the sites.
- 2) A backdoor connection is shared between CE1 and CE2. Since that connection is intended for backup link only, make sure that the L3VPN connection is preferred for sending traffic between CE1 and CE2. The backdoor link is configured with MED of 300.
- 3) The routes received by sites CE3 and CE4 must be marked uniquely per site. All routes received from MPLS backbone must be advertised to the OSPF sites as external. However, those routes must not be recognizable as coming from the AS 5673.873 when examining the CE3 and CE4 routing tables.

Task 6.2: VPNB configuration

- 4) Configure Layer 3 VPNB connecting sites: VPNB-CE1, VPNB-CE2 and VPNB-CE3. Site CE3 is connected through 3rd party ISP1 interconnecting with your AS at R1. Use BGP as PE-CE protocol for both sites, connected directly to your network, using AS 65100. Make sure you use your local AS number for configuring the route-target communities and the loopback IP address for route-distinguisher. Ensure you receive a number of prefixes from each site.
- 5) You have to provide connectivity to and from AS 87.109 for site CE1. Make sure you do not advertise additional routes to the other VPN sites.
- 6) Since site CE3 is connected to 3rd party provider ISP1, you must extend the VPNB using inter-provider method that will provide you with a way to filter routes received from ISP-1 based on target-community. However, do not use routing policy on R1 providing connectivity to CE3. The remote PE is using route-target community *target:4356:500*.
- 7) Ensure that all BGP sessions in VPNB can send and receive packets larger than the default 512-byte maximum segment size.
- 8) Make sure that traffic coming from VPNB sites with precedence bits *010*, *011*, *100* and *100* is classified as *assured-forwarding*. Ensure 20% high priority bandwidth reservation for that traffic in your network, while keeping the default reservation for rest of the traffic.

Task 6.3: VPLS configuration


- 9) Configure a VPLS between sites VPLS-1 and VPLS-2, using the interfaces as shown in the physical diagram. Since site VPLS-2 has connectivity to both R6 and R8, dedicate R8 as primary PE and R6 as backup PE. You have to use LDP signaled VPLS with automatic site discovery, hence you are not allowed to configure the sites explicitly.
- 10) One more site attached to R2 should be connected to the VPLS VPN. Assume that R2 supports only L2 circuit VPN. Interconnect the L2VPN-3 site with the VPLS VPN at R8 as primary and R6 as back-up connection.

- 11) Protect the VPLS VPN from excessive broadcast and multicast traffic send by the users attached behind the VPLS sites. Ensure that not more than 30Mbps of such traffic can be send by a single site.
- 12) Make sure all 3 sites are able to communicate. However, you are not allowed to manipulate the VLAN tags directly on the interfaces.

Appendix 1: Additional Theory

OSPF adjacency troubleshooting

In this section we will 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;
        }
      }
    }
  }
}
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@MX1# 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 if you have an OSPF adjacency with MX2 on interface ge-0/0/1.0

```
root@MX1# run show ospf neighbor interface ge-0/0/1.0
[edit]
```

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

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

```
root@MX1# set protocols ospf traceoptions file ospf
root@MX1# set protocols ospf traceoptions flag all

root@MX1# run monitor start ospf

[edit]
root@MX1#

*** 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
```

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

Change MX1's OSPF configuration to reflect MX2's settings

<http://www.inetzero.com> - Copyright 2018 iNET ZERO. All rights reserved

For personal non commercial use only – do not distribute - Steganoid=ON

```

root@MX1# rename protocols ospf area 0 to area 99
root@MX1# set protocols ospf area 0.0.0.99 interface ge-0/0/1.0 hello-interval 2
root@MX1# 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 in our example although we did configure the dead-interval it is actually not needed.

5) Verify OSPF adjacency with MX2(Venus)

```

root@MX1# 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@MX1# 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

```

It looks like interface ge-0/0/1.0 on MX2 has been configured with the link type to “p2p”, since no DR/BDR election is desired on ge-0/0/1.0 interface. After all it’s a direct connection between the devices. Let’s change the OSPF interface type to “p2p” on our ge-0/0/1.0 interface.

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

```

root@MX1# set protocols ospf area 0.0.0.99 interface ge-0/0/1.0 interface-type p2p

```

```

root@MX1# 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 others 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 where 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

```

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

```
root@MX1# 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

9) 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@MX1# set interfaces ge-0/0/1 mtu 9192

or

root@MX1# set interfaces ge-0/0/1.0 family inet mtu 9178

root@MX1# commit
commit complete
```

10) Verify if the OSPF adjacency is established

```
root@MX1# 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.

11) Verify if we receive OSPF routes from MX2

```
root@MX1# 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, 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@MX1# 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 is able to 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@MX1# show protocols bgp
traceoptions {
  file bgp;
  flag open;
}
group ebgp {
  neighbor 172.30.0.2 {
    peer-as 1;
  }
}
```

Check BGP adjacency with MX2

```
root@MX1# 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 neighborhood with MX2 is not established.

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

```
root@MX1# 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@MX1# 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@MX1# run show bgp summary
Groups: 1 Peers: 1 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...
172.30.0.2      64555      33      33      0      2      3:36 0/0/0/0      0/0/0/0
```

```
root@MX1# 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!

BGP IPV6 NLRI over IPV4 peering

In the following example we will 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@MX1#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;
      }
    }
    family inet6 {
      address 2001:1111:1111:1111::1/128;
    }
  }
}

root@MX1# show protocols bgp
group ebgp {
  type external;
```



```

family inet {
  unicast;
}
family inet6 {
  unicast;
}
export myloopback;
neighbor 172.30.0.2 {
  peer-as 2222;
}
}

root@MX1# show policy-options policy-statement myloopback
from interface lo0.0;
then accept;

root@MX1# show routing-options
autonomous-system 1111;

```

Please note that we configured an IPv4 neighborhood 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@MX1# 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
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 in the same way 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

```

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

```

root@MX1# 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@MX1# set protocols bgp traceoptions file bgp
root@MX1# set protocols bgp traceoptions flag route
root@MX1# set protocols bgp traceoptions flag
root@MX1# 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 a 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 scenario's how to configure IPv6 NLRI exchange over and IPv4 peering.

Option 1:

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@MX1# 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@MX1#set system allow-v4mapped-packets
```

Check if we receive the ipv4 loopback address from MX2.

```
root@MX1# 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
```

<http://www.inetzero.com> - Copyright 2018 iNET ZERO. All rights reserved

For personal non commercial use only – do not distribute - Steganoid=ON

```
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@MX1#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
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@MX1#delete interfaces ge-0/0/1.0 family inet6 address ::ffff:172.30.0.1/126
root@MX1# delete system allow-v4mapped-packets
Root@MX1#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@MX1# 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@MX1# 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.

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@MX1# set protocols bgp group ebgp accept-remote-nexthop
root@MX1# 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., but at least we should see the prefix learned from MX2, but hidden as the next-hop :ffff:172.30.0.2 is not reachable.

```

root@MX1# 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@MX1# show policy-options policy-statement fixnexthop
from protocol bgp;
then {
    next-hop 2001:aaaa:bbbb::2;
}
[edit]
root@MX1# set protocols bgp group ebgp import fixnexthop

[edit]
root@MX1# 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@MX1# 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.

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 requirement(s):

- **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 signalling 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@MX1# 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
  
```

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@MX1# 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

```

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 and import inet.0 table 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 192.168.1.0/24 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
```

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 are able to 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 it 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@Inetzero# show protocols bgp group ibgp
type internal;
local-address 192.168.1.1;
family inet {
    unicast;
}
neighbor 192.168.1.2;

lab@Inetzero# 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.

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@Inetzero# 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@Inetzero# 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@inetzero# 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.

Another great JUNOS feature is “**apply-flags omit**”. With this feature its 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 its just 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 our re-protect firewall filter.

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

Verify our re-protect filter

```
lab@Inetzero# show firewall family inet
```

```
filter re-protect { /* OMITTED */ };
```

```
lab@Inetzero#
```

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@Inetzero# 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@Inetzero# 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 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 2 : Topology diagrams

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































Chapter 5 - L3VPN 1





Chapter 6 - L2VPN and VPLS 1


<http://www.inetzero.com> - Copyright 2018 iNET ZERO. All rights reserved
For personal non commercial use only – do not distribute - Steganoid=ON


<https://t.me/learningnets>













Chapter One solutions: General System Features

This chapter provides a detailed walkthrough focusing on the configuration of initial system services, management access to the devices and initial interface configuration. All 8 devices have to 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 plane 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 plane (RE) running JUNOS and the data plane (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 login without a password. Login with user *root* places you in to the UNIX shell.

```
Amnesiac (ttyu0)
Login: root
--- JUNOS 11.4R8.4 built 2013-05-23 05:53:07 UTC
root@mx1%
```

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@mx1% cli
root@mx1>
```

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, the prompt changes to #.

```
root@mx1> configure
Entering configuration mode

[edit]
root@mx1#
```


In *configuration* mode the configuration statements are hierarchically organized so they are easy 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.

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 txt 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) Using iNET ZERO rack rental service, the initial JNCIE-SP workbook configurations have to be loaded on the devices before you can start Configure the tasks. Ensure you also load the VR device configurations.

```
root# load override terminal
[Type ^D at a new line to end input]
system {
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
```

```

}
security {
  forwarding-options {
    family {
      inet6 {
        mode packet-based;
      }
      mpls {
        mode packet-based;
      }
      iso {
        mode packet-based;
      }
    }
  }
}
^D
load complete

[edit]
root#

```

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 terminal** is used to replace the complete configuration with the initial iNET ZERO configuration file pasted from the terminal.

NOTE: You can save the complete candidate configuration or parts of it with the **save** command .

- 2) This step requires Configure hostnames according the 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** command to load the configuration on the other devices after changing the router specific parameters.

Configure host name with following command:

```

[edit]
root# set system host-name Sun

```

NOTE: We will not commit the configuration after every single changes configured. Commit will be performed after several commands entered.

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

To access the CLI via ssh and telnet you must first configure a OoB network interface. Most Juniper Networks devices have a dedicated interface for OoB management. The interfaces on iNET ZEROs rack rental service that should be used for this purpose are listed in Table 2.

First we will configure the interface descriptions.

```
root# set interfaces ge-0/0/0 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.

```
root# set interfaces ge-0/0/0.0 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 *system services* stanza.

```
[edit]
root# set system services ssh

[edit]
root# set system services telnet

[edit]
root# set system services telnet

[edit]
root# commit

commit complete

[edit]
root@Sun#
```

NOTE: At this point we perform the **commit** command. The host name changes to the one configured.

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.

```
root@Sun# 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** stanza. 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.

```
root@Sun# set routing-options static route 10.10.10.0/24 next-hop
10.10.1.2 no-readvertise
```

```
[edit]
root@Sun# commit
commit complete
```

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

```
[edit]
root@Sun# 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 ge-0/0/0.0, 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** stanza has to be configured. When the RPD process has started this route is removed from the local routing and forwarding tables.

```
root@Sun# set system backup-router 10.10.1.254 destination
10.10.10.0/24
```

```
[edit]
root@Sun# 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.

```
root@Sun# set system name-server 10.10.1.100
```

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

- 7) Set the time zone location of the device.

```
root@Sun# set system time-zone Europe/Amsterdam
```

- 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.

```
root@Sun# set system ntp boot-server 10.10.1.100
```

```
[edit]
root@Sun# set system ntp authentication-key 1 type md5 value
workbook
```

```
[edit]
root@Sun# set system ntp server 10.10.1.100 key 1
```

```
[edit]
root@Sun# set system ntp trusted-key 1
```

```
[edit]
root@Sun# 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:

```
root@Sun# run show ntp status
status=0664 leap_none, sync_ntp, 6 events, event_peer/strat_chg,
version="ntpd 4.2.0-a Sat Aug 25 08:03:34 UTC 2012 (1)",
processor="octeon", system="JUNOS11.4R5.5", leap=00, stratum=4,
precision=-17, rootdelay=13.660, rootdispersion=280.927,
peer=13220,
refid=10.10.1.100,
```

```

reftime=d5649d66.4532aca9 Thu, Jun 13 2013 21:35:02.270, poll=6,
clock=d5649de9.248a6fe7 Thu, Jun 13 2013 21:37:13.142, state=4,
offset=-0.029, frequency=15.367, jitter=36.987, stability=0.000

root@Sun# run show ntp associations
remote          refid    st t when poll reach  delay  offset
jitter
=====
*10.10.1.100    46.19.33.5 3 - 10  64  17  1.255 -0.029
1.848

```

The * shows that server is selected for synchronization.

NOTE: The outputs above are *operational* display commands, but 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 OS provides the capability of automatically uploading the configuration to a FTP or SCP server. To meet the task requirements the configuration must be archived each time the configuration is committed.

```

root@Sun# set system archival configuration transfer-on-commit

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

[edit]
root@Sun# 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 as they have an other 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.

- 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** stanza 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]
root@Sun# set system authentication-order [ radius password ]
```

Then the RADIUS server configuration with provided settings.

```
[edit]
root@Sun# set system radius-server 10.10.1.100 secret workbook

[edit]
root@Sun# set system radius-server 10.10.1.100 timeout 2

[edit]
root@Sun# set system radius-server 10.10.1.100 retry 1

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



```
root@Sun# commit
commit complete
```

NOTE: Another powerful tool of the JUNOS CLI is demonstrated. **show | compare** 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. **show | compare rollback #** 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 Configure a local user also a **class** must 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 be member of single class only.

Four predefined classes are available in JUNOS CLI. 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

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 user and assign it to super-user class.

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

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

```
[edit]
root@Sun# 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
```

From this point we recommended to use user *lab* when accessing the iNET ZERO devices.

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

The first user should have any name with *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 accounts. Local user template account 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@Sun# set system login class limited permissions [ view view-
configuration ]
```

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

```
[edit]
lab@Sun# 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@Sun# 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. **deny-configuration** and **allow-configuration** can be used for *configuration* commands. Regular expressions for matching the respecting commands can be used. JUNOS applies the **deny** statements before **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@Sun# set system login class privileged permissions all
```

First you attach all permissions set to new class.

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

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

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

```
[edit]
lab@Sun# 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@Sun> show cli authorization
Current user: 'noc' class 'privileged'
Permissions:
  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

```

13) Configure syslog as shown in 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 */var/log/*. 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 in the *system syslog* stanza and delete the whole *syslog* configuration.

```

[edit]
lab@Sun# show system syslog

```

```

user * {
    any emergency;
}
user ops {
    any warning;
}
file messages {
    any notice;
    authorization info;
}
file interactive-commands {
    interactive-commands any;
}

```

```

[edit]
lab@Sun# delete system syslog

```

```

[edit]
lab@Sun# commit
commit complete

```

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@Sun# 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@Sun# set system syslog host 10.10.1.100 interactive-commands any

```

```

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

```

```

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

```

```
[edit]
lab@Sun# set system syslog user * any critical
```

NOTE: The asterisk * sign represents all users.

14) Last step of this 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 needs to be adjusted. The log files have to be rotated until 3 files are created with a size of 100kb.

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* stanza. For this task you can use the global configuration.

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

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

This was the last step from Task 1. All changes made so far have been done on a single router (Sun) only. To speed up the configuration on rest of the devices you can use the active configuration file from Sun. 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 build 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 an 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.

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

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

```
[edit]
lab@Sun# 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 SNMP v3 must provide read access only is accomplished with **read-view** configured under the **security-level** stanza.

- 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.

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

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

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

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

- b. Configure the destination parameters of the management station.

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

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

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

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

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

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

[edit]
lab@Sun# commit
commit complete

```

Display all SNMPv3 configurations.

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

Local engine ID: 80 00 0a 4c 01 0a 0a 01 01
Engine boots:          3
Engine time:          12780 seconds
Max msg size:         65507 bytes

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

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

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
S1           10.10.1.100          162  S1-paramete type
active                               nonvolatile

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
traps         tag-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
```


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

```
[edit]
lab@Sun# 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
```

At the end of this step you can again load the configuration to the other devices to save time.

Task 3: Firewall filters

This task requires the configuration 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 a group of match conditions and one or multiple actions or non-terminating actions. If no traffic is matched from any of the terms there is an implicit rule that discards all traffic. This last rule is hidden and not displayed in 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 recognizes it as a *family inet* filter matching IPv4 traffic by default.

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@Sun# set firewall family inet filter protect-re term ah from
protocol ah

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

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

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

[edit]
```

```
lab@Sun# 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@Sun# set firewall family inet filter protect-re term vrrp from
protocol vrrp
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

NOTE: prefix-lists are configured under the **policy-options** stanza. They are also used for routing-policies. The wildcard "*" can be used to match "all" BGP groups and match "all" neighbor IP address.

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

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

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

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

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

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

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

```
[edit]
```

```

lab@Sun# set firewall family inet filter protect-re term ntp from
port ntp

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

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

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

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

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

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

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

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

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

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

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

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

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

[edit]

```

```

lab@Sun# set firewall family inet filter protect-re term ssh from
source-address 10.10.1.0/24

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

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

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

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

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

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

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

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

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

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

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

[edit]
lab@Sun# 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 $action$;
  }
}

```

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

```

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

```

```

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

```

```

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

```

```

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

```

```

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

```

```

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

```

```

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

```

```

[edit]

```

```
lab@Sun# set firewall family inet filter protect-re term
traceroute then policer re-policer
```

```
[edit]
```

```
lab@Sun# set firewall family inet filter protect-re term
traceroute then accept
```

```
[edit]
```

```
lab@Sun# commit
commit complete
```

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

Verify policer operation. Policar should rate-limit ping and traceroute traffic above 100kbit/s. You can use ping with big packets to simulate increased icmp traffic. Pinging from R2 interface IP address of R1.

```
lab@Sirius> ping 172.30.0.1 size 1400 rapid count 100
PING 172.30.0.1 (172.30.0.1): 1400 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 172.30.0.1 ping statistics ---
100 packets transmitted, 84 packets received, 16% packet loss
round-trip min/avg/max/stddev = 1.265/2.590/27.761/3.434 ms

lab@Sirius> ping 172.30.0.1 size 1400 rapid count 100
PING 172.30.0.1 (172.30.0.1): 1400 data bytes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
--- 172.30.0.1 ping statistics ---
100 packets transmitted, 85 packets received, 15% packet loss
round-trip min/avg/max/stddev = 1.244/2.051/7.872/1.738 ms
```

You 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@Sun# run show firewall filter protect-re

Filter: protect-re
Counters:
Name                               Bytes
Packets
dropped-packets                    900
18
Policers:
Name                               Bytes
Packets
re-policer-icmp                    28
re-policer-traceroute              0
```

- 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@Sun# set firewall family inet filter protect-re term last then
count dropped-packets
```

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

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

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

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

NOTE: When applying a filter to the RE always use **commit confirm <min.>**. This way if you make mistake, the router will rollback the configuration after the specified time if a second **commit** is not issued.

Verify the operation of 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@Canopus> telnet 10.10.1.1
Trying 10.10.1.1...
Connected to 10.10.1.1.
Escape character is '^]'.

Sun (ttypl)

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@Sun# run show firewall filter counter dropped-packets protect-
re

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

```
[edit]
lab@Sun# run show firewall log
Log :
Time      Filter  Action Interface  Protocol  Src Addr  Dest
Addr
18:03:46  pfe      D      ge-0/0/0.0   ICMP      10.10.1.2
10.10.1.1
18:03:45  pfe      D      ge-0/0/0.0   ICMP      10.10.1.2
10.10.1.1
18:03:44  pfe      D      ge-0/0/0.0   ICMP      10.10.1.2
10.10.1.1
18:03:43  pfe      D      ge-0/0/0.0   ICMP      10.10.1.2
10.10.1.1
18:03:42  pfe      D      ge-0/0/0.0   ICMP      10.10.1.2
10.10.1.1
18:03:41  pfe      D      ge-0/0/0.0   ICMP      10.10.1.2
10.10.1.1
18:03:40  pfe      D      ge-0/0/0.0   ICMP      10.10.1.2
10.10.1.1
18:03:39  pfe      D      ge-0/0/0.0   ICMP      10.10.1.2
10.10.1.1
```

NOTE: Do not forget to activate the *icmp* term after you verified the filter.
When you are finished with this task you can add the configuration to the other devices in the lab.

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.

JUNOS Interfaces

The main function of JUNOS based devices is to forward traffic from one interface to one other. 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 **ge0/0/0.0** – 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**.

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 is 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 availability 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 – “Sun”

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

NOTE: By default no aggregated interfaces exist, they have to be created. Under the *chassis* stanza 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 LAG member interfaces. *Vlan-tagging* is supported on aggregated interfaces but have to be configured under the *ae0* group interface.

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

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

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

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

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

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

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@Sun# run show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 501
  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: Sendbroadcast-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:fefd: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.

Display the operation of the LACP LAG interface.

```
[edit]
lab@Sun# 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 have to configure the interface terminating two VLANs by creating different logical interfaces for each VLAN.

```
[edit]
lab@Sun# set interfaces ge-0/0/4 vlan-tagging

[edit]
lab@Sun# set interfaces ge-0/0/4 unit 114 description "R4
connection"

[edit]
lab@Sun# set interfaces ge-0/0/4 unit 114 vlan-id 114

[edit]
lab@Sun# set interfaces ge-0/0/4 unit 114 family inet address
172.30.0.5/30

[edit]
lab@Sun# set interfaces ge-0/0/4 unit 118 description "R8
connection"

[edit]
lab@Sun# set interfaces ge-0/0/4 unit 118 vlan-id 118

[edit]
lab@Sun# set interfaces ge-0/0/4 unit 118 family inet address
172.30.0.9/30

[edit]
lab@Sun# set interfaces ge-0/0/4 unit 118 family inet6
```

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 Loopback0 interface.

```
[edit]
lab@Sun# set interfaces lo0 unit 0 family inet address
172.30.5.1/32
```

```
[edit]
```

```
lab@Sun# set interfaces lo0 unit 0 family inet6 address
fd17:f0f4:f691:5::1/128
```

```
[edit]
lab@Sun# 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 Loopback0 interface has two IPv6 addresses as shown below.

```
[edit]
lab@Sun# 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 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@Canopus# set interfaces ge-0/0/4 vlan-tagging

[edit]
lab@Canopus# set interfaces ge-0/0/4 unit 200 description "DC1 LAN
1"

[edit]
lab@Canopus# set interfaces ge-0/0/4 unit 200 vlan-id 200

[edit]
lab@Canopus# 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@Canopus# set interfaces ge-0/0/4 unit 200 family inet address
172.30.1.1/24 vrrp-group 1 priority 150

[edit]
lab@Canopus# set interfaces ge-0/0/4 unit 200 family inet address
172.30.1.1/24 vrrp-group 1 authentication-type md5

[edit]
lab@Canopus# set interfaces ge-0/0/4 unit 200 family inet address
172.30.1.1/24 vrrp-group 1 authentication-key workbook

[edit]
lab@Canopus# 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@Canopus# 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: unit 200 has priority 150, which is higher than the default. 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 then the default value of 100 advertised by R4. In
that case the VRRP master router will switch to R4.

[edit]
lab@Canopus# set interfaces ge-0/0/4 unit 201 description "DC1 LAN
2"

[edit]
lab@Canopus# set interfaces ge-0/0/4 unit 201 vlan-id 201

[edit]
```

```
lab@Canopus# 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@Canopus# set interfaces ge-0/0/4 unit 201 family inet address
172.30.2.1/24 vrrp-group 2 authentication-type md5
```

```
[edit]
```

```
lab@Canopus# set interfaces ge-0/0/4 unit 201 family inet address
172.30.2.1/24 vrrp-group 2 authentication-key workbook
```

```
[edit]
```

```
lab@Canopus# 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@Arcturus# set interfaces ge-0/0/4 vlan-tagging
```

```
[edit]
```

```
lab@Arcturus# set interfaces ge-0/0/4 unit 200 description "DC1
LAN 1"
```

```
[edit]
```

```
lab@Arcturus# set interfaces ge-0/0/4 unit 200 vlan-id 200
```

```
[edit]
```

```
lab@Arcturus# 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@Arcturus# set interfaces ge-0/0/4 unit 200 family inet address
172.30.1.2/24 vrrp-group 1 authentication-type md5
```

```
[edit]
```

```
lab@Arcturus# 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@Arcturus# set interfaces ge-0/0/4 unit 201 description "DC1
LAN 2"
```

```
[edit]
```

```
lab@Arcturus# set interfaces ge-0/0/4 unit 201 vlan-id 201
```

```
[edit]
```

```
lab@Arcturus# 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@Arcturus# set interfaces ge-0/0/4 unit 201 family inet address
172.30.2.2/24 vrrp-group 2 priority 150
```

```
[edit]
```

```
lab@Arcturus# set interfaces ge-0/0/4 unit 201 family inet address
172.30.2.2/24 vrrp-group 2 authentication-type md5
```

```
[edit]
```

```
lab@Arcturus# set interfaces ge-0/0/4 unit 201 family inet address
172.30.2.2/24 vrrp-group 2 authentication-key workbook
```

```
[edit]
```

```
lab@Arcturus# 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
```

```
[edit]
```

```
lab@Arcturus# 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@Arcturus # 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@Canopus> 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 vip	172.30.1.1
172.30.1.254							
ge-0/0/4.201	up	2	backup	Active	D 3.407	lcl vip	172.30.2.1
172.30.2.254						mas	172.30.2.2

```
lab@Canopus> 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@Arcturus> 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 vip	172.30.1.2
172.30.1.254							
ge-0/0/4.201	up	2	master	Active	A 0.519	lcl vip	172.30.1.1
172.30.2.254							

```
lab@Arcturus> 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

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@Canopus# deactivate interfaces ge-0/0/4.123

[edit]
lab@Canopus# deactivate interfaces ge-0/0/4.136

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

On R3 both VRRP groups are in backup state:

```
[edit]
lab@Canopus# 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 vip	172.30.1.1
172.30.1.254							
ge-0/0/4.201	up	2	backup	Active	D 3.282	lcl vip	172.30.1.2
172.30.2.254							
						mas	172.30.2.2

```
[edit]
lab@Canopus# 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@Arcturus> 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 vip	172.30.1.2
172.30.1.254							

ge-0/0/4.201	up	2	master	Active	A	0.783	lcl	172.30.2.2
172.30.2.254							vip	

NOTE: Don't forget to activate both interfaces to rollback to the requested configuration.

For this task using the **load merge terminal** command to configure the other routers is not recommended. It is better that each device interface is configured separately because many parameters are different.

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.
- 8)
- *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 triggered 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@Sun> file copy ftp://lab:lab123@10.10.1.100/show-  
interfaces.slax /var/db/scripts/op/show-interfaces.slax
```

NOTE: You need to be logged in with *root* user to copy files to /var/db/scripts folder.

- 2) Download the commit script.

```
lab@Sun> 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.

```
lab@Sun> 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@Sun# set system scripts op file show-interfaces.slax
```

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

Verify the script:

This operation script modifies the output of the command *show interface terse*. The description of the interfaces are added to the output.

```
lab@Sun# run op show-interfaces
```

Interface	Admin Link	Proto	Local	Remote
ge-0/0/0.0	OoB management			
		inet	10.10.1.1/24	
sp-0/0/0.0		inet		
sp-0/0/0.16383		inet	10.0.0.1	--> 10.0.0.16
			10.0.0.6	--> 0/0
			128.0.0.1	--> 128.0.1.16
			128.0.0.6	--> 0/0
ge-0/0/1.0		aenet	--> ae0.0	
ge-0/0/2.0		aenet	--> ae0.0	
ge-0/0/4.114	R4 connection	inet	172.30.0.5/30	
ge-0/0/4.118	R8 connection	inet	172.30.0.9/30	
		inet6	fe80::226:8800:76ef:7584/64	
ge-0/0/4.32767				
ae0.0	R2 connection	inet	172.30.0.1/30	
		inet6	fe80::226:88ff:feef:7600/64	

b. Enable the commit script.

```
[edit]
lab@Sun# set system scripts commit file interface-mask-check.slax
```

```
[edit]
lab@Sun# 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@Sun# set interfaces ge-0/0/4.115 family inet address
172.30.4.1/20
```

```
[edit]
```



```
lab@Sun# commit
warning: The address of 172.30.4.1 has a mask of /20
on interface ge-0/0/4 unit 115
commit 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 a action 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

```
[edit]
lab@Sun# 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@Sun# 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----
*
*
```

NOTE: The operational command **file show <file name>** displays the content of the file.

Configure the event-policy.

```
[edit]
lab@Sun# set event-options policy sylog_if_description events
snmp_trap_link_up

[edit]
lab@Sun# set event-options policy sylog_if_description events
snmp_trap_link_down

[edit]
lab@Sun# set event-options policy sylog_if_description then event-
script syslog-int-desc-on-link-change.slax

[edit]
```

```
lab@Sun# set event-options policy snmptrap_if_description events
SYSTEM
```

```
[edit]
```

```
lab@Sun# set event-options policy snmptrap_if_description
attributes-match SYSTEM.message matches NEW_SNMP_TRAP_LINK
```

```
[edit]
```

```
lab@Sun# set event-options policy snmptrap_if_description then
raise-trap
```

```
[edit]
```

```
lab@Sun# commit
commit complete
```

NOTE: The first policy `syslog_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.

```
[edit]
```

```
lab@Sun# set interfaces ae0 disable
```

```
[edit]
```

```
lab@Sun# commit
commit complete
```

```
lab@Sun> show log jncie-sp-messages
Jun 20 20:31:22 Sun mgd[3470]: UI_CHILD_START: Starting child '/sbin/ifinfo'
Jun 20 20:31:22 Sun mgd[3470]: UI_CHILD_STATUS: Cleanup child '/sbin/ifinfo',
PID 3479, status 0
Jun 20 20:31:22 Sun mgd[3471]: UI_CHILD_STATUS: Cleanup child '/sbin/ifinfo',
PID 3478, status 0
Jun 20 20:31:22 Sun mgd[3471]: UI_LOGOUT_EVENT: User 'root' logout
Jun 20 20:31:22 Sun cscript: NEW_SNMP_TRAP_LINK_DOWN, Sun, , , ,
Jun 20 20:31:22 Sun mgd[3470]: UI_LOGOUT_EVENT: User 'root' logout
Jun 20 20:31:22 Sun cscript: NEW_SNMP_TRAP_LINK_DOWN, Sun, ae0.0, up, down, R2
connection
Jun 20 20:32:14 Sun mgd[1934]: UI_CMDLINE_READ_LINE: User 'lab', command 'exit '
Jun 20 20:32:14 Sun mgd[1934]: UI_DBASE_LOGOUT_EVENT: User 'lab' exiting
configuration mode
Jun 20 20:32:37 Sun 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 save the configuration.

```
[edit]
```

```
lab@Sun# save my_baseline.conf
```


Chapter Two solutions: 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, 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 a LSA received from a neighbor back to the same neighbor. If there is a need to limit the flooding of LSAs and SPF 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 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.
- **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 a 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 100Mbps or faster will have cost of 1 by default. The cost of each interface can be configured manually also.

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 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, as 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. Again

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. **Don't forget 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 segment. By default Keep-alive/Hello messages are send 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 from the files provided with the JNCIE workbook for Chapter Two.

NOTE: If you are using iNET ZERO rack rental service, don't forget to load the configuration on the **vr-device** as well.

```
[edit]
lab@Sun# load override terminal
```

2) Verify the OSPF configuration and adjacencies.

a. R1

```
[edit]
lab@Sun# 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;
    }
  }
}
```

NOTE: Protocols in JUNOS are configured under the `protocols` stanza.

OSPF interfaces configuration is applied under the `area` stanza. 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.

```
[edit]
lab@Sun# 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@Sun# 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.

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@Sirius# 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@Sirius# 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*.

```
[edit]
lab@Sirius# 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@Canopus# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/4.123 {
        authentication {
            md5 1 key "$9$T3A0MWx-b2y1vLNboaTz39tO"; ## 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 "$9$B6BE1Mg4ZDHmVw2aUH5TBIEyeW"; ## SECRET-DATA
        }
    }
}
```

List the active neighbors on R3.

```
[edit]
lab@Canopus# 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.

```
[edit]
lab@Canopus# 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@Arcturus# 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 "$9$GGiPQ1IcrK8tuBEyKx7GDf.fz"; ## SECRET-DATA
        }
    }
    interface lo0.0;
}
```

List the active neighbors on R4.

```
[edit]
lab@Arcturus# run show ospf neighbor
```

List all interfaces configured for OSPF.

```
[edit]
lab@Arcturus# 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.

e. R5 OSPF configuration.

```
[edit]
lab@A-Centauri# 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@A-Centauri# 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@A-Centauri# 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

Display the OSPF configuration.

```
[edit]
lab@Vega# 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 {
    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@Vega# 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@Vega# **run show ospf interface**

Interface	State	Area	DR ID	BDR ID
Nbrs				
ge-0/0/4.136	DR	0.0.0.0	172.30.5.2	0.0.0.0
0				
ge-0/0/4.167	BDR	0.0.0.0	172.30.5.7	172.30.5.2
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	172.30.5.2
1				

Two neighbors in *Full* state, one in Area 0 and one in Area 4.

g. R7

Display the OSPF configuration.

```

[edit]
lab@Rigel# 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 {
        authentication {
            md5 1 key "$9$0eHXRrv2gJUi.dbYoGif50lRS1M"; ## SECRET-DATA
        }
    }
    interface lo0.0;
}
area 0.0.0.1 {
    interface ge-0/0/4.178 {
        authentication {
            md5 1 key "$9$Nab4ZzF/AulmfQnCuEhNdb2oG"; ## SECRET-DATA
        }
    }
}

```

Display the active OSPF neighbors.

```

[edit]
lab@Rigel# 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@Rigel# 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	172.30.5.2
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				

Two neighbors are in the *Full* state, one in Area 0 and one in Area 1.

h. R8

Display the OSPF configuration.

```
[edit]
lab@Procyon# 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
        }
    }
    interface lo0.0;
}
```

As with R1, R8 has interfaces in two non-backbone areas.

Display the active OSPF neighbors.

```
[edit]
lab@Procyon# 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@Procyon# 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
```

100.0 0	DR	0.0.0.2	172.30.5.8	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 have to add additional configurations to the protocol stanza and commit the changes to start tracking and record debug information in specified file.

a. R1 – R2 adjacency.

Let's configure traceoptions for OSPF on R1 for all errors with option *detail*.

```
[edit]
lab@Sun# set protocols ospf traceoptions file ospf.log

[edit]
lab@Sun# set protocols ospf traceoptions flag error detail

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

NOTE: You can enable one or multiple flags for the same tracefile. You can use the *detail* option to get more detailed information.

Now you can see the content of the file:

```
[edit]
lab@Sun# 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@Sun# run show ospf interface ae0.0
Interface      State   Area      DR ID      BDR ID
Nbrrs
ae0.0          DR      0.0.0.3    172.30.5.1  0.0.0.0
0
```

```
[edit]
lab@Sirius# run show ospf interface ae0.0
Interface      State   Area      DR ID      BDR ID
Nbrrs
ae0.0          DR      0.0.0.33   172.30.5.2  0.0.0.0
0
```

Fix the configuration on R2.

```
[edit]
lab@Sirius# rename protocols ospf area 33 to area 3
```

```
[edit]
lab@Sirius# 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@Sun# run show ospf interface ae0.0
Interface      State   Area      DR ID      BDR ID
Nbrrs
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@Sirius# run show ospf interface ae0.0
Interface      State   Area      DR ID      BDR ID
Nbrrs
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.

b. R2 – R3 adjacency.

Interface `ge-0/0/4.123` on R2 is connected to R3. From the previous Task you know that it reached the OSPF *ExStart* state only. This is the state where the database exchange starts. For this state the MTU must match on each side of the link.

Again you can use traceoptions on R2 to troubleshoot neighbor establishment.

```
[edit]
lab@Sirius# set protocols ospf traceoptions file ospf.log

[edit]
lab@Sirius# set protocols ospf traceoptions flag error detail

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

Display the content of the trace file.

```
[edit]
lab@Sirius# 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 a 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@Sirius# 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

[edit]
```

```
lab@Canopus# 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@Sirius# 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@Sirius# delete interfaces ge-0/0/4.123 family inet mtu

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

After equalizing the MTU on both ends, the OSPF neighbors reach the *Full* state.

```
[edit]
lab@Sirius# 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

c. R2 – R7 adjacency.

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@Sirius# 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@Sirius# run show ospf interface ge-0/0/4.127 extensive
Interface      State  Area      DR ID      BDR ID
Nbrrs
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@Rigel# run show ospf interface ge-0/0/4.127 detail
Interface      State  Area      DR ID      BDR ID
Nbrrs
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
```

Using again the `rename` feature you can change the interface IP address on R7 with just one command.

```
[edit]
lab@Rigel# rename interface ge-0/0/4.127 family inet address
172.30.1.18/30 to address 172.30.0.18/30
```

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

Fixing the interface address configuration puts the neighbors to *Full* state.

```
[edit]
lab@Sirius# run show ospf interface
```

Interface	State	Area	DR ID	BDR ID
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@Rigel# run show ospf interface
```

Interface	State	Area	DR ID	BDR ID
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				

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@Canopus# set protocols ospf traceoptions file ospf.log

[edit]
lab@Canopus# set protocols ospf traceoptions flag error detail

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

The interface connected to R4 reports a different type of Area configured on the other end.

```
[edit]
```

```
lab@Canopus# 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@Canopus# show protocols ospf area 4
nssa {
    area-range 172.30.32.0/20;
}
interface ge-0/0/4.134 {
    authentication {
        md5 1 key "$9$B6BE1Mg4ZDHmVw2aUH5TBIEyeW"; ## SECRET-
DATA
    }
}
```

```
[edit]
lab@Arcturus# show protocols ospf area 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;
```

Change the Area type on R4 to *NSSA*.

```
[edit]
lab@Arcturus# set protocols ospf area 4 nssa

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

This results in one more neighbor in state *Full* on R3.

```
[edit]
lab@Canopus# 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 neighborhood with R6 is not established.

```
[edit]
lab@Canopus# 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.

The first method is not very secured because the password is exchanged in clear text format in every message, but it is ofcourse 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@Vega# set protocols ospf area 0 interface ge-0/0/4.136
authentication md5 1 key "$9$L3KNs4f5F6CuHqPQnCB1LxNbYo"
```

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

Fixing the authentication issue resolves all neighbor adjacencies issues on R3 and R6.

```
[edit]
lab@Canopus# 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  172.30.5.2
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@Vega# 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  172.30.5.2
1
ge-0/0/4.167   BDR     0.0.0.0    172.30.5.7  172.30.5.2
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  172.30.5.2
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!

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 have to examine the content of the link state database on all devices.

a. R1.

```
[edit]
lab@Sun# 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
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@Sun# 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
```

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@Sirius# 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
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@Canopus# 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						
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 36	172.30.32.0	172.30.5.4	0x80000002	309	0x28	0x1df7
NSSA 36	172.30.32.0	172.30.5.5	0x80000004	839	0x28	0x21ef
NSSA 36	172.30.33.0	172.30.5.4	0x80000002	135	0x28	0x1202
NSSA 36	172.30.33.0	172.30.5.5	0x80000004	681	0x28	0x16f9
NSSA 36	172.30.34.0	172.30.5.4	0x80000001	554	0x28	0x90b
NSSA 36	172.30.34.0	172.30.5.5	0x80000004	523	0x28	0xb04
NSSA 36	172.30.35.0	172.30.5.4	0x80000001	554	0x28	0xfd15
NSSA 36	172.30.35.0	172.30.5.5	0x80000004	365	0x28	0xff0e
NSSA 36	172.30.36.0	172.30.5.4	0x80000001	554	0x28	0xf21f
NSSA 36	172.30.36.0	172.30.5.5	0x80000004	207	0x28	0xf418
NSSA 36	172.30.37.0	172.30.5.4	0x80000001	554	0x28	0xe729
NSSA 36	172.30.37.0	172.30.5.5	0x80000004	50	0x28	0xe922
NSSA 36	172.30.38.0	172.30.5.4	0x80000001	554	0x28	0xdc33
NSSA 36	172.30.38.0	172.30.5.5	0x80000003	2576	0x28	0xe02b
NSSA 36	172.30.39.0	172.30.5.4	0x80000001	554	0x28	0xd13d
NSSA 36	172.30.39.0	172.30.5.5	0x80000003	2418	0x28	0xd535
NSSA 36	172.30.40.0	172.30.5.4	0x80000003	553	0x28	0xc249
NSSA 36	172.30.40.0	172.30.5.5	0x80000003	2260	0x28	0xca3f
NSSA 36	172.30.41.0	172.30.5.4	0x80000003	553	0x28	0xb753
NSSA 36	172.30.41.0	172.30.5.5	0x80000003	2103	0x28	0xbf49
NSSA 36	172.30.42.0	172.30.5.4	0x80000003	553	0x28	0xac5d
NSSA 36	172.30.42.0	172.30.5.5	0x80000003	1945	0x28	0xb453
NSSA 36	172.30.43.0	172.30.5.4	0x80000003	553	0x28	0xa167
NSSA 36	172.30.43.0	172.30.5.5	0x80000003	1787	0x28	0xa95d
NSSA 36	172.30.44.0	172.30.5.4	0x80000003	553	0x28	0x9671
NSSA 36	172.30.44.0	172.30.5.5	0x80000003	1629	0x28	0x9e67
NSSA 36	172.30.45.0	172.30.5.4	0x80000003	553	0x28	0x8b7b
NSSA 36	172.30.45.0	172.30.5.5	0x80000003	1471	0x28	0x9371
NSSA 36	172.30.46.0	172.30.5.4	0x80000003	553	0x28	0x8085

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 that 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@Canopus# show protocols ospf area 4
nssa {
    area-range 172.30.32.0/20;
}
interface ge-0/0/4.134 {
    authentication {
        md5 1 key "$9$B6BE1Mg4ZDHmVw2aUH5TBIEyeW"; ## SECRET-
DATA
    }
}
```

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` stanza of the area configuration. This way you aggregate only the NSSA, external routes for the domain.

d. R4.

```
[edit]
lab@Arcturus# run show ospf database

OSPF database, Area 0.0.0.4
Type      ID          Adv Rtr      Seq      Age  Opt  Cksum
Len
```

Router 36	172.30.5.2	172.30.5.2	0x8000000c	424	0x20	0xf8b1
Router 36	172.30.5.3	172.30.5.3	0x8000000f	539	0x20	0xb30b
Router 60	*172.30.5.4	172.30.5.4	0x80000005	586	0x20	0xb051
Router 60	172.30.5.5	172.30.5.5	0x80000016	626	0x20	0xcd08
Network 32	172.30.0.21	172.30.5.3	0x80000003	621	0x20	0xb132
Network 32	172.30.0.30	172.30.5.5	0x80000003	619	0x20	0x5f77
Network 32	172.30.0.33	172.30.5.5	0x80000007	1069	0x20	0x1db4
Summary 28	0.0.0.0	172.30.5.2	0x80000006	1184	0x20	0xcf8e
Summary 28	172.30.0.12	172.30.5.3	0x8000000b	896	0x20	0xace0
Summary 28	172.30.0.16	172.30.5.3	0x80000001	2	0x20	0xa2ef
Summary 28	172.30.0.24	172.30.5.3	0x8000000a	423	0x20	0x364c
Summary 28	172.30.0.40	172.30.5.3	0x80000002	2	0x20	0xb9be
Summary 28	172.30.0.44	172.30.5.3	0x80000002	2	0x20	0x91e2
Summary 28	172.30.5.2	172.30.5.3	0x8000001c	11	0x20	0xc9b4
Summary 28	172.30.5.3	172.30.5.3	0x80000009	374	0x20	0xdbb5
Summary 28	172.30.5.7	172.30.5.3	0x80000002	2	0x20	0xd5bc
NSSA 36	*172.30.32.0	172.30.5.4	0x80000002	382	0x28	0x1df7
NSSA 36	172.30.32.0	172.30.5.5	0x80000004	911	0x28	0x21ef
NSSA 36	*172.30.33.0	172.30.5.4	0x80000002	207	0x28	0x1202
NSSA 36	172.30.33.0	172.30.5.5	0x80000004	753	0x28	0x16f9
NSSA 36	*172.30.34.0	172.30.5.4	0x80000002	32	0x28	0x70c
NSSA 36	172.30.34.0	172.30.5.5	0x80000004	595	0x28	0xb04
NSSA 36	*172.30.35.0	172.30.5.4	0x80000001	626	0x28	0xfd15
NSSA 36	172.30.35.0	172.30.5.5	0x80000004	438	0x28	0xff0e
NSSA 36	*172.30.36.0	172.30.5.4	0x80000001	626	0x28	0xf21f
NSSA 36	172.30.36.0	172.30.5.5	0x80000004	280	0x28	0xf418
NSSA 36	*172.30.37.0	172.30.5.4	0x80000001	626	0x28	0xe729
NSSA 36	172.30.37.0	172.30.5.5	0x80000004	123	0x28	0xe922
NSSA 36	*172.30.38.0	172.30.5.4	0x80000001	626	0x28	0xdc33
NSSA 36	172.30.38.0	172.30.5.5	0x80000003	2648	0x28	0xe02b

NSSA 36	*172.30.39.0	172.30.5.4	0x80000001	626	0x28	0xd13d
NSSA 36	172.30.39.0	172.30.5.5	0x80000003	2490	0x28	0xd535
NSSA 36	*172.30.40.0	172.30.5.4	0x80000003	625	0x28	0xc249
NSSA 36	172.30.40.0	172.30.5.5	0x80000003	2332	0x28	0xca3f
NSSA 36	*172.30.41.0	172.30.5.4	0x80000003	625	0x28	0xb753
NSSA 36	172.30.41.0	172.30.5.5	0x80000003	2175	0x28	0xbf49
NSSA 36	*172.30.42.0	172.30.5.4	0x80000003	625	0x28	0xac5d
NSSA 36	172.30.42.0	172.30.5.5	0x80000003	2017	0x28	0xb453
NSSA 36	*172.30.43.0	172.30.5.4	0x80000003	625	0x28	0xa167
NSSA 36	172.30.43.0	172.30.5.5	0x80000003	1859	0x28	0xa95d
NSSA 36	*172.30.44.0	172.30.5.4	0x80000003	625	0x28	0x9671
NSSA 36	172.30.44.0	172.30.5.5	0x80000003	1701	0x28	0x9e67
NSSA 36	*172.30.45.0	172.30.5.4	0x80000003	625	0x28	0x8b7b
NSSA 36	172.30.45.0	172.30.5.5	0x80000003	1543	0x28	0x9371
NSSA 36	*172.30.46.0	172.30.5.4	0x80000003	625	0x28	0x8085
NSSA 36	172.30.46.0	172.30.5.5	0x80000003	1385	0x28	0x887b
NSSA 36	*172.30.47.0	172.30.5.4	0x80000003	625	0x28	0x758f
NSSA 36	172.30.47.0	172.30.5.5	0x80000003	1227	0x28	0x7d85

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@A-Centauri# 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						
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 36	*172.30.41.0	172.30.5.5	0x80000003	2432	0x28	0xbf49
NSSA 36	172.30.42.0	172.30.5.4	0x80000003	883	0x28	0xac5d
NSSA 36	*172.30.42.0	172.30.5.5	0x80000003	2274	0x28	0xb453
NSSA 36	172.30.43.0	172.30.5.4	0x80000003	883	0x28	0xa167
NSSA 36	*172.30.43.0	172.30.5.5	0x80000003	2116	0x28	0xa95d
NSSA 36	172.30.44.0	172.30.5.4	0x80000003	883	0x28	0x9671
NSSA 36	*172.30.44.0	172.30.5.5	0x80000003	1958	0x28	0x9e67
NSSA 36	172.30.45.0	172.30.5.4	0x80000003	883	0x28	0x8b7b
NSSA 36	*172.30.45.0	172.30.5.5	0x80000003	1800	0x28	0x9371
NSSA 36	172.30.46.0	172.30.5.4	0x80000003	883	0x28	0x8085
NSSA 36	*172.30.46.0	172.30.5.5	0x80000003	1642	0x28	0x887b
NSSA 36	172.30.47.0	172.30.5.4	0x80000003	883	0x28	0x758f
NSSA 36	*172.30.47.0	172.30.5.5	0x80000003	1484	0x28	0x7d85

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@Vega# 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
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 28	172.30.0.28	172.30.5.3	0x80000001	920	0x22	0xc78	
Summary 28	*172.30.0.32	172.30.5.2	0x80000bd3	3	0x22	0x198b	
Summary 28	172.30.0.32	172.30.5.3	0x80000001	920	0x22	0xed91	
Summary 28	172.30.0.44	172.30.5.7	0x8000000a	911	0x22	0x3731	
Summary 28	*172.30.5.4	172.30.5.2	0x80000219	1	0x22	0xa7d3	
Summary 28	172.30.5.4	172.30.5.3	0x80000001	920	0x22	0xcdc7	
Summary 28	*172.30.5.5	172.30.5.2	0x80000bd2	3	0x22	0x5b3	
Summary 28	172.30.5.5	172.30.5.3	0x80000001	920	0x22	0xcdc5	
OSPF database, Area 0.0.0.4							
Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	
Len							
Router 36	*172.30.5.2	172.30.5.2	0x8000000c	719	0x20	0xf8b1	
Router 36	172.30.5.3	172.30.5.3	0x8000000f	838	0x20	0xb30b	
Router 60	172.30.5.4	172.30.5.4	0x80000005	886	0x20	0xb051	
Router 60	172.30.5.5	172.30.5.5	0x80000016	924	0x20	0xcd08	
Network 32	172.30.0.21	172.30.5.3	0x80000004	34	0x20	0xaf33	
Network 32	172.30.0.30	172.30.5.5	0x80000003	916	0x20	0x5f77	
Network 32	172.30.0.33	172.30.5.5	0x80000007	1366	0x20	0x1db4	
Summary 28	*0.0.0.0	172.30.5.2	0x80000007	39	0x20	0xcd8f	
Summary 28	172.30.0.12	172.30.5.3	0x8000000c	200	0x20	0xaae1	
Summary 28	172.30.0.16	172.30.5.3	0x80000001	3600	0x20	0xa2ef	
Summary 28	172.30.0.24	172.30.5.3	0x8000000a	722	0x20	0x364c	
Summary 28	172.30.0.40	172.30.5.3	0x80000005	3600	0x20	0xa9cc	
Summary 28	172.30.0.44	172.30.5.3	0x80000005	3600	0x20	0x8be5	
Summary 28	172.30.5.2	172.30.5.3	0x8000002d	5	0x20	0xa7c5	
Summary 28	172.30.5.3	172.30.5.3	0x80000009	673	0x20	0xdbb5	
Summary 28	172.30.5.6	172.30.5.3	0x80000004	3600	0x20	0xd1c0	
Summary 28	172.30.5.7	172.30.5.3	0x80000005	3600	0x20	0xcfbf	
NSSA 36	172.30.32.0	172.30.5.4	0x80000002	682	0x28	0x1df7	
NSSA 36	172.30.32.0	172.30.5.5	0x80000004	1208	0x28	0x21ef	
NSSA 36	172.30.33.0	172.30.5.4	0x80000002	506	0x28	0x1202	

NSSA 36	172.30.33.0	172.30.5.5	0x80000004	1051	0x28	0x16f9	
NSSA 36	172.30.34.0	172.30.5.4	0x80000002	331	0x28	0x70c	
NSSA 36	172.30.34.0	172.30.5.5	0x80000004	893	0x28	0xb04	
NSSA 36	172.30.35.0	172.30.5.4	0x80000002	157	0x28	0xfb16	
NSSA 36	172.30.35.0	172.30.5.5	0x80000004	735	0x28	0xff0e	
NSSA 36	172.30.36.0	172.30.5.4	0x80000001	926	0x28	0xf21f	
NSSA 36	172.30.36.0	172.30.5.5	0x80000004	577	0x28	0xf418	
NSSA 36	172.30.37.0	172.30.5.4	0x80000001	926	0x28	0xe729	
NSSA 36	172.30.37.0	172.30.5.5	0x80000004	420	0x28	0xe922	
NSSA 36	172.30.38.0	172.30.5.4	0x80000001	926	0x28	0xdc33	
NSSA 36	172.30.38.0	172.30.5.5	0x80000004	261	0x28	0xde2c	
NSSA 36	172.30.39.0	172.30.5.4	0x80000001	926	0x28	0xd13d	
NSSA 36	172.30.39.0	172.30.5.5	0x80000004	103	0x28	0xd336	
NSSA 36	172.30.40.0	172.30.5.4	0x80000003	924	0x28	0xc249	
NSSA 36	172.30.40.0	172.30.5.5	0x80000003	2630	0x28	0xca3f	
NSSA 36	172.30.41.0	172.30.5.4	0x80000003	924	0x28	0xb753	
NSSA 36	172.30.41.0	172.30.5.5	0x80000003	2472	0x28	0xbf49	
NSSA 36	172.30.42.0	172.30.5.4	0x80000003	924	0x28	0xac5d	
NSSA 36	172.30.42.0	172.30.5.5	0x80000003	2314	0x28	0xb453	
NSSA 36	172.30.43.0	172.30.5.4	0x80000003	924	0x28	0xa167	
NSSA 36	172.30.43.0	172.30.5.5	0x80000003	2156	0x28	0xa95d	
NSSA 36	172.30.44.0	172.30.5.4	0x80000003	924	0x28	0x9671	
NSSA 36	172.30.44.0	172.30.5.5	0x80000003	1998	0x28	0x9e67	
NSSA 36	172.30.45.0	172.30.5.4	0x80000003	924	0x28	0x8b7b	
NSSA 36	172.30.45.0	172.30.5.5	0x80000003	1840	0x28	0x9371	
NSSA 36	172.30.46.0	172.30.5.4	0x80000003	924	0x28	0x8085	
NSSA 36	172.30.46.0	172.30.5.5	0x80000003	1682	0x28	0x887b	
NSSA 36	172.30.47.0	172.30.5.4	0x80000003	924	0x28	0x758f	
NSSA 36	172.30.47.0	172.30.5.5	0x80000003	1524	0x28	0x7d85	
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 a 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@Canopus# run 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@Canopus# run show routing-options
router-id 172.30.5.3;
```

NOTE: The Router-id is configured under the **routing-options** stanza and is used for whole router.

Check the same command outputs on R6.

```
[edit]
lab@Vega# run 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@Vega# run show routing-options
router-id 172.30.5.2;
```

The Router-id is different then 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@Sirius# run show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-re;
        }
        address 172.30.5.2/32;
    }
    family inet6 {
        address fd17:f0f4:f691:5::2/128;
    }
}
```

```
[edit]
lab@Sirius# run show routing-options
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@Vega# set routing-options router-id 172.30.5.6
```

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

Fixing this error changes the OSPF database significantly.

```
lab@Vega# run show ospf database

    OSPF database, Area 0.0.0.0
Type      ID              Adv Rtr          Seq           Age    Opt    Cksum
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 28	172.30.0.20	172.30.5.3	0x8000000a	883	0x22	0x4044
Summary 28	*172.30.0.20	172.30.5.6	0x80000001	80	0x22	0x5434
Summary 28	172.30.0.28	172.30.5.3	0x80000002	234	0x22	0xa79
Summary 28	*172.30.0.28	172.30.5.6	0x80000001	80	0x22	0xf987
Summary 28	172.30.0.32	172.30.5.3	0x80000002	68	0x22	0xeb92
Summary 28	*172.30.0.32	172.30.5.6	0x80000002	80	0x22	0xc5b7
Summary 28	172.30.0.44	172.30.5.7	0x8000000b	44	0x22	0x3532
Summary 28	172.30.5.4	172.30.5.3	0x80000001	1287	0x22	0xcdc7
Summary 28	*172.30.5.4	172.30.5.6	0x80000001	80	0x22	0xc5cb
Summary 28	172.30.5.5	172.30.5.3	0x80000001	1287	0x22	0xcdc5
Summary 28	*172.30.5.5	172.30.5.6	0x80000001	80	0x22	0xb1df
OSPF database, Area 0.0.0.4						
Type	ID	Adv Rtr	Seq	Age	Opt	Cksum
Len						
Router 36	172.30.5.2	172.30.5.2	0x8000000c	1088	0x20	0xf8b1
Router 36	172.30.5.3	172.30.5.3	0x8000000f	1205	0x20	0xb30b
Router 60	172.30.5.4	172.30.5.4	0x80000005	1252	0x20	0xb051
Router 60	172.30.5.5	172.30.5.5	0x80000018	86	0x20	0xc90a
Router 36	*172.30.5.6	172.30.5.6	0x80000004	80	0x20	0xc8e1
Network 32	172.30.0.21	172.30.5.3	0x80000004	400	0x20	0xaf33
Network 32	172.30.0.30	172.30.5.5	0x80000003	1283	0x20	0x5f77
Network 32	172.30.0.33	172.30.5.5	0x80000009	86	0x20	0x517a
Summary 28	0.0.0.0	172.30.5.2	0x80000007	407	0x20	0xcd8f
Summary 28	*0.0.0.0	172.30.5.6	0x80000001	85	0x20	0xc19d
Summary 28	172.30.0.0	172.30.5.3	0x80000001	88	0x20	0x435f
Summary 28	172.30.0.12	172.30.5.3	0x8000000c	567	0x20	0xaae1
Summary 28	172.30.0.16	172.30.5.3	0x80000001	83	0x20	0xa2ef
Summary 28	172.30.0.24	172.30.5.3	0x8000000c	87	0x20	0x324e
Summary 28	172.30.0.40	172.30.5.3	0x80000002	83	0x20	0xafc9
Summary 28	172.30.0.44	172.30.5.3	0x80000002	83	0x20	0x91e2
Summary 28	172.30.5.2	172.30.5.3	0x80000036	95	0x20	0x95ce

Summary 28	172.30.5.3	172.30.5.3	0x80000009	1040	0x20	0xdbb5
Summary 28	172.30.5.6	172.30.5.3	0x80000002	83	0x20	0xd5be
Summary 28	172.30.5.7	172.30.5.3	0x80000002	83	0x20	0xd5bc
NSSA 36	172.30.32.0	172.30.5.4	0x80000002	1048	0x28	0x1df7
NSSA 36	172.30.32.0	172.30.5.5	0x80000004	1575	0x28	0x21ef
NSSA 36	172.30.33.0	172.30.5.4	0x80000002	873	0x28	0x1202
NSSA 36	172.30.33.0	172.30.5.5	0x80000004	1417	0x28	0x16f9
NSSA 36	172.30.34.0	172.30.5.4	0x80000002	698	0x28	0x70c
NSSA 36	172.30.34.0	172.30.5.5	0x80000004	1259	0x28	0xb04
NSSA 36	172.30.35.0	172.30.5.4	0x80000002	523	0x28	0xfb16
NSSA 36	172.30.35.0	172.30.5.5	0x80000004	1101	0x28	0xff0e
NSSA 36	172.30.36.0	172.30.5.4	0x80000002	349	0x28	0xf020
NSSA 36	172.30.36.0	172.30.5.5	0x80000004	943	0x28	0xf418
NSSA 36	172.30.37.0	172.30.5.4	0x80000002	174	0x28	0xe52a
NSSA 36	172.30.37.0	172.30.5.5	0x80000004	786	0x28	0xe922
NSSA 36	172.30.38.0	172.30.5.4	0x80000001	1292	0x28	0xdc33
NSSA 36	172.30.38.0	172.30.5.5	0x80000004	629	0x28	0xde2c
NSSA 36	172.30.39.0	172.30.5.4	0x80000001	1293	0x28	0xd13d
NSSA 36	172.30.39.0	172.30.5.5	0x80000004	471	0x28	0xd336
NSSA 36	172.30.40.0	172.30.5.4	0x80000003	1291	0x28	0xc249
NSSA 36	172.30.40.0	172.30.5.5	0x80000004	313	0x28	0xc840
NSSA 36	172.30.41.0	172.30.5.4	0x80000003	1291	0x28	0xb753
NSSA 36	172.30.41.0	172.30.5.5	0x80000004	155	0x28	0xbd4a
NSSA 36	172.30.42.0	172.30.5.4	0x80000003	1291	0x28	0xac5d
NSSA 36	172.30.42.0	172.30.5.5	0x80000003	2682	0x28	0xb453
NSSA 36	172.30.43.0	172.30.5.4	0x80000003	1291	0x28	0xa167
NSSA 36	172.30.43.0	172.30.5.5	0x80000003	2524	0x28	0xa95d
NSSA 36	172.30.44.0	172.30.5.4	0x80000003	1291	0x28	0x9671
NSSA 36	172.30.44.0	172.30.5.5	0x80000003	2366	0x28	0x9e67
NSSA 36	172.30.45.0	172.30.5.4	0x80000003	1291	0x28	0x8b7b

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.

```
[edit]
lab@Arcturus# 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 36	172.30.37.0	172.30.5.5	0x80000004	123	0x28	0xe922
NSSA 36	*172.30.38.0	172.30.5.4	0x80000001	626	0x28	0xdc33
NSSA 36	172.30.38.0	172.30.5.5	0x80000003	2648	0x28	0xe02b
NSSA 36	*172.30.39.0	172.30.5.4	0x80000001	626	0x28	0xd13d
NSSA 36	172.30.39.0	172.30.5.5	0x80000003	2490	0x28	0xd535
NSSA 36	*172.30.40.0	172.30.5.4	0x80000003	625	0x28	0xc249
NSSA 36	172.30.40.0	172.30.5.5	0x80000003	2332	0x28	0xca3f
NSSA 36	*172.30.41.0	172.30.5.4	0x80000003	625	0x28	0xb753
NSSA 36	172.30.41.0	172.30.5.5	0x80000003	2175	0x28	0xbf49
NSSA 36	*172.30.42.0	172.30.5.4	0x80000003	625	0x28	0xac5d
NSSA 36	172.30.42.0	172.30.5.5	0x80000003	2017	0x28	0xb453
NSSA 36	*172.30.43.0	172.30.5.4	0x80000003	625	0x28	0xa167
NSSA 36	172.30.43.0	172.30.5.5	0x80000003	1859	0x28	0xa95d
NSSA 36	*172.30.44.0	172.30.5.4	0x80000003	625	0x28	0x9671
NSSA 36	172.30.44.0	172.30.5.5	0x80000003	1701	0x28	0x9e67
NSSA 36	*172.30.45.0	172.30.5.4	0x80000003	625	0x28	0x8b7b
NSSA 36	172.30.45.0	172.30.5.5	0x80000003	1543	0x28	0x9371
NSSA 36	*172.30.46.0	172.30.5.4	0x80000003	625	0x28	0x8085
NSSA 36	172.30.46.0	172.30.5.5	0x80000003	1385	0x28	0x887b
NSSA 36	*172.30.47.0	172.30.5.4	0x80000003	625	0x28	0x758f
NSSA 36	172.30.47.0	172.30.5.5	0x80000003	1227	0x28	0x7d85

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.

```
[edit]
lab@Canopus# set protocols ospf area 4 interface ge-0/0/4.134
interface-type p2p
```

```
[edit]
```

```
lab@Canopus# commit
commit complete
```

```
[edit]
lab@Arcturus# set protocols ospf area 4 interface ge-0/0/4.134
interface-type p2p
```

```
[edit]
lab@Arcturus# set protocols ospf area 4 interface ge-0/0/4.145
interface-type p2p
```

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

```
[edit]
lab@A-Centauri# set protocols ospf area 4 interface ge-0/0/4.145
interface-type p2p
```

```
[edit]
lab@A-Centauri# set protocols ospf area 4 interface ae0.0
interface-type p2p
```

```
[edit]
lab@A-Centauri# commit
commit complete
```

```
[edit]
lab@Vega# set protocols ospf area 4 interface ae0.0 interface-type
p2p
```

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

Check the results of the interface type change on R5.

```
[edit]
lab@A-Centauri# 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@Arcturus# run show ospf database area 4
```

OSPF database, Area 0.0.0.4							
Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len

Router 36	172.30.5.2	172.30.5.2	0x8000000c	2508	0x20	0xf8b1
Router 48	172.30.5.3	172.30.5.3	0x80000012	61	0x20	0x4596
Router 84	*172.30.5.4	172.30.5.4	0x8000000a	60	0x20	0x23c
Router 84	172.30.5.5	172.30.5.5	0x8000001f	254	0x20	0xc744
Router 48	172.30.5.6	172.30.5.6	0x8000000a	301	0x20	0xfeed4
Summary 28	0.0.0.0	172.30.5.2	0x80000007	1828	0x20	0xcd8f
NSSA 36	0.0.0.0	172.30.5.3	0x80000001	61	0x20	0xabaa
NSSA 36	0.0.0.0	172.30.5.5	0x80000002	246	0x28	0xbe3e
NSSA 36	0.0.0.0	172.30.5.6	0x80000001	648	0x20	0x99b9
NSSA 36	*172.30.32.0	172.30.5.4	0x80000002	2466	0x28	0x1df7
NSSA 36	172.30.32.0	172.30.5.5	0x80000005	469	0x28	0x1ff0
NSSA 36	*172.30.33.0	172.30.5.4	0x80000002	2291	0x28	0x1202
NSSA 36	172.30.33.0	172.30.5.5	0x80000005	311	0x28	0x14fa
NSSA 36	*172.30.34.0	172.30.5.4	0x80000002	2116	0x28	0x70c
NSSA 36	172.30.34.0	172.30.5.5	0x80000005	41	0x28	0x905
NSSA 36	*172.30.35.0	172.30.5.4	0x80000002	1942	0x28	0xfb16
NSSA 36	172.30.35.0	172.30.5.5	0x80000004	2522	0x28	0xff0e
NSSA 36	*172.30.36.0	172.30.5.4	0x80000002	1767	0x28	0xf020
NSSA 36	172.30.36.0	172.30.5.5	0x80000004	2364	0x28	0xf418
NSSA 36	*172.30.37.0	172.30.5.4	0x80000002	1592	0x28	0xe52a
NSSA 36	172.30.37.0	172.30.5.5	0x80000004	2207	0x28	0xe922
NSSA 36	*172.30.38.0	172.30.5.4	0x80000002	1417	0x28	0xda34
NSSA 36	172.30.38.0	172.30.5.5	0x80000004	2048	0x28	0xde2c
NSSA 36	*172.30.39.0	172.30.5.4	0x80000002	1242	0x28	0xcf3e
NSSA 36	172.30.39.0	172.30.5.5	0x80000004	1890	0x28	0xd336
NSSA 36	*172.30.40.0	172.30.5.4	0x80000004	1068	0x28	0xc04a
NSSA 36	172.30.40.0	172.30.5.5	0x80000004	1732	0x28	0xc840
NSSA 36	*172.30.41.0	172.30.5.4	0x80000004	893	0x28	0xb554
NSSA 36	172.30.41.0	172.30.5.5	0x80000004	1574	0x28	0xbd4a
NSSA 36	*172.30.42.0	172.30.5.4	0x80000004	718	0x28	0xaa5e

NSSA 36	172.30.42.0	172.30.5.5	0x80000004	1416	0x28	0xb254
NSSA 36	*172.30.43.0	172.30.5.4	0x80000004	543	0x28	0x9f68
NSSA 36	172.30.43.0	172.30.5.5	0x80000004	1258	0x28	0xa75e
NSSA 36	*172.30.44.0	172.30.5.4	0x80000004	369	0x28	0x9472
NSSA 36	172.30.44.0	172.30.5.5	0x80000004	1100	0x28	0x9c68
NSSA 36	*172.30.45.0	172.30.5.4	0x80000004	194	0x28	0x897c
NSSA 36	172.30.45.0	172.30.5.5	0x80000004	943	0x28	0x9172
NSSA 36	*172.30.46.0	172.30.5.4	0x80000004	19	0x28	0x7e86
NSSA 36	172.30.46.0	172.30.5.5	0x80000004	785	0x28	0x867c
NSSA 36	*172.30.47.0	172.30.5.4	0x80000003	2709	0x28	0x758f
NSSA 36	172.30.47.0	172.30.5.5	0x80000004	627	0x28	0x7b86

```
[edit]
lab@Arcturus# run show ospf database area 4 summary
Area 0.0.0.4:
    5 Router LSAs
    1 Summary 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:
```

There are no LSAs Type 2 present anymore in the link state database.

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` stanza 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@Canopus# set protocols ospf area 4 nssa default-lsa default-
metric 10 type-7
```

```
[edit]
lab@Canopus# set protocols ospf area 4 nssa no-summaries
```

```
[edit]
lab@Canopus# 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@Vega# run show protocols area 4 nssa
default-lsa default-metric 10;
no-summaries;
```

```
[edit]
lab@Vega# set protocols ospf area 4 nssa default-lsa type-7
```

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

Checking the content of the database in NSSA 4.

```
[edit]
lab@Arcturus# 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
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
```

```

NSSA      172.30.36.0      172.30.5.5      0x80000005      184  0x28 0xf219
36
NSSA      *172.30.37.0      172.30.5.4      0x80000004      55  0x28 0xe12c
36
NSSA      172.30.37.0      172.30.5.5      0x80000005      16  0x28 0xe723
36
NSSA      *172.30.38.0      172.30.5.4      0x80000003      1308 0x28 0xd835
36
NSSA      172.30.38.0      172.30.5.5      0x80000004      1293 0x28 0xde2c
36

[edit]
lab@Arcturus# run show ospf database summary
Area 0.0.0.4:
    4 Router LSAs
    34 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:

```

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@Arcturus# run show protocols rip
group rip {
    export default-to-rip;
    neighbor ge-0/0/4.202;
}

[edit]
lab@A-Centauri# run 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@Arcturus# run show rip neighbor
```

In Neighbor Met	Local State	Source Address	Destination Address	Send Mode	Receive Mode
---	----	-----	-----	----	-----
ge-0/0/4.202	Up	172.30.0.49	224.0.0.9	mcast	both

```
1
```

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@Arcturus# 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: Since 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@Arcturus# run show route 0/0 exact
```

inet.0: 39 destinations, 55 routes (39 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```
0.0.0.0/0          *[OSPF/150] 00:09:10, 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@Arcturus# run show route receive-protocol rip 172.30.0.50

inet.0: 36 destinations, 36 routes (36 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
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@Arcturus# run show route protocol rip
```

```

inet.0: 39 destinations, 55 routes (39 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
                   > 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@Arcturus# run show ospf database

    OSPF database, Area 0.0.0.4

```

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum
Router Len	172.30.5.3	172.30.5.3	0x8000000a	490	0x20	0x558e
Router 48	*172.30.5.4	172.30.5.4	0x8000000e	489	0x20	0xf940
Router 84	172.30.5.5	172.30.5.5	0x80000014	1161	0x20	0xdd39
Router 84	172.30.5.6	172.30.5.6	0x80000009	150	0x20	0xf0d3
NSSA 48	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						
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@A-Centauri# run show rip neighbor

In      Local   Source      Destination   Send   Receive
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@A-Centauri# run show route advertising-protocol rip
172.30.0.57
```

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@A-Centauri# 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 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@A-Centauri# run show route receive-protocol rip 172.30.0.58

inet.0: 33 destinations, 33 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.

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 a 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 behaviour 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;
      }
    }
  }
}
```

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 *accept* or *reject*. 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. *Accept*, *reject* and *next-policy* are available terminating actions.

The *To* statement can be used as a match condition too. It is used to match routes which are sent 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.

```
[edit]
lab@Arcturus# set policy-options policy-statement rip-filter term
1 from protocol rip

[edit]
lab@Arcturus# set policy-option policy-statement rip-filter term 1
from route-filter 0/0 exact

[edit]
lab@Arcturus# set policy-otions policy-statement rip-filter term 1
then reject

[edit]
lab@Arcturus# set protocols rip group rip import rip-filter

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

```
[edit]
lab@A-Centauri# set policy-option policy-statement rip-filter term
1 from protocol rip
```

```
[edit]
lab@A-Centauri# set policy-option policy-statement rip-filter term
1 from route-filter 0/0 exact

[edit]
lab@A-Centauri# set policy-option policy-statement rip-filter term
1 then reject

[edit]
lab@A-Centauri# set protocols rip group rip import rip-filter

[edit]
lab@A-Centauri# 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.

```
[edit]
lab@Arcturus# run show route 0/0 exact

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

[edit]
lab@A-Centauri# run show route 0/0 exact

inet.0: 35 destinations, 51 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@Arcturus# run show route protocol rip

inet.0: 36 destinations, 36 routes (36 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
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@A-Centauri# run show route protocol rip

inet.0: 33 destinations, 33 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
> 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) Advertising 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@Arcturus# show protocols rip
group rip {
  export default-to-rip;
  import rip-filter;
  neighbor ge-0/0/4.202;
}

[edit]
lab@A-Centauri# 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@Arcturus# 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@A-Centauri# show policy-options policy-statement default-to-rip
term 1 {

```

```

    from {
        protocol ospf;
        route-filter 0.0.0.0/0 exact;
    }
    then {
        metric 5;
        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@Arcturus# run 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@A-Centauri# run 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;
    }
}

```

```

        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@Arcturus# run 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@Vega# run show ospf database external

```

Type	OSPF AS SCOPE ID	link state database Adv Rtr	Seq	Age	Opt	Cksum
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	0xfc1036
Extern	*172.30.47.0	172.30.5.6	0x80000001	116	0x22	0xf11a36

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 chooses the ABR with the highest Router ID. So R6 should be the ABR used for this purpose.

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@Vega# 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** stanza. Configured under the Area directly this command will summarize the only LSAs Type 5 but not LSAs Type 7, but applied under the **nssa** stanza only the NSSA Type 7 LSAs are summarized.

Fix the mistake.

First to remove the unneeded command.

```
[edit]
lab@Vega# delete protocols ospf area 4 area-range 172.30.32.0/20

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


Configure the **area-range** command under the correct **nssa** stanza.

```
[edit]
lab@Vega# set protocols ospf area 4 nssa area-range 172.30.32.0/20

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

```
[edit]
lab@Vega# run show ospf database external
```

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum
Extern	*172.30.32.0	172.30.5.6	0x80000003	14	0x22	0xf506

After the change, the summary route is originated by R6 and more specific routes are suppressed.

11) Verify Loopback reachability.

a. R1.

```
[edit]
lab@Sun# run show route 172.30.5/24 terse
```

inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A	Destination	P	Prf	Metric 1	Metric 2	Next hop
	AS path					
*	172.30.5.1/32	D	0			>100.0
*	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.

```
[edit]
lab@Sirius# run show route 172.30.5/24 terse
```

inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination	P	Prf	Metric 1	Metric 2	Next hop
AS path					
* 172.30.5.2/32	D	0			>100.0
* 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@Canopus# run show route 172.30.5/24 terse

inet.0: 43 destinations, 43 routes (43 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

A 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           >100.0
* 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.

```
[edit]
lab@Rigel# run show route 172.30.5/24 terse

inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

A 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	>100.0

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@Procyon# run show route 172.30.5/24 terse

inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A 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             >100.0
```

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 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@Sirius# 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@Sirius# 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@Rigel# 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@Rigel# 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@Sun# run show 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;
    }
  }
}

```

Although is not a recommended design to have discontinuous Area's, 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@Sun# edit protocols ospf area 0

[edit protocols ospf area 0.0.0.0]

```

```
lab@Sun# set virtual-link neighbor-id 172.30.5.2 transit-area 3

[edit protocols ospf area 0.0.0.0]
lab@Sun# 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@Sirius# edit protocols ospf

[edit protocols ospf]
lab@Sirius# set area 0 virtual-link neighbor-id 172.30.5.1
transit-area 3

[edit protocols ospf]
lab@Sirius# commit
commit complete
```

c. R8.

```
[edit]
lab@Procyon# edit protocols ospf area 0

[edit protocols ospf area 0.0.0.0]
lab@Procyon# set virtual-link neighbor-id 172.30.5.7 transit-area
1

[edit protocols ospf area 0.0.0.0]
lab@Procyon# commit
commit complete
```

d. R7.

```
[edit]
lab@Rigel# edit protocols ospf area 0

[edit protocols ospf area 0.0.0.0]
lab@Rigel# set virtual-link neighbor-id 172.30.5.8 transit-area 1

[edit protocols ospf area 0.0.0.0]
lab@Rigel# commit
commit complete
```

Verify the result of the virtual link configuration on R1.

```
[edit]
lab@Sun# show 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;
}
area 0.0.0.0 {
    virtual-link neighbor-id 172.30.5.2 transit-area 0.0.0.3;
}

[edit]
lab@Sun# 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@Sirius# show protocols ospf
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@Sirius# 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.

Verify the results of the virtual link configuration on R8.

```

[edit]
lab@Procyon# 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@Rigel# run show route 172.30.5/24 terse

inet.0: 26 destinations, 26 routes (26 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A 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 protocols ospf area 0.0.0.0]
lab@Rigel# 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

    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 have to 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@Procyon# 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.

```
[edit]
lab@Vega# deactivate protocols ospf

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

```
[edit protocols ospf area 0.0.0.0]
lab@Procyon# 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!

Don't forget to activate OSPF on R6.

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. ISIS 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 ISIS 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.
- 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 are 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 out dated 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 to 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 ISIS metric of 10. You can configure a manual metric on each ISIS interface.

- 1) Let's start with loading the configuration on all eight devices from the files provided with the JNCIE workbook for Chapter Two.

[edit]

```
lab@Sun# load override terminal
```

- 2) Verify IS-IS adjacencies according to Figure 5. Pay attention that only one adjacency is configured on each IS-IS interface.

a. R1.

```
[edit]
lab@Sun# show interface ae0.0
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@Sun# show interfaces lo0.0
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
 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 have to configure *family iso* together with adding it under the *protocol isis* stanza.

```

[edit]
lab@Sun# show protocols isis
interface ge-0/0/4.118 {

```

```

    level 2 disable;
    level 1 {
        hello-authentication-key "$9$bB2Zj6/t01hTznCulyrbs2aGi";
## 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 have to 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@Sun# run show isis adjacency
Interface          System          L   State          Hold (secs)  SNPA
ae0.0              1720.3000.5002  2   Initializing    22
ge-0/0/4.118      Procyon        1   Up              8
f8:c0:1:dc:2e:84

[edit]
lab@Sun# 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 Procyon.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@Sirius# show interface lo0.0
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@Sirius# 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;

[edit]
lab@Sirius# 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@Sirius# 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@Canopus# show interface lo0.0
unit 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@Canopus# show protocols isis
export [ inject-externals ll-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@Canopus# 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 Canopus.00        Disabled
10/10
ge-0/0/4.136       1    0x1 Canopus.00        Disabled
10/10
lo0.0              0    0x1 Passive           Passive
0/0

[edit]
lab@Canopus# run show isis adjacency
Interface          System          L State          Hold (secs)
SNPA
ge-0/0/4.123       Sirius          2  Up              20

```

R3 has just one of the three neighbor adjacencies established.

d. R4.

```

[edit]
lab@Arcturus# 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@Arcturus# 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@Arcturus# run show isis adjacency
[edit]
lab@Arcturus# 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 Arcturus.00    Disabled
10/10
ge-0/0/4.145       1    0x1 Arcturus.00    Disabled
10/10
lo0.0               0    0x1 Passive       Passive
0/0

```

R4 has no established adjacencies.

e. R5.

```

[edit]
lab@A-Centauri# show interface lo0.0
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@A-Centauri# 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;
    }
}
interface lo0.0;

[edit]
lab@A-Centauri# run show isis adjacency
Interface           System           L State           Hold (secs)
SNPA
ae0.0                Vega                1  Up                20
f8:c0:1:dc:2c:80

[edit]
lab@A-Centauri# run show isis interface
IS-IS interface database:
Interface           L CirID Level 1 DR           Level 2 DR           L1/L2
Metric
ae0.0                1    0x2 A-Centauri.02         Disabled
10/10
ge-0/0/4.14          1    0x1 A-Centauri.00         Disabled
10/10
lo0.0                0    0x1 Passive              Passive
0/0

```

R5 has one adjacency established with R6.

f. R6.

```

[edit]
lab@Vega# show interface lo0.0
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@Vega# 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;
    }
}
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@Vega# run show isis adjacency
Interface      System      L State      Hold (secs) SNPA
ae0.0          A-Centauri    1 Up          7
f8:c0:1:dd:4:0

[edit]
lab@Vega# run show isis interface
IS-IS interface database:
Interface      L CirID Level 1 DR      Level 2 DR      L1/L2
Metric
ae0.0          1    0x1 A-Centauri.02    Disabled
10/10
ge-0/0/4.136   1    0x1 Vega.00          Disabled
10/10
ge-0/0/4.167   2    0x1 Disabled         Vega.00
10/10
lo0.0          0    0x1 Passive          Passive
0/0

```

R6 has only one adjacency established to R5.

g. R7.

```
[edit]
lab@Rigel# show interface lo0.0
unit 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;
    }
}
```

```
[edit]
lab@Rigel# 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@Rigel# run show isis adjacency
```

Interface	System	L State	Hold (secs)
SNPA			
ge-0/0/4.127	Sirius	2 Up	20

```
[edit]
```

```
lab@Rigel# 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
```

R7s adjacency with R2 is established. The other two interfaces configured for IS-IS have no adjacencies established.

h. R8.

```
[edit]
lab@Procyon# 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@Procyon# 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;
```

```
[edit]
lab@Procyon# run show isis adjacency
Interface          System    L State      Hold (secs) SNPA
ge-0/0/4.118       Sun      1 Up         18
0:26:88:ef:75:84
```

```
[edit]
lab@Procyon# 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 Procyon.02      Disabled
10/10
lo0.0              0      0x1 Passive        Passive
0/0
```

R8 has one established adjacency with R1.

NOTE: `run show isis interface` 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@Sun# run show interfaces ae0.0
Logical interface ae0.0 (Index 78) (SNMP ifIndex 585)
Description: R2 connection
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Statistics          Packets          pps          Bytes
bps
Bundle:
  Input :           9           0           812
0
  Output:          86           0          3686
0
Security: Zone: Null
Protocol inet, MTU: 1386
Flags: Sendbroadcast-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
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::226:88ff:feef:7600

```

```

[edit]
lab@Sirius# run show interfaces ae0.0
  Logical interface ae0.0 (Index 69) (SNMP ifIndex 557)
  Description: R1 connection
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Statistics          Packets          pps          Bytes
bps
  Bundle:
    Input :           2539             0         596504
0
    Output:           3115             0         708597
0
  Security: Zone: Null
  Protocol inet, MTU: 1500
    Flags: Sendbcst-pkt-to-re, Is-Primary
    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
  Flags: Is-Primary
  Addresses, Flags: Is-Preferred
  Destination: fe80::/64, Local: fe80::fac0:1ff:fedd:280

```

R1 is not accepting the padded IIHs because they are too big.

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@Sun# 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@Sun# delete interfaces ae0 mtu

```

```

[edit]
lab@Sun# commit
commit complete

```

After committing the changes R1 has two neighbors in the “Up” state.

```

lab@Sun# run show isis adjacency

```

Interface	System	L	State	Hold (secs)	SNPA
ae0.0	Sirius	2	Up	24	
ge-0/0/4.118	Procyon	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.

Activate IS-IS traceoptions on R4 for collecting errors generated by the protocol.

```

[edit]
lab@Arcturus# set protocols isis traceoptions flag error detail

```

```

[edit]
lab@Arcturus# commit
commit complete

```

Analyze the output in the isis-debug file.

```

lab@Arcturus# run show log isis-debug
Jun 29 14:27:28 trace_on: Tracing to "/var/log/isis-debug" 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 Canopus 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 Canopus 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@Vega# set protocols isis traceoptions flag error detail

[edit]
lab@Vega# commit
commit complete

```

Analyze the output from the file.

```

lab@Vega# 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.

```

lab@Canopus# 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@Canopus# rename interface lo0.0 family iso address
49.0001.1720.3000.5003.00 to address 49.0002.1720.3000.5003.00

```



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

Verify the adjacency table on R3 after changing the Area ID reveals that R3 has established adjacencies with R4 and R7.

```
[edit]
lab@Canopus# run show isis adjacency
```

Interface	System	L	State	Hold (secs)	SNPA
ge-0/0/4.123	Sirius	2	Up	25	
ge-0/0/4.134	Arcturus	1	Up	6	
f8:c0:1:dc:31:84					
ge-0/0/4.136	Vega	1	Up	7	
f8:c0:1:dc:2c:4					

c. R4 – R5 adjacency.

```
[edit]
lab@Arcturus# run show isis adjacency
```

Interface	System	L	State	Hold (secs)	SNPA
ge-0/0/4.134	Canopus	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.

```
lab@Arcturus# run show log isis-debug
Jun 29 14:27:28 trace on: Tracing to "/var/log/isis-debug" 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 Canopus 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 Canopus 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 an address mismatch on interface ge-0/0/4.145. Verifying the interface configuration on both routers reveals the mistake.

```
[edit]
lab@Arcturus# 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@A-Centauri# 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 command **rename**.

```

[edit]
lab@A-Centauri# edit interfaces ge-0/0/4.145

[edit interfaces ge-0/0/4 unit 145]
lab@A-Centauri# rename family inet address 172.30.1.30/30 to
address 172.30.0.30/30

[edit interfaces ge-0/0/4 unit 145]
lab@A-Centauri# commit
commit complete

```

After this change, R4 has established an adjacency with R5 also.

```

[edit]
lab@Arcturus# run show isis adjacency

```

Interface	System	L	State	Hold (secs)	SNPA
ge-0/0/4.134	Canopus	1	Up	23	
2c:21:72:cd:26:84					
ge-0/0/4.145	A-Centauri	1	Up	8	
f8:c0:1:dd:3:84					

d. R6 – R7 adjacency.

Activate traceoptions on R6 and R7 with the **hello** flag.

```

[edit]
lab@Vega# set protocols isis traceoptions file isis.log

[edit]
lab@Vega# set protocols isis traceoptions flag hello detail

[edit]
lab@Vega# commit

```

```
commit complete
```

```
[edit]
lab@Rigel# set protocols isis traceoptions file isis.log
```

```
[edit]
lab@Rigel# set protocols isis traceoptions flag hello detail
```

```
lab@Rigel# commit
commit complete
```

Verify the trace files on both routers.

```
[edit]
lab@Vega# 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
Vega.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
Vega.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
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@Rigel# 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 12
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@Vega# run show isis statistics
IS-IS statistics for Vega:
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

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@Rigel# run show isis statistics
IS-IS statistics for Rigel:
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, ISIS 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@Vega# set protocols isis interface ge-0/0/4.167 point-to-point

[edit]
lab@Vega# commit
commit complete

```

```

[edit]
lab@Vega# run show isis adjacency
Interface          System          L State          Hold (secs) SNPA
ae0.0              A-Centauri      1 Up              6
f8:c0:1:dd:4:0
ge-0/0/4.136       Canopus         1 Up              26
2c:21:72:cd:26:84
ge-0/0/4.167       Rigel           2 Up              23

```

After the change, R6 established an adjacency with R7.

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@Rigel# 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@Procyon# 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 Procyon.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.

This might be a problem with the configuration on R8. Verify if this interface is added to **protocols isis** stanza.

```

[edit]
lab@Procyon# 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@Procyon# 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@Procyon# set interfaces ge-0/0/4.178 family iso

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

```
[edit]
lab@Procyon# run show isis adjacency
```

Interface	System	L	State	Hold (secs)	SNPA
ge-0/0/4.118	Rigel	1	Up	20	
0:26:88:ef:75:84					
ge-0/0/4.178	Rigel	2	Up	23	

Enable 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 “Sun” on R1. But after the change, the adjacency table shows R7’s hostname “Rigel” 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@Sun# run show isis database
```

IS-IS level 1 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
Sun.00-00	0xd	0x78ad	540	L1 L2
Attached				
Procyon.00-00	0x7	0x113a	758	L1 L2
Attached				
Procyon.02-00	0x5	0xe814	755	L1 L2

3 LSPs

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
Sun.00-00	0x10	0x8690	929	L1 L2
Sirius.00-00	0x11	0xe48d	446	L1 L2

2 LSPs

As Level 1 and Level 2 router, it has separate databases for both Levels. LSP Procyon.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@Sirius# run show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
Sirius.00-00          0xb    0xf8ea    1055 L1 L2
1 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
Sirius.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.

```
lab@Canopus# run show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
Canopus.00-00         0xa    0x43a6    1134 L1 L2
Attached
Arcturus.00-00        0xa    0x163     1195 L1 L2
Arcturus.02-00        0x7    0x296c    867 L1 L2
A-Centauri.00-00      0x9    0xb32a    1185 L1 L2
A-Centauri.02-00      0x5    0x43bc    537 L1 L2
A-Centauri.03-00      0x3    0xcc30    1185 L1 L2
Vega.00-00            0x6    0x2967    383 L1 L2
Attached
Vega.02-00            0x5    0x25d9    383 L1 L2
8 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
Sirius.00-00          0x12   0xc84b    1149 L1 L2
Canopus.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@Sirius# 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
```


Verify the IS-IS statistics for dropped PDUs.

```
[edit]
lab@Sirius# run show isis statistics
IS-IS statistics for Sirius:
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 checksums 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 (Hello, 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** stanza, 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@Sirius# 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

[edit]
lab@Sirius# 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@Sirius# delete protocols isis level 2 authentication-type

[edit]
lab@Sirius# delete protocols isis level 2 authentication-key

[edit]
lab@Sirius# commit
commit complete

```

Check the LSDB after this change.

```

[edit]
lab@Sirius# run show isis database
IS-IS level 1 link-state database:

```

```

LSP ID                      Sequence Checksum Lifetime Attributes
Sirius.00-00                0x14    0xeeee3      815 L1 L2
Attached
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
Sun.00-00                   0x61    0xe3e1       517 L1 L2
Sirius.00-00                0x24    0xfe2a       822 L1 L2
Canopus.00-00               0xd     0x5c08      1089 L1 L2
Vega.00-00                  0x8     0          0 L1 L2
Procyon.00-00               0x9     0          0 L1 L2
  5 LSPs

```

Now the L2 database looks different. 5 LSPs are stored and used in IS-IS.

You can see that the R7 LSP is missing in the database. Checking the IS-IS adjacency shows the problem with R7.

```

[edit]
lab@Sirius# run show isis adjacency
Interface          System          L State          Hold (secs)
SNPA
ae0.0              Sun             2 Up             26
ge-0/0/4.123       Canopus         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@Sirius# edit protocols isis

[edit protocols isis]
lab@Sirius# set interface ge-0/0/4.127 level 2 hello-
authentication-type md5

[edit protocols isis]
lab@Sirius# set interface ge-0/0/4.127 level 2 hello-
authentication-key "$9$MusXVY.mTF6ADiqlfz6u0M8X-b2"

[edit protocols isis]
lab@Sirius# commit
commit complete

```

Configure Hello authentication under the interface to R7 fixes the problem.

```

[edit]
lab@Sirius# run show isis database
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
Sirius.00-00                0x4     0xfd3       1172 L1 L2
Attached
  1 LSPs

```

```
IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
Sun.00-00              0x65   0xdbe5    1180 L1 L2
Sirius.00-00           0x2e   0xdc81    1178 L1 L2
Canopus.00-00          0x18   0x80d8     528 L1 L2
Vega.00-00             0xe    0xf093     462 L1 L2
Rigel.00-00            0x9    0x734c    1170 L1 L2
Procyon.00-00          0xe    0xd7b1    1007 L1 L2
6 LSPs
```

All six LSPs of the L1 routers are in the database.

- 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@Sun# run show isis hostname
IS-IS hostname database:
System ID      Hostname      Type
1720.3000.5001 Sun          Static
1720.3000.5002 Sirius       Dynamic
1720.3000.5003 Canopus    Dynamic
1720.3000.5006 Vega       Dynamic
1720.3000.5008 Procyon    Dynamic
```

The output displays the SysID with the mapped hostname and the way it is learned. R7 is missing in the database.

- b. R2.

```
[edit]
lab@Sirius# run show isis hostname
IS-IS hostname database:
System ID      Hostname      Type
1720.3000.5001 Rigel        Dynamic
1720.3000.5002 Sirius       Static
1720.3000.5003 Canopus    Dynamic
1720.3000.5006 Vega       Dynamic
1720.3000.5008 Procyon    Dynamic
```

The output reveals that the SysID of R1 is mapped to the hostname of R7.

- c. R7.

```
[edit]
lab@Rigel# 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 'Rigel'.

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@Rigel# rename interface lo0.0 family iso address
49.0002.1720.3000.5001.00 to address 49.0002.1720.3000.5007.00

[edit]
lab@Rigel# commit
commit complete

```

The new SysID on R7 is now displayed in the hostname database.

```

[edit]
lab@Rigel# run show isis hostname
IS-IS hostname database:
System ID      Hostname      Type
1720.3000.5001 Sun
Dynamic
1720.3000.5002 Sirius
Dynamic
1720.3000.5003 Canopus
Dynamic
1720.3000.5006 Vega
Dynamic
1720.3000.5007 Rigel      Static
1720.3000.5008 Procyon
Dynamic

```

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@Arcturus# run show route

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] 00:32:20, 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.

```
[edit]
lab@@Arcturus# show protocols rip
group rip {
    export isis-to-rip;
    neighbor ge-0/0/4.202;
```

```

    }

[edit]
lab@Arcturus# run show rip neighbor

```

In Neighbor Met	Local State	Source Address	Destination Address	Send Mode	Receive Mode
ge-0/0/4.202 1	Up	172.30.0.49	224.0.0.9	mcast	both

The routing policy **isis-to-rip** exports all IS-IS routes according to the task requirements.

```

[edit]
lab@Arcturus# 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.

```

lab@Arcturus# run show route receive-protocol rip 172.30.0.50

inet.0: 48 destinations, 50 routes (48 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
> 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@Arcturus# run show route advertising-protocol rip 172.30.0.49

inet.0: 48 destinations, 50 routes (48 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
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@A-Centauri# show protocols rip
group rip {
    export isis-to-rip;
    neighbor ge-0/0/4.204;
}

[edit]
lab@A-Centauri# run show rip neighbor
```

In Neighbor Met	Local State	Source Address	Destination Address	Send Mode	Receive Mode
-----	-----	-----	-----	----	-----

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@ A-Centauri# show policy-options policy-statement isis-to-rip
term 1 {
    from protocol isis;
    then {
        metric 5;
        accept;
    }
}
```

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@A-Centauri# run show route receive-protocol rip 172.30.0.58

inet.0: 44 destinations, 46 routes (44 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
                   > 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.

```
lab@A-Centauri# 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] 00:00:43, metric 10
                    > to 172.30.0.34 via ae0.0
172.30.0.20/30     *[IS-IS/15] 00:25:51, metric 20
                    > to 172.30.0.29 via ge-0/0/4.145
172.30.0.24/30     *[IS-IS/15] 00:00:43, 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] 00:00:43, 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] 00:25:51, metric 10
                    > to 172.30.0.29 via ge-0/0/4.145
```

R5 is not advertising prefixes in the range 192.168.8.0/21. Let's check the routing table for these prefixes.

```
[edit]
lab@A-Centauri# 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] 02:30:23, 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@A-Centauri# run show isis database level 1 Canopus.00-00
extensive
IS-IS level 1 link-state database:

Canopus.00-00 Sequence: 0xc, Checksum: 0x9fc9, Lifetime: 1128 secs
  IS neighbor: Arcturus.00                      Metric: 10
  Two-way fragment: Arcturus.00-00, Two-way first fragment:
Arcturus.00-00
  IS neighbor: Vega.00                          Metric: 10
  Two-way fragment: Vega.00-00, Two-way first fragment:
Vega.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
  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@A-Centauri# run show isis database level 1 Vega.00-00
extensive
IS-IS level 1 link-state database:

Vega.00-00 Sequence: 0xa, Checksum: 0x808e, Lifetime: 344 secs
  IS neighbor: Canopus.00              Metric:           10
  Two-way fragment: Canopus.00-00, Two-way first fragment:
Canopus.00-00
  IS neighbor: A-Centauri.00            Metric:           10
  Two-way fragment: A-Centauri.00-00, Two-way first fragment:
A-Centauri.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.

Other 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@Arcturus# show policy-options policy-statement isis-to-rip
term 1 {
    from protocol isis;
    then {
        metric 1;
    }
}

```

```

    accept;
}
}

```

```

[edit]
lab@Arcturus# set policy-options policy-statement isis-to-rip term
1 then tag 1234

```

```

[edit]
lab@Arcturus# commit
commit complete

```

The routing policy **isis-to-rip** exports the IS-IS prefixes with metric 1 and tag of 1234.

```

[edit]
lab@Arcturus# 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@A-Centauri# set policy-options policy-statement isis-to-rip
term 1 then tag 1234

```

```

[edit]
lab@A-Centauri# commit
commit complete

```

The result is that R5 receives the RIP routes advertised by R4 with a tag of 1234.

```

lab@A-Centauri# 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
                   > 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.

```

[edit]
lab@Arcturus# set policy-options policy-statement rip-filter term
1 from protocol rip

[edit]
lab@Arcturus# set policy-options policy-statement rip-filter term
1 from tag 1234

[edit]
lab@Arcturus# set policy-options policy-statement rip-filter term
1 then reject

[edit]
lab@Arcturus# set protocols rip group rip import rip-filter

[edit]
lab@Arcturus# commit
commit complete

```

After filtering the routes, R5 advertises the IS-IS routes.

```

lab@A-Centauri# 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@Arcturus# show policy-options policy-statement rip-to-isis
term 1 {
    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@Canopus# 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@Canopus# 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@Canopus# run show isis database level 2 Canopus.00-00
extensive
IS-IS level 2 link-state database:

Canopus.00-00 Sequence: 0xd, Checksum: 0xb8ab, Lifetime: 1092 secs
  IS neighbor: Sirius.00                      Metric:          10
    Two-way fragment: Sirius.00-00, Two-way first fragment:
Sirius.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
  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@Canopus# run show isis database level 2 Vega.00-00 extensive
IS-IS level 2 link-state database:

Vega.00-00 Sequence: 0xb, Checksum: 0xcb74, Lifetime: 1167 secs
  IS neighbor: Rigel.00                Metric:      10
  Two-way fragment: Rigel.00-00, Two-way first fragment:
Rigel.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.

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@Canopus# run show isis database level 1 A-Centauri.00-00
extensive
IS-IS level 1 link-state database:

```

```

A-Centauri.00-00 Sequence: 0xc, Checksum: 0xb386, Lifetime: 1138
secs
  IS neighbor: Arcturus.00                      Metric:      10
    Two-way fragment: Arcturus.00-00, Two-way first fragment:
Arcturus.00-00
  IS neighbor: Vega.00                          Metric:      10
    Two-way fragment: Vega.00-00, Two-way first fragment:
Vega.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 readvertising those routes to Level 2 by the L1/L2 routers.

```

[edit]
lab@Canopus# run show isis database level 1 Arcturus.00-00
extensive

```

```

IS-IS level 1 link-state database:

Arcturus.00-00 Sequence: 0xb, Checksum: 0xfedf, Lifetime: 1101
secs
  IS neighbor: Canopus.00                      Metric:      10
    Two-way fragment: Canopus.00-00, Two-way first fragment:
Canopus.00-00
  IS neighbor: A-Centauri.00                    Metric:      10
    Two-way fragment: A-Centauri.00-00, Two-way first fragment:
A-Centauri.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
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 ISIS 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@Arcturus# show protocols isis level 1
wide-metrics-only;
```

As 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@Arcturus# delete protocols isis level 1 wide-metrics-only
```

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

Verify the level 2 LSP of R3 if the more specific routes are suppressed.

```
[edit]
lab@Canopus# run show isis database level 2 Canopus.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.

```

lab@Canopus# run show isis database level 1 Arcturus.00-00
extensive | no-more
IS-IS level 1 link-state database:

Arcturus.00-00 Sequence: 0xc, Checksum: 0x8240, Lifetime: 1136
secs
  IS neighbor: Canopus.00                      Metric:      10
  Two-way fragment: Canopus.00-00, Two-way first fragment:
Canopus.00-00
  IS neighbor: A-Centauri.00                    Metric:      10
  Two-way fragment: A-Centauri.00-00, Two-way first fragment:
A-Centauri.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		

The redistributed prefixes are now advertised as external.

13) Verify Loopback reachability.

a. R1.

```
[edit]
lab@Sun# 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@Canopus# 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
                  > 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@Arcturus# 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@Arcturus# 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.

```
[edit]
lab@Canopus# run show isis database level 1 Vega.00-00 extensive
IS-IS level 1 link-state database:

Vega.00-00 Sequence: 0xb, Checksum: 0x7e8f, Lifetime: 1187 secs
  IS neighbor: Canopus.00 Metric: 10
    Two-way fragment: Canopus.00-00, Two-way first fragment:
Canopus.00-00
  IS neighbor: A-Centauri.00 Metric: 10
    Two-way fragment: A-Centauri.00-00, Two-way first fragment:
A-Centauri.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@Canopus# run show isis database level 2 Vega.00-00 extensive
IS-IS level 2 link-state database:

Vega.00-00 Sequence: 0xb, Checksum: 0xcb74, Lifetime: 1167 secs
  IS neighbor: Rigel.00 Metric: 10
    Two-way fragment: Rigel.00-00, Two-way first fragment:
Rigel.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@Vega# show protocols isis export
export [ inject-externals l1-to-l2 ];

[edit]
lab@Vega# show policy-options policy-statement inject-externals
term 1 {
    from {
        protocol static;
        route-filter 192.168.12.0/22 longer;
    }
    then accept;
}

[edit]
lab@Vega# 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@Vega# delete policy-options policy-statement l1-to-l2 term 2

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

Verify if the change has effect and the Loopback IP address of R6 is reachable.

On R5 Level 1.

```
[edit]
lab@A-Centauri# 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
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@Sirius# 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 **11-to-12** export policy. After this change, R6 LSP for Level 2 holds all other Internal prefixes that were missing before.

```
lab@Canopus# run show isis database level 2 Vega.00-00 extensive
IS-IS level 2 link-state database:

Vega.00-00 Sequence: 0xc, Checksum: 0x8956, Lifetime: 1137 secs
  IS neighbor: Rigel.00 Metric: 10
  Two-way fragment: Rigel.00-00, Two-way first fragment:
Rigel.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
  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@Procyon# 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@Procyon# 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** stanza on R6.

```

[edit]
lab@Vega# deactivate protocols isis

[edit]
lab@Vega# commit
commit complete

```

Again perform a traceroute from R8.

```

[edit]
lab@Procyon# 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@Vega# rollback 1
load complete

[edit]
lab@Vega# commit
commit complete

```

As a last step don't forget 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 of chapter One is used as base configuration. Figure 6 describes the topology used for the IGP network.

- 1) Load the previous saved configuration from Chapter One.

```
[edit]
root@mx1# load override terminal
```

Again using the **load override terminal** command helps with the loading of the configuration from Chapter One on all 8 devices.

- 2) Configure the additional interfaces on R4 and R5.

- a. R4.

Nothing special here. You have to configure two new subinterfaces connected to DC2 and DC3 routers according table 11.

```
[edit]
lab@Arcturus# set interfaces ge-0/0/4.202 vlan-id 202

[edit]
lab@Arcturus# set interfaces ge-0/0/4.202 description "DC2
connection"

[edit]
lab@Arcturus# set interfaces ge-0/0/4.202 family inet address
172.30.0.49/30

[edit]
lab@Arcturus# set interfaces ge-0/0/4.203 vlan-id 203

[edit]
lab@Arcturus# set interfaces ge-0/0/4.203 description "DC3
connection"

[edit]
```

```
lab@Arcturus# set interfaces ge-0/0/4.203 family inet address
172.30.0.53/30
```

```
[edit]
```

```
lab@Arcturus# set interfaces ge-0/0/4.203 family inet6
```

```
[edit groups]
```

```
lab@Arcturus# 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.

```
[edit]
lab@Arcturus# 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@Arcturus# 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@A-Centauri# set interfaces ge-0/0/4.204 vlan-id 204
```

```
[edit]
```

```
lab@A-Centauri# set interfaces ge-0/0/4.204 description "DC2
connection"
```

```
[edit]
```

```
lab@A-Centauri# set interfaces ge-0/0/4.204 family inet address
172.30.0.57/30
```

```
[edit]
```

```
lab@A-Centauri# set interfaces ge-0/0/4.205 vlan-id 205
```

```
[edit]
```

```
lab@A-Centauri# set interfaces ge-0/0/4.205 description "DC3
connection"
```

```
[edit]
```

```
lab@A-Centauri# set interfaces ge-0/0/4.205 family inet address
172.30.0.61/30
```



```
[edit]
lab@A-Centauri# set interfaces ge-0/0/4.205 family inet6

[edit groups]
lab@A-Centauri# commit
commit complete
```

Again, verify the connectivity with ping.

```
[edit]
lab@A-Centauri# 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@A-Centauri# 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 12 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 subinterfaces of ge-0/0/4 and aggregated Ethernet. So you can easily match them.

```
[edit]
lab@Sun# set groups if-families interfaces ge-0/0/4 unit <*>
family iso

[edit]
lab@Sun# set groups if-families interfaces <ae*> unit <*> family
iso
```

NOTE: For match conditions you can use wildcards similar to regular expressions. You have to use angle brackets "<>" to define the match condition with wildcards. "*" is very common 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 stanza **groups**.

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 have to apply the configured group and commit the configuration.

```
[edit]
lab@Sun# set apply-groups if-families

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

To verify the apply-groups, you have to display the configuration with an additional option.

```
[edit]
lab@Sun# 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 {
      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;
        }
    }
}
}

```

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 12. All 8 routers are in the same AreaID. The SysID part from the NET address is built from the Loopback IP address.

```
[edit]
lab@Sun# set interface lo0.0 family iso address
49.0001.1720.3000.5001.00
```

Configure all other router Loopback interfaces with the NET addresses specified in Table 12 the same way.

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.

```
[edit]
lab@Sun# set routing-options router-id 172.30.5.1
```

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.

```
[edit]
lab@Sun# set protocols isis level 2 disable
```

The interface metric should be calculated based on the bandwidth of the interfaces.

```
[edit]
lab@Sun# 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.

```
[edit]
lab@Sun# 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.

```
[edit]
lab@Sun# set protocols isis level 1 authentication-type md5
```

```
[edit]
lab@Sun# set protocols isis level 1 authentication-key workbook
```

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.

```
[edit]
lab@Sun# 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 500ms. IS-IS internal Hello mechanism supports hold timer of minimum 3 seconds and a hello interval of minimum 1 second. JUNOS allows to a configured hold timer of 1 second, which means the hello interval is 333ms. However the requirement states to detect adjacency loss in less than 500ms.

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 have to configure **bfd-liveness-detection** under the interface configuration.

```
[edit]
lab@Sun# set protocol isis interface all bfd-liveness-detection
minimum-interval 150
```

```
[edit]
lab@Sun# set protocols isis interface all bfd-liveness-detection
multiplier 3
```

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.

```
[edit]
lab@Canopus# set protocols isis interface ge-0/0/4.200 passive

[edit]
lab@Canopus# set protocols isis interface ge-0/0/4.201 passive

[edit]
lab@Arcturus# set protocols isis interface ge-0/0/4.200 passive

[edit]
lab@Arcturus# set protocols isis interface ge-0/0/4.201 passive
```

f. Verify the configuration.

First you have to commit the configuration and configure the other routers in the same way. As the configuration is the same you can again use **load merge terminal**.

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

First you have to 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@Sun# run show isis adjacency
```

Interface	System	L	State	Hold (secs)
SNPA				
ae0.0	Sirius	1	Up	20
ge-0/0/4.114	Arcturus	1	Up	26
ge-0/0/4.118	Procyon	1	Up	25

```
[edit]
lab@Canopus# run show isis adjacency
```

Interface	System	L	State	Hold (secs)
SNPA				
ge-0/0/4.123	Sirius	1	Up	23
ge-0/0/4.134	Arcturus	1	Up	25
ge-0/0/4.136	Vega	1	Up	26

```
[edit]
lab@Vega# run show isis adjacency
```

Interface	System	L	State	Hold (secs)
SNPA				
ae0.0	A-Centauri	1	Up	25
ge-0/0/4.136	Canopus	1	Up	21
ge-0/0/4.167	Rigel	1	Up	24

```
[edit]
```

```
lab@Procyon# run show isis adjacency
```

Interface	System	L	State	Hold (secs)
SNPA				
ge-0/0/4.118	Sun	1	Up	22
ge-0/0/4.158	A-Centauri	1	Up	20
ge-0/0/4.178	Rigel	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@Sun# run show isis database
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
Sun.00-00                    0xd     0x16f1    1178 L1
Sirius.00-00                  0xe     0xbd4f    1034 L1
Canopus.00-00                 0xe     0xe1a2    1174 L1
Arcturus.00-00                0x8     0xefae    1176 L1
A-Centauri.00-00              0xc     0xa426    1150 L1
Vega.00-00                    0x8     0x45bb    1057 L1
Rigel.00-00                   0x6     0x614a    1073 L1
Procyon.00-00                 0x6     0x5fe1    1081 L1
8 LSPs
```

```
IS-IS level 2 link-state database:
0 LSPs
```

```
[edit]
lab@Arcturus# run show isis database
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
Sun.00-00                    0x19    0xc3d5     962 L1
Sirius.00-00                  0x1b    0x8904    1015 L1
Canopus.00-00                 0x13    0x4b46    1008 L1
Arcturus.00-00                0x1b    0x8843    1035 L1
A-Centauri.00-00              0x1d    0xbb26    1196 L1
Vega.00-00                    0xc     0xd7dd     885 L1
Rigel.00-00                   0xb     0x5fca     871 L1
Procyon.00-00                 0xb     0x142b     824 L1
8 LSPs

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@Sun# run show isis database level 1 Sirius.00-00 extensive
IS-IS level 1 link-state database:

Sirius.00-00 Sequence: 0xe, Checksum: 0xbd4f, Lifetime: 955 secs
```

```

IS neighbor: Sun.00                      Metric:      5
  Two-way fragment: Sun.00-00, Two-way first fragment: Sun.00-00
IS neighbor: Canopus.00                  Metric:      10
  Two-way fragment: Canopus.00-00, Two-way first fragment: Canopus.00-00
IS neighbor: Rigel.00                    Metric:      10
  Two-way fragment: Rigel.00-00, Two-way first fragment: Rigel.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 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@Canopus# run show route 172.30.5/24

inet.0: 50 destinations, 50 routes (50 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
                   > 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@A-Centauri# 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/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@Sirius# run show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	
Multiplier					
172.30.0.1	Up	ae0.0	0.450	0.150	3

```
Client ISIS L1, TX interval 0.150, RX interval 0.150
Session up time 01:29:44
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Min async interval 0.150, min slow interval 1.000
Adaptive async TX interval 0.150, RX interval 0.150
Local min TX interval 0.150, minimum RX interval 0.150,
multiplier 3
Remote min TX interval 0.150, min RX interval 0.150, multiplier 3
Local discriminator 1, remote discriminator 1
Echo mode disabled/inactive
```

Address	State	Interface	Detect Time	Transmit Interval	
Multiplier					
172.30.0.14	Up	ge-0/0/4.123	0.450	0.150	3

```
Client ISIS L1, TX interval 0.150, RX interval 0.150
Session up time 01:28:24
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Min async interval 0.150, min slow interval 1.000
```

```

Adaptive async TX interval 0.150, RX interval 0.150
Local min TX interval 0.150, minimum RX interval 0.150,
multiplier 3
Remote min TX interval 0.150, min RX interval 0.150, multiplier 3
Local discriminator 2, remote discriminator 1
Echo mode disabled/inactive

Address          State      Interface      Detect    Transmit
Multiplier
172.30.0.18      Up         ge-0/0/4.127   0.450     0.150      3
Client ISIS L1, TX interval 0.150, RX interval 0.150
Session up time 00:03:46
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Min async interval 0.150, min slow interval 1.000
Adaptive async TX interval 0.150, RX interval 0.150
Local min TX interval 0.150, minimum RX interval 0.150,
multiplier 3
Remote min TX interval 0.150, min RX interval 0.150, multiplier 3
Local discriminator 3, remote discriminator 1
Echo mode disabled/inactive

3 sessions, 3 clients
Cumulative transmit rate 20.0 pps, cumulative receive rate 20.0
pps

```

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 ISIS L1, transmit and receive intervals are 150ms. The multiplier is 3, which results in a protocol detection time of approximately of 450ms.

Verify the advertisement of directly connected VRRP interfaces on R3 and R4.

```
lab@Sun# run show isis database level 1 Canopus.00-00 extensive
IS-IS level 1 link-state database:

Canopus.00-00 Sequence: 0xe, Checksum: 0xela2, Lifetime: 1096 secs
  IS neighbor: Sirius.00 Metric: 10
    Two-way fragment: Sirius.00-00, Two-way first fragment:
Sirius.00-00
  IS neighbor: Arcturus.00 Metric: 10
    Two-way fragment: Arcturus.00-00, Two-way first fragment:
Arcturus.00-00
  IS neighbor: Vega.00 Metric: 10
    Two-way fragment: Vega.00-00, Two-way first fragment:
Vega.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@Sun# run show isis database level 1 Arcturus.00-00 extensive
IS-IS level 1 link-state database:

Arcturus.00-00 Sequence: 0x8, Checksum: 0xefae, Lifetime: 1098
secs
  IS neighbor: Sun.00 Metric: 10
    Two-way fragment: Sun.00-00, Two-way first fragment: Sun.00-
00
  IS neighbor: Canopus.00 Metric: 10
    Two-way fragment: Canopus.00-00, Two-way first fragment:
Canopus.00-00
  IS neighbor: A-Centauri.00 Metric: 10
    Two-way fragment: A-Centauri.00-00, Two-way first fragment:
A-Centauri.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	

R3 and R4 advertise the VRRP prefixes in their LSPs. The interfaces are listed as passive in the IS-IS interface table.

```
[edit]
lab@Canopus# 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   0    0x1 Passive         Passive
10/10
ge-0/0/4.201   0    0x1 Passive         Passive
10/10
lo0.0          0    0x1 Passive         Passive
0/0
lo0.32768      0    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@Arcturus# set protocols rip group dc2 neighbor ge-0/0/4.202

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

Verify RIP neighbor.

```
[edit]
lab@Arcturus# 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.

```
[edit]
lab@Arcturus# run show route advertising-protocol rip 172.30.0.49
```

```
[edit]
lab@Arcturus# run show route receive-protocol rip 172.30.0.50

inet.0: 54 destinations, 54 routes (54 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

172.30.32.0/24    *[RIP/100] 00:03:29, 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.

```
[edit]
lab@A-Centauri# set protocols rip group dc2 neighbor ge-0/0/4.204
[edit]
lab@A-Centauri# commit
commit complete
```

Verify RIP operation.

```
[edit]
lab@A-Centauri# run show rip neighbor
```

In Neighbor Met	Local State	Source Address	Destination Address	Send Mode	Receive Mode
---	----	-----	-----	----	-----

ge-0/0/4.204	Up	172.30.0.57	224.0.0.9	mcast	both
1					

```
lab@A-Centauri# run show route receive-protocol rip 172.30.0.58
```

```
inet.0: 52 destinations, 52 routes (52 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
                  > 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 have to create the aggregate default route.

```
[edit]
lab@Arcturus# set routing-options aggregate route 0.0.0.0/0

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

Verify if the aggregate route exist in the routing table.

```
[edit]
lab@Arcturus# 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 have to 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@Arcturus# set policy-options policy-statement agg-to-rip term
1 from protocol aggregate

[edit]
lab@Arcturus# set policy-options policy-statement agg-to-rip term
1 from route-filter 0.0.0.0/0 exact

[edit]
lab@Arcturus# set policy-options policy-statement agg-to-rip term
1 then tag 123

[edit]
lab@Arcturus# 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@Arcturus# set protocols rip group dc2 export agg-to-rip

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

Verify the export of the default route to DC2.

```
[edit]
lab@Arcturus# 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 have to configure the aggregate default route and the same export policy as on R4.

```
[edit]
lab@A-Centauri# set routing-options aggregate route 0.0.0.0/0

[edit]
lab@A-Centauri# set policy-options policy-statement agg-to-rip
term 1 from protocol aggregate

[edit]
lab@A-Centauri# set policy-options policy-statement agg-to-rip
term 1 from route-filter 0.0.0.0/0 exact

[edit]
lab@A-Centauri# set policy-options policy-statement agg-to-rip
term 1 then tag 123

[edit]
lab@A-Centauri# 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@A-Centauri# set protocols rip group dc2 export agg-to-rip
```

```
[edit]
lab@A-Centauri# commit
commit complete
```

Verify the result of the RIP export policy.

```
[edit]
lab@A-Centauri# run show route advertising-protocol rip
172.30.0.57
```

R5 is not advertising anything to DC2. Verify the routing table for the aggregate default route.

```
[edit]
lab@A-Centauri# 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 have to filter the default route coming from RIP at R5 and R4.

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@Arcturus# set policy-options policy-statement filter-rip term
1 from protocol rip
```

```
[edit]
lab@Arcturus# set policy-options policy-statement filter-rip term
1 from tag 123
```

```
[edit]
lab@Arcturus# 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@Arcturus# set protocols rip group dc2 import filter-rip

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

b. R5.

On R5 you have to 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@A-Centauri# set policy-options policy-statement filter-rip
term 1 from protocol rip

[edit]
lab@A-Centauri# set policy-options policy-statement filter-rip
term 1 from tag 123

[edit]
lab@A-Centauri# set policy-options policy-statement filter-rip
term 1 then reject
```

Apply the policy and commit.

```
[edit]
lab@A-Centauri# set protocols rip group dc2 import filter-rip

[edit]
lab@A-Centauri# commit
commit complete
```

Verify if the default route from protocol RIP is filtered.

```
[edit]
lab@A-Centauri# 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@A-Centauri# 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@Arcturus# 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.

On R4 and R5 you have to redistribute the other RIP routes into IS-IS.

a. R4.

Use an ISIS export policy to match all RIP routes .

```
[edit]
lab@Arcturus# set policy-options policy-statement rip-to-isis term
1 from protocol rip
```

```
[edit]
lab@Arcturus# set policy-options policy-statement rip-to-isis term
1 then accept
```

```
[edit]
lab@Arcturus# set protocols isis export rip-to-isis
```

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

b. R5.

Implement the same export policy on R5.

```
[edit]
lab@A-Centauri# set policy-options policy-statement rip-to-isis
term 1 from protocol rip
```

```
[edit]
lab@A-Centauri# set policy-options policy-statement rip-to-isis
term 1 then accept
```

```
[edit]
lab@A-Centauri# set protocols isis export rip-to-isis
```

```
[edit]
lab@A-Centauri# commit
commit complete
```

Verify the redistribution of the RIP routes to IS-IS.

```
[edit]
lab@Arcturus# run show isis database Arcturus.00-00 extensive
IS-IS level 1 link-state database:

Arcturus.00-00 Sequence: 0x1b, Checksum: 0x8843, Lifetime: 1017
secs
  IS neighbor: Sun.00                      Metric:          10
    Two-way fragment: Sun.00-00, Two-way first fragment: Sun.00-
00
  IS neighbor: Canopus.00                  Metric:          10
    Two-way fragment: Canopus.00-00, Two-way first fragment:
Canopus.00-00
  IS neighbor: A-Centauri.00               Metric:          10
    Two-way fragment: A-Centauri.00-00, Two-way first fragment:
A-Centauri.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@A-Centauri# run show isis database A-Centauri.00-00 extensive
IS-IS level 1 link-state database:

A-Centauri.00-00 Sequence: 0x1e, Checksum: 0x3a9a, Lifetime: 1144
secs
  IS neighbor: Arcturus.00                      Metric:          10
  Two-way fragment: Arcturus.00-00, Two-way first fragment:
Arcturus.00-00
  IS neighbor: Vega.00                          Metric:          5
  Two-way fragment: Vega.00-00, Two-way first fragment:
Vega.00-00
  IS neighbor: Procyon.00                      Metric:          10
  Two-way fragment: Procyon.00-00, Two-way first fragment:
Procyon.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
  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@A-Centauri# 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.

```
[edit]
lab@Arcturus# set protocols rip group dc2 preference 14

[edit]
lab@Arcturus# commit
commit complete

[edit]
lab@A-Centauri# set protocols rip group dc2 preference 14

[edit]
lab@A-Centauri# commit
commit complete
```

Verify the results of the change.

```
[edit]
lab@A-Centauri# 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@A-Centauri# run show isis database A-Centauri.00-00 extensive
IS-IS level 1 link-state database:

A-Centauri.00-00 Sequence: 0x22, Checksum: 0xa3b0, Lifetime: 1134
secs
  IS neighbor: Arcturus.00                      Metric:      10
  Two-way fragment: Arcturus.00-00, Two-way first fragment:
Arcturus.00-00
  IS neighbor: Vega.00                          Metric:      5
  Two-way fragment: Vega.00-00, Two-way first fragment:
Vega.00-00
  IS neighbor: Procyon.00                       Metric:      10
  Two-way fragment: Procyon.00-00, Two-way first fragment:
Procyon.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.

Check the IPv6 routing table on R1.

```
[edit]
lab@Sun# run show route table inet6 | no-more

inet6.0: 11 destinations, 12 routes (11 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691:5::1/128
    *[Direct/0] 02:47:31
    > via lo0.0
fd17:f0f4:f691:5::2/128
    *[IS-IS/15] 01:31:43, metric 5
    > to fe80::fac0:1ff:fedd:280 via ae0.0
fd17:f0f4:f691:5::3/128
    *[IS-IS/15] 01:31:43, metric 15
    > to fe80::fac0:1ff:fedd:280 via ae0.0
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
fd17:f0f4:f691:5::6/128
    *[IS-IS/15] 00:07:00, metric 25
    > to fe80::fac0:1ff:fedd:280 via ae0.0
    > to fe80::fac0:100:76dc:2e84 via ge-0/0/4.118
fd17:f0f4:f691:5::7/128
    *[IS-IS/15] 00:07:47, metric 15
    > to fe80::fac0:1ff:fedd:280 via ae0.0
fd17:f0f4:f691:5::8/128
    *[IS-IS/15] 00:07:00, metric 10
    > to fe80::fac0:100:76dc:2e84 via ge-0/0/4.118
fe80::/64
    *[Direct/0] 02:46:16
    > via ge-0/0/4.118
    [Direct/0] 02:46:12
    > via ae0.0
fe80::226:8800:76ef:7584/128
    *[Local/0] 02:46:32
    Local via ge-0/0/4.118
fe80::226:880f:fcef:7580/128
    *[Direct/0] 02:47:31
    > via lo0.0
fe80::226:88ff:feef:7600/128
    *[Local/0] 02:46:48
    Local via ae0.0
```

R1 has routes for some of the IPv6 Loopback interfaces but not for all. For example the R4 Loopback IPv6 address is not present.

Even some of the routes that are present in routing table are not reachable.

```
[edit]
lab@Sun# 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 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 IPv6 Loopback address is missing in the routing table. First check if the IPv6 prefix is present in the ISIS LSP sent from R4.

```
[edit]
lab@Sun# run show isis database Arcturus.00-00 extensive |match V6

V6 prefix: fd17:f0f4:f691:5::4/128          Metric:          0
Internal Up
```

The IPv6 route is present in the LSP sent from R4.

```
[edit]
lab@Sun# 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**.

```
[edit]
lab@Sun# 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@Sun# 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@Sun# 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@Sun# 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 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@Sun# set protocols isis topologies ipv6-unicast

[edit]
lab@Sun# 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@Sun# run show route table inet6 | no-more

inet6.0: 12 destinations, 13 routes (12 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

fd17:f0f4:f691:5::1/128
    *[Direct/0] 04:07:25
    > via lo0.0
fd17:f0f4:f691:5::2/128
    *[IS-IS/15] 00:01:08, metric 5
    > to fe80::fac0:1ff:fedd:280 via ae0.0
fd17:f0f4:f691:5::3/128
    *[IS-IS/15] 00:01:00, metric 15
    > to fe80::fac0:1ff:fedd:280 via ae0.0
fd17:f0f4:f691:5::4/128
    *[IS-IS/15] 00:00:50, metric 25
    > to fe80::fac0:1ff:fedd:280 via ae0.0
fd17:f0f4:f691:5::5/128
    *[IS-IS/15] 00:00:04, metric 35
    to fe80::fac0:1ff:fedd:280 via ae0.0
    > to fe80::fac0:100:76dc:2e84 via ge-0/0/4.118
fd17:f0f4:f691:5::6/128
    *[IS-IS/15] 00:00:04, metric 30
    > to fe80::fac0:100:76dc:2e84 via ge-0/0/4.118
fd17:f0f4:f691:5::7/128
    *[IS-IS/15] 00:00:04, metric 20
    > to fe80::fac0:100:76dc:2e84 via ge-0/0/4.118
fd17:f0f4:f691:5::8/128
    *[IS-IS/15] 00:00:04, metric 10
    > to fe80::fac0:100:76dc:2e84 via ge-0/0/4.118
fe80::/64
    *[Direct/0] 04:06:10
    > via ge-0/0/4.118
    [Direct/0] 04:06:06
    > via ae0.0
fe80::226:8800:76ef:7584/128
    *[Local/0] 04:06:26
    Local via ge-0/0/4.118
fe80::226:880f:fcef:7580/128
    *[Direct/0] 04:07:25
    > via lo0.0
fe80::226:88ff:feef:7600/128
    *[Local/0] 04:06:42
    Local via ae0.0
```

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.

```
lab@Sun# run show isis spf brief
IPV4 Unicast IS-IS level 1 SPF results:
Node           Metric   Interface   NH   Via       SNPA
Vega.00        25       ae0.0       IPV4 Sirius
f8:c0:1:dd:2:80
                ge-0/0/4.114 IPV4 Arcturus
f8:c0:1:dc:31:84
A-Centauri.00  20       ge-0/0/4.114 IPV4 Arcturus
f8:c0:1:dc:31:84
Rigel.00       15       ae0.0       IPV4 Sirius
f8:c0:1:dd:2:80
Canopus.00     15       ae0.0       IPV4 Sirius
f8:c0:1:dd:2:80
Arcturus.00    10       ge-0/0/4.114 IPV4 Arcturus
f8:c0:1:dc:31:84
Sirius.00      5        ae0.0       IPV4 Sirius
f8:c0:1:dd:2:80
Procyon.00     10       ge-0/0/4.118 IPV4 Procyon
f8:c0:1:dd:24:81
Sun.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
Rigel.00       50       ge-0/0/4.118 IPV6 Procyon
f8:c0:1:dd:24:81
Vega.00        40       ge-0/0/4.118 IPV6 Procyon
f8:c0:1:dd:24:81
A-Centauri.00  35       ge-0/0/4.118 IPV6 Procyon
f8:c0:1:dd:24:81
Arcturus.00    25       ae0.0       IPV6 Sirius
f8:c0:1:dd:2:80
Canopus.00     15       ae0.0       IPV6 Sirius
f8:c0:1:dd:2:80
Sirius.00      5        ae0.0       IPV6 Sirius
f8:c0:1:dd:2:80
Procyon.00     10       ge-0/0/4.118 IPV6 Procyon
f8:c0:1:dd:24:81
Sun.00         0
    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@Sun# 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@Sun# 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. So you can use the same RouterID for OSPFv2 and OSPFv3 if they are configured in the same time.

The neighbors 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.

- **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 any more.
- **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.

a. Configure OSPFv3 on R4.

```
lab@Arcturus# set protocols ospf3 area 0 interface ge-0/0/4.203

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

NOTE: OSPFv3 is configured under the `protocols ospf3` stanza.

```
[edit]
lab@Arcturus# 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: `ospf3` is used in operational commands to display OSPFv3 information. The neighbor-address is the IPv6 link-local address of the neighbor.

```
[edit]
lab@Arcturus# 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
        *[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@Arcturus# 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       0x80000002    1420    0x8a91
40
Router      0.0.0.0          172.30.5.5       0x80000002    1212    0x9683
40
Router      0.0.0.0          172.30.63.1      0x80000004    1216    0x605b
56
Network     0.0.0.2          172.30.63.1      0x80000001    1426    0x54c8
32
Network     0.0.0.3          172.30.63.1      0x80000001    1216    0x58c2
32
IntraArPfx  0.0.0.1          172.30.63.1      0x80000007    1216    0xb20a
52
IntraArPfx  0.0.0.3          172.30.63.1      0x80000001    1426    0x2686
52
IntraArPfx  0.0.0.4          172.30.63.1      0x80000001    1216    0x376b
52

    OSPF3 AS SCOPE link state database
  Type      ID          Adv Rtr          Seq          Age    Cksum
Len
Extern     0.0.0.1          172.30.63.1      0x80000004    510     0x6dad
40
Extern     0.0.0.2          172.30.63.1      0x80000004    391     0x77a1
40
Extern     0.0.0.3          172.30.63.1      0x80000004    271     0x8195
40
Extern     0.0.0.4          172.30.63.1      0x80000004    151     0x8b89
40
Extern     0.0.0.5          172.30.63.1      0x80000004    32      0x957d
40
Extern     0.0.0.6          172.30.63.1      0x80000003    2447    0xa170
40
Extern     0.0.0.7          172.30.63.1      0x80000003    2297    0xab64
40
Extern     0.0.0.8          172.30.63.1      0x80000003    2147    0xb558
40
Extern     0.0.0.9          172.30.63.1      0x80000003    1997    0xbf4c
40
Extern     0.0.0.10         172.30.63.1      0x80000003    1848    0xc940
40
Extern     0.0.0.11         172.30.63.1      0x80000003    1698    0xd334
40
Extern     0.0.0.12         172.30.63.1      0x80000003    1548    0xdd28
40
Extern     0.0.0.13         172.30.63.1      0x80000003    1398    0xe71c
40
Extern     0.0.0.14         172.30.63.1      0x80000003    1248    0xf110
40
```

Extern 40	0.0.0.15	172.30.63.1	0x80000003	1118	0xfb04
Extern 40	0.0.0.16	172.30.63.1	0x80000003	989	0x6f7
OSPF3 Link-Local database, interface ge-0/0/4.203 Area 0.0.0.0					
Type	ID	Adv Rtr	Seq	Age	Cksum
Len					
Link 44	*0.0.0.1	172.30.5.4	0x80000001	1425	0x6cba
Link 64	0.0.0.2	172.30.63.1	0x80000004	630	0xe3fe

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@Arcturus# 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      0x80000007    1403    0xb20a
    52
    Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id
    172.30.63.1
    Prefix-count 1
    Prefix fd17:f0f4:f691:f::1/128
    Prefix-options 0x2, Metric 0
    Aging timer 00:36:36
    Installed 00:23:22 ago, expires in 00:36:37
    Last changed 00:23:22 ago, Change count: 3
```

The intra Area Prefix LSA hold one prefix originated by DC3. This is the IPv6 Loopback address of DC3.

```
[edit]
lab@Arcturus# run show ospf3 database 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    771    0x6dad
    40
    Prefix fd17:f0f4:f691:f::/80
    Prefix-options 0x0, Metric 0, Type 2,
    Aging timer 00:47:09
    Installed 00:12:48 ago, expires in 00:47:09
    Last changed 00:28:06 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@Arcturus# set protocols ospf3 realm ipv4-unicast area 0
interface ge-0/0/4.203
```

```
[edit]
lab@Arcturus# 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@Arcturus# 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
                   > 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@Arcturus# run show ospf3 database realm ipv4-unicast
```

OSPF3 database, Area 0.0.0.0						
Type	ID	Adv Rtr	Seq	Age	Cksum	
Router	*0.0.0.0	172.30.5.4	0x80000002	159	0x8992	
40						
Router	0.0.0.0	172.30.5.5	0x80000002	116	0x9584	
40						
Router	0.0.0.0	172.30.63.1	0x80000004	120	0x5f5c	
56						
Network	0.0.0.2	172.30.63.1	0x80000002	164	0x51ca	
32						
Network	0.0.0.3	172.30.63.1	0x80000002	118	0x55c4	
32						
IntraArPfx	0.0.0.1	172.30.63.1	0x80000008	120	0xc6f8	
40						
IntraArPfx	0.0.0.3	172.30.63.1	0x80000002	164	0xd0ff	
40						
IntraArPfx	0.0.0.4	172.30.63.1	0x80000002	118	0x8145	
40						
OSPF3 AS SCOPE link state database						
Type	ID	Adv Rtr	Seq	Age	Cksum	
Extern	0.0.0.1	172.30.63.1	0x80000004	1194	0xedb4	
32						
Extern	0.0.0.2	172.30.63.1	0x80000004	1044	0xf0af	
32						
Extern	0.0.0.3	172.30.63.1	0x80000004	894	0xf3aa	
32						
Extern	0.0.0.4	172.30.63.1	0x80000004	744	0xf6a5	
32						
Extern	0.0.0.5	172.30.63.1	0x80000004	594	0xf9a0	
32						
Extern	0.0.0.6	172.30.63.1	0x80000004	444	0xfc9b	
32						
Extern	0.0.0.7	172.30.63.1	0x80000004	294	0xff96	
32						
Extern	0.0.0.8	172.30.63.1	0x80000004	144	0x391	
32						
Extern	0.0.0.9	172.30.63.1	0x80000003	2993	0x88b	
32						
Extern	0.0.0.10	172.30.63.1	0x80000003	2844	0xb86	
32						
Extern	0.0.0.11	172.30.63.1	0x80000003	2694	0xe81	
32						
Extern	0.0.0.12	172.30.63.1	0x80000003	2544	0x117c	
32						
Extern	0.0.0.13	172.30.63.1	0x80000003	2394	0x1477	
32						
Extern	0.0.0.14	172.30.63.1	0x80000003	2244	0x1772	
32						
Extern	0.0.0.15	172.30.63.1	0x80000003	2094	0x1a6d	
32						
Extern	0.0.0.16	172.30.63.1	0x80000003	1944	0x1d68	
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	164	0xa1f8
52					
Link	0.0.0.2	172.30.63.1	0x80000004	1344	0x3825
52					

Display the external LSAs carrying the IPv4 prefix.

```
[edit]
lab@Arcturus# 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
```

b. Configure OSPFv3 on R5.

Configure OSPFv3 on R5 the same way as on R4.

```
[edit]
lab@A-Centauri# set protocols ospf3 area 0 interface ge-0/0/4.205

[edit]
lab@A-Centauri# set protocols ospf3 realm ipv4-unicast area 0
interface ge-0/0/4.205

[edit]
lab@A-Centauri# commit
commit complete
```

Verify the operation of OSPFv3 on R5.

```
[edit]
lab@A-Centauri# 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@A-Centauri# 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      0x80000002    2000    0x8a91
40
Router    *0.0.0.0        172.30.5.5      0x80000002    1788    0x9683
40
Router    0.0.0.0        172.30.63.1     0x80000004    1794    0x605b
56
Network   0.0.0.2        172.30.63.1     0x80000001    2004    0x54c8
32
Network   0.0.0.3        172.30.63.1     0x80000001    1794    0x58c2
32
IntraArPfx 0.0.0.1        172.30.63.1     0x80000007    1794    0xb20a
52
IntraArPfx 0.0.0.3        172.30.63.1     0x80000001    2004    0x2686
52
IntraArPfx 0.0.0.4        172.30.63.1     0x80000001    1794    0x376b
52

      OSPF3 AS SCOPE link state database
Type      ID          Adv Rtr          Seq          Age    Cksum
Len
Extern    0.0.0.1        172.30.63.1     0x80000004    1088    0x6dad
40
Extern    0.0.0.2        172.30.63.1     0x80000004    969     0x77a1
40
Extern    0.0.0.3        172.30.63.1     0x80000004    849     0x8195
40
Extern    0.0.0.4        172.30.63.1     0x80000004    730     0x8b89
40
Extern    0.0.0.5        172.30.63.1     0x80000004    610     0x957d
40
Extern    0.0.0.6        172.30.63.1     0x80000004    490     0x9f71
40
Extern    0.0.0.7        172.30.63.1     0x80000004    371     0xa965
40
Extern    0.0.0.8        172.30.63.1     0x80000004    251     0xb359
40
Extern    0.0.0.9        172.30.63.1     0x80000004    132     0xbd4d
40
Extern    0.0.0.10       172.30.63.1     0x80000004    12      0xc741
40
Extern    0.0.0.11       172.30.63.1     0x80000003    2276    0xd334
40
```

Extern 40	0.0.0.12	172.30.63.1	0x80000003	2126	0xdd28
Extern 40	0.0.0.13	172.30.63.1	0x80000003	1976	0xe71c
Extern 40	0.0.0.14	172.30.63.1	0x80000003	1826	0xf110
Extern 40	0.0.0.15	172.30.63.1	0x80000003	1696	0xfb04
Extern 40	0.0.0.16	172.30.63.1	0x80000003	1567	0x6f7
OSPF3 Link-Local database, interface ge-0/0/4.205 Area 0.0.0.0					
Type	ID	Adv Rtr	Seq	Age	Cksum
Len					
Link 44	*0.0.0.1	172.30.5.5	0x80000001	1793	0xd67a
Link 64	0.0.0.3	172.30.63.1	0x80000004	1328	0x7166

R5 receives IPv4 and IPv6 prefixes from DC3.

```
[edit]
lab@A-Centauri# 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
                   *[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
    > 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

```

11) 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@Arcturus# set policy-options policy-statement isis-to-ospf3
term 1 from protocol isis
```

```
[edit]
lab@Arcturus# set policy-options policy-statement isis-to-ospf3
term 1 then tag 123
```

```
[edit]
lab@Arcturus# set policy-options policy-statement isis-to-ospf3
term 1 then accept
```

Apply the policy for both IPv4 and IPv6 routes.

```
[edit]
lab@Arcturus# set protocols ospf3 export isis-to-ospf3
```

```
[edit]
lab@Arcturus# set protocols ospf3 realm ipv4-unicast export isis-
to-ospf3
```

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

Verify the redistribution of IPv6 and IPv4 prefixes to DC3.

```
[edit]
lab@Arcturus# 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   0x80000002   368    0xddbc
48
Extern      *0.0.0.2      172.30.5.4   0x80000002   368    0x422d
48
Extern      *0.0.0.3      172.30.5.4   0x80000002   368    0x290
48
Extern      *0.0.0.21     172.30.5.4   0x80000001   353    0xe2a6
48
Extern      *0.0.0.22     172.30.5.4   0x80000001   339    0xfe7c
48
Extern      *0.0.0.23     172.30.5.4   0x80000001   339    0xf08d
48
Extern      *0.0.0.24     172.30.5.4   0x80000001   331    0x87eb
48

[edit]
lab@Arcturus# 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 36	*0.0.0.14	172.30.5.4	0x80000001	839	0x811f
Extern 36	*0.0.0.15	172.30.5.4	0x80000001	839	0x841a
Extern 36	*0.0.0.16	172.30.5.4	0x80000001	839	0x8715

R4 originates External LSAs for IPv6 and IPv4 prefixes after applying the policy in the export direction.

```
[edit]
lab@A-Centauri# 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@A-Centauri# 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
```

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@A-Centauri# set policy-options policy-statement isis-to-ospf3
term 1 from protocol isis

[edit]
lab@A-Centauri# set policy-options policy-statement isis-to-ospf3
term 1 then tag 123

[edit]
lab@A-Centauri# set policy-options policy-statement isis-to-ospf3
term 1 then accept

[edit]
lab@A-Centauri# set protocols ospf3 export isis-to-ospf3
```



```
[edit]
lab@A-Centauri# set protocols ospf3 realm ipv4-unicast export
isis-to-ospf3

[edit]
lab@A-Centauri# commit
commit complete
```

Verify the redistribution of IPv6 and IPv4 prefixes to DC3.

```
[edit]
lab@A-Centauri# 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    1061    0xbfdb
48
Extern *0.0.0.2          172.30.5.5        0x80000002    945     0xd59
48
Extern *0.0.0.3          172.30.5.5        0x80000001    966     0xc1da
48
Extern *0.0.0.4          172.30.5.5        0x80000001    940     0xc3cb
48
Extern *0.0.0.5          172.30.5.5        0x80000001    926     0xdfa1
48
Extern *0.0.0.6          172.30.5.5        0x80000001    926     0xd1b2
48
Extern *0.0.0.7          172.30.5.5        0x80000001    918     0xef9d
48
```

R5 originates External LSAs to DC3. The IPv6 loopback addresses are redistributed to OSPFv3.

```
[edit]
lab@A-Centauri# run show ospf3 database ospf3 database external
advertising-router 172.30.5.5 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          172.30.5.5        0x80000001    724
0xbfdb 48
Prefix fd17:f0f4:f691:5::4/128
Prefix-options 0x0, Metric 10, Type 2,
Tag 0.0.0.123
Gen timer 00:07:46
Aging timer 00:47:56
Installed 00:12:04 ago, expires in 00:47:56, sent 00:12:02 ago
Last changed 00:12:04 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@Arcturus# 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@Arcturus# 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.

12) 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).

a. Configure OSPFv3 to IS-IS redistribution policy on R4.

Configure and apply the export policy matching all OSPFv3 routes.

```

[edit]
lab@Arcturus# set policy-options policy-statement ospf3-to-isis
term 1 from protocol ospf3

[edit]
lab@Arcturus# set policy-options policy-statement ospf3-to-isis
term 1 then accept

[edit]
lab@Arcturus# set protocols isis export ospf3-to-isis

[edit]
lab@Arcturus# commit
commit complete

```

Verify the redistribution of OSPFv3 IPv4 and IPv6 routes in the IS-IS domain.

```

[edit]
lab@Arcturus# run show isis database level 1 Arcturus.00-00
extensive
IS-IS level 1 link-state database:

Arcturus.00-00 Sequence: 0x11, Checksum: 0x926f, Lifetime: 939
secs

```

IPV4 Unicast IS neighbor: Sun.00	Metric:	10
Two-way fragment: Sun.00-00, Two-way first fragment: Sun.00-00		
IPV4 Unicast IS neighbor: Canopus.00	Metric:	10
Two-way fragment: Canopus.00-00, Two-way first fragment: Canopus.00-00		
IPV4 Unicast IS neighbor: A-Centauri.00	Metric:	10
Two-way fragment: A-Centauri.00-00, Two-way first fragment: A-Centauri.00-00		
IPV6 Unicast IS neighbor: Canopus.00	Metric:	10
Two-way fragment: Canopus.00-00, Two-way first fragment: Canopus.00-00		
IPV6 Unicast IS neighbor: A-Centauri.00	Metric:	10
Two-way fragment: A-Centauri.00-00, Two-way first fragment: A-Centauri.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		
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		

ubleshooting

```

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

```

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@A-Centauri# set policy-options policy-statement ospf3-to-isis
term 1 from protocol ospf3

[edit]
lab@A-Centauri# set policy-options policy-statement ospf3-to-isis
term 1 then accept

[edit]
lab@A-Centauri# set protocols isis export ospf3-to-isis

[edit]
lab@A-Centauri# commit
commit complete

```

Verify the redistribution on R5 from OSPFv3 to IS-IS domain.

```

[edit]
lab@A-Centauri# run show isis database level 1 A-Centauri.00-00
extensive
IS-IS level 1 link-state database:

A-Centauri.00-00 Sequence: 0xc, Checksum: 0x94e1, Lifetime: 940
secs
  IPV4 Unicast IS neighbor: Arcturus.00      Metric:      10
    Two-way fragment: Arcturus.00-00, Two-way first fragment:
Arcturus.00-00
  IPV4 Unicast IS neighbor: Vega.00          Metric:      5
    Two-way fragment: Vega.00-00, Two-way first fragment:
Vega.00-00
  IPV6 Unicast IS neighbor: Arcturus.00      Metric:      10
    Two-way fragment: Arcturus.00-00, Two-way first fragment:
Arcturus.00-00
  IPV6 Unicast IS neighbor: Vega.00          Metric:      5
    Two-way fragment: Vega.00-00, Two-way first fragment:
Vega.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
  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@A-Centauri# 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
                  > to 172.30.0.62 via ge-0/0/4.205

[edit]
lab@A-Centauri# 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 ISIS routes with an internal route preference.

13) 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.

```
[edit]
lab@Arcturus# set protocols ospf3 external-preference 13

[edit]
lab@Arcturus# set protocols ospf3 realm ipv4-unicast external-
preference 13
```

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

[edit]
lab@A-Centauri# set protocols ospf3 external-preference 13

[edit]
lab@A-Centauri# set protocols ospf3 realm ipv4-unicast external-
preference 13

[edit]
lab@A-Centauri# 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@A-Centauri# 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@A-Centauri# run show isis database level 1 A-Centauri.00-00
extensive | match 172.30.132.0/24

IP IPV4 Unicast prefix: 172.30.132.0/24    Metric:      0 Internal
Up
```

14) Configure RIP to OSPFv3 redistribution policy at R4 and R5.

You need to mutual redistribute between RIP and OSPFv3 routes.

a. R4.

Creating policy matching all RIP routes on R4 and apply it in the export direction to OSPFv3.

```
[edit]
lab@Arcturus# set policy-options policy-statement rip-to-ospf3
term 1 from protocol rip
```

```
[edit]
lab@Arcturus# set policy-options policy-statement rip-to-ospf3
term 1 then tag 123
```

```
[edit]
lab@Arcturus# 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@Arcturus# set protocols ospf3 realm ipv4-unicast export rip-
to-ospf3
```

```
[edit]
lab@Arcturus# 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@Arcturus# 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 36	*0.0.0.13	172.30.5.4	0x80000001	839	0x7e24
Extern 36	*0.0.0.14	172.30.5.4	0x80000001	839	0x811f
Extern 36	*0.0.0.15	172.30.5.4	0x80000001	839	0x841a
Extern 36	*0.0.0.16	172.30.5.4	0x80000001	839	0x8715
Extern 36	*0.0.0.17	172.30.5.4	0x80000001	817	0xa8e5
Extern 36	*0.0.0.18	172.30.5.4	0x80000001	817	0xb7c1
Extern 36	*0.0.0.20	172.30.5.4	0x80000001	817	0xe9c3
Extern 36	*0.0.0.21	172.30.5.4	0x80000001	793	0x684b
Extern 36	*0.0.0.22	172.30.5.4	0x80000001	793	0xecb8
Extern 36	*0.0.0.23	172.30.5.4	0x80000001	793	0x931b
Extern 36	*0.0.0.24	172.30.5.4	0x80000002	696	0xf8f
Extern 36	*0.0.0.25	172.30.5.4	0x80000001	783	0x5d38
Extern 36	*0.0.0.26	172.30.5.4	0x80000001	783	0xa104
Extern 36	*0.0.0.27	172.30.5.4	0x80000001	737	0xed9f
Extern 36	*0.0.0.28	172.30.5.4	0x80000002	696	0x8ffe
Extern 36	*0.0.0.29	172.30.5.4	0x80000001	722	0x86f2
Extern 36	*0.0.0.30	172.30.5.4	0x80000001	722	0xb1eb
Extern 36	*0.0.0.49	172.30.5.4	0x80000001	697	0x3225
Extern 36	*0.0.0.50	172.30.5.4	0x80000001	697	0x334b
Extern 3	*0.0.0.51	172.30.5.4	0x80000001	696	0x96f9

b. R5.

The same policy has to be applied to R5 as well for redundancy requirement.

```
[edit]
lab@A-Centauri# set policy-options policy-statement rip-to-ospf3
term 1 from protocol rip
```

```
[edit]
lab@A-Centauri# set policy-options policy-statement rip-to-ospf3
term 1 then tag 123
```

```
[edit]
lab@A-Centauri# set policy-options policy-statement rip-to-ospf3
term 1 then accept
```

```
[edit]
lab@A-Centauri# set protocols ospf3 realm ipv4-unicast export rip-
to-ospf3
```

```
[edit]
lab@A-Centauri# commit
commit complete
```

Verify the advertisement of RIP routes to OSPFv3.

```
[edit]
lab@A-Centauri# 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
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.

15) 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@A-Centauri# run show route 172.30.32.0/24

inet.0: 69 destinations, 136 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.

```
[edit]
lab@Arcturus# set policy options policy-statement ospf3-filter
term 1 from protocol ospf3
```

```
[edit]
lab@Arcturus# set policy options policy-statement ospf3-filter
term 1 from tag 123
```

```
[edit]
lab@Arcturus# 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@Arcturus# set protocols ospf3 import ospf3-filter
```

```
[edit]
lab@Arcturus# set protocols ospf3 realm ipv4-unicast import ospf3-
filter
```

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

Verify again the routing table of R5.

```
[edit]
lab@A-Centauri# 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@A-Centauri# 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
Extern	*0.0.0.1	172.30.5.5	0x80000002	280	0x5266
Extern	*0.0.0.2	172.30.5.5	0x80000002	199	0x5561
Extern	*0.0.0.3	172.30.5.5	0x80000002	118	0x585c
Extern	*0.0.0.4	172.30.5.5	0x80000002	36	0x5b57
Extern	*0.0.0.5	172.30.5.5	0x80000001	670	0x6051
Extern	*0.0.0.6	172.30.5.5	0x80000001	670	0x634c
Extern	*0.0.0.7	172.30.5.5	0x80000001	670	0x6647
Extern	*0.0.0.8	172.30.5.5	0x80000001	670	0x6942
Extern	*0.0.0.9	172.30.5.5	0x80000001	670	0x6c3d
Extern	*0.0.0.10	172.30.5.5	0x80000001	670	0x6f38
Extern	*0.0.0.11	172.30.5.5	0x80000001	670	0x7233
Extern	*0.0.0.12	172.30.5.5	0x80000001	670	0x752e
Extern	*0.0.0.13	172.30.5.5	0x80000001	670	0x7829
Extern	*0.0.0.14	172.30.5.5	0x80000001	670	0x7b24
Extern	*0.0.0.15	172.30.5.5	0x80000001	670	0x7e1f
Extern	*0.0.0.16	172.30.5.5	0x80000001	670	0x811a
Extern	*0.0.0.17	172.30.5.5	0x80000001	656	0xe0cc

Extern 36	*0.0.0.18	172.30.5.5	0x80000001	656	0xb7e4
Extern 36	*0.0.0.19	172.30.5.5	0x80000001	656	0x3748
Extern 36	*0.0.0.20	172.30.5.5	0x80000001	656	0x6d46
Extern 36	*0.0.0.21	172.30.5.5	0x80000001	656	0x7041
Extern 36	*0.0.0.22	172.30.5.5	0x80000001	656	0xc1e9
Extern 36	*0.0.0.23	172.30.5.5	0x80000001	632	0x8a1c
Extern 36	*0.0.0.24	172.30.5.5	0x80000001	632	0xf89
Extern 36	*0.0.0.25	172.30.5.5	0x80000001	632	0xb5eb
Extern 36	*0.0.0.26	172.30.5.5	0x80000003	520	0x1184
Extern 36	*0.0.0.27	172.30.5.5	0x80000002	513	0x4150
Extern 36	*0.0.0.28	172.30.5.5	0x80000001	622	0xc3d4
Extern 36	*0.0.0.29	172.30.5.5	0x80000001	561	0x632f
Extern 36	*0.0.0.30	172.30.5.5	0x80000001	561	0x3a47
Extern 36	*0.0.0.31	172.30.5.5	0x80000001	561	0x6540
Extern 36	*0.0.0.32	172.30.5.5	0x80000002	513	0x98d7
Extern 36	*0.0.0.33	172.30.5.5	0x80000002	513	0x99fd
Extern 36	*0.0.0.34	172.30.5.5	0x80000002	520	0x5743

Filtering for routes with tag 123 in routing table. No routes with tag 123 is matched.

```
[edit]
lab@Arcturus# run show route | match "tag 123"
```

b. Apply the filter policy on R5

Configure the same policy on R5.

```
[edit]
lab@A-Centauri# set policy-options policy-statement ospf3-filter
term 1 from protocol ospf3
```

```
[edit]
lab@A-Centauri# set policy-options policy-statement ospf3-filter
term 1 from tag 123
```

```
[edit]
```



```
lab@A-Centauri# set policy-options policy-statement ospf3-filter
term 1 then reject
```

```
[edit]
```

```
lab@A-Centauri# set protocols ospf3 import ospf3-filter
```

```
[edit]
```

```
lab@A-Centauri# set protocols ospf3 realm ipv4-unicast import
ospf3-filter
```

```
[edit]
```

```
lab@A-Centauri# commit
commit complete
```

Verify the result of the import policy on R5.

```
[edit]
```

```
lab@A-Centauri# run show route | match "tag 123"
```

No routes with tag 123 are present in the routing table.

16) 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@Arcturus# set policy-options policy-statement ospf3-to-rip
term 1 from protocol ospf3
```

```
[edit]
```

```
lab@Arcturus# set policy-options policy-statement ospf3-to-rip
term 1 then tag 123
```

```
[edit]
```

```
lab@Arcturus# 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@Arcturus# set protocols rip group dc2 export ospf3-to-rip
```

```
[edit]
```

```
lab@Arcturus# commit
commit complete
```

Policy is applied in export direction to the RIP group on R4.

Verify the result of the applied policy.

```
[edit]
lab@Arcturus# run show route advertising-protocol rip 172.30.0.49

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@A-Centauri# set policy-options policy-statement ospf3-to-rip
term 1 from protocol ospf3
```

```
[edit]
```

```
lab@A-Centauri# set policy-options policy-statement ospf3-to-rip
term 1 then tag 123
```

```
[edit]
```

```
lab@A-Centauri# set policy-options policy-statement ospf3-to-rip
term 1 then accept
```

```
[edit]
```

```
lab@A-Centauri# set protocols rip group dc2 export ospf3-to-rip
```

```
[edit]
```

```
lab@A-Centauri# commit
commit complete
```

Verify the redistribution of the OSPFv3 routes to RIP on R5.

```
lab@A-Centauri# 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
                   > to 172.30.0.62 via ge-0/0/4.205
```

```
172.30.133.0/24    [IS-IS/15] 00:41:24, metric 10
                  > to 172.30.0.29 via ge-0/0/4.145
                  *[OSPF3/13] 00:21:38, metric 0, tag 0
                  > to 172.30.0.62 via ge-0/0/4.205
172.30.134.0/24    [IS-IS/15] 00:41:24, metric 10
                  > to 172.30.0.29 via ge-0/0/4.145
                  *[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

```

R5 is successfully advertising the OSPFv3 routes to DC3 via RIP.

17) 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@Sun# 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@Sun# 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@Sun# 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** stanza.

```
[edit]
lab@Arcturus# deactivate protocols
```

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

Now when R4 is out of the routing path, you can verify again the path to the same IP addresses.

```
[edit]
lab@Sun# 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
```

Path to RIP route is now from R1, R8, R5 and ends at DC2.

```
[edit]
lab@Sun# 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@Sun# 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.

Don't forget to activate the **protocols** stanza on R4.

Chapter Three solutions: 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: IBGP and Confederation

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 other 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.
- 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 an 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 message 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.
- *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 a 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.
- *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** knob the algorithm stops before step 7 and all remaining routes are used for forwarding.

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 were 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.

9)

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.

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.

```
[edit]
lab@Sun# set routing-options autonomous-system 65000

[edit]
lab@Sun# set routing-options confederation 54591 members [ 65000
65001 65002 65003 ]

[edit]
lab@Sun# 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.

```
[edit]
lab@Canopus# set routing-options autonomous-system 65001

[edit]
lab@Canopus# set routing-options confederation 54591 members [
65000 65001 65002 65003 ]

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

NOTE: Configuration on R4 is the same, as it participates in the same sub-AS.

c. R5.

```
[edit]
lab@A-Centauri# set routing-options autonomous-system 65002

[edit]
lab@A-Centauri# set routing-options confederation 54591 members [
65000 65001 65002 65003 ]

[edit]
lab@A-Centauri# commit
commit complete
```

NOTE: Configuration on R6 is the same, as it participates in the same sub-AS.

d. R7.

```
[edit]
lab@Rigel# set routing-options autonomous-system 65003

[edit]
lab@Rigel# set routing-options confederation 54591 members [ 65000
65001 65002 65003 ]

[edit]
lab@Rigel# 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@Sun# set protocols bgp group ibgp type internal

[edit]
lab@Sun# set protocols bgp group ibgp local-address 172.30.5.1

[edit]
lab@Sun# set protocols bgp group ibgp neighbor 172.30.5.2

[edit]
lab@Sun# set protocols bgp group ibgp authentication-key workbook

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

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@Sun# set protocols bgp group cbgp type external

[edit]
lab@Sun# set protocols bgp group cbgp local-address 172.30.5.1

[edit]
lab@Sun# set protocols bgp group cbgp peer-as 65003
```

```
[edit]
lab@Sun# set protocols bgp group cbgp neighbor 172.30.5.8

[edit]
lab@Sun# set protocols bgp group cbgp multihop

[edit]
lab@Sun# set protocols bgp group cbgp authentication-key workbook

[edit]
lab@Sun# 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.

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@Sirius# set protocols bgp group ibgp type internal

[edit]
lab@Sirius# set protocols bgp group ibgp local-address 172.30.5.2

[edit]
lab@Sirius# set protocols bgp group ibgp neighbor 172.30.5.1

[edit]
lab@Sirius# set protocols bgp group ibgp authentication-key
workbook

[edit]
lab@Sirius# set protocols bgp group cbgp type external

[edit]
lab@Sirius# set protocols bgp group cbgp local-address 172.30.5.2

[edit]
lab@Sirius# set protocols bgp group cbgp peer-as 65001

[edit]
lab@Sirius# set protocols bgp group cbgp neighbor 172.30.5.3

[edit]
```

```
lab@Sirius# set protocols bgp group cbgp authentication-key
workbook
```

```
[edit]
```

```
lab@Sirius# set protocols bgp group cbgp multihop
```

```
[edit]
```

```
lab@Sirius# commit
commit complete
```

The configuration on R2 is similar to R1.

c. R8.

R8 is in sub-AS 65003 and has an IBGP neighborhood with R7 and also an CBGP neighborhood with R1 in sub-AS 65000.

```
[edit]
```

```
lab@Procyon# set protocols bgp group ibgp type internal
```

```
[edit]
```

```
lab@Procyon# set protocols bgp group ibgp local-address 172.30.5.8
```

```
[edit]
```

```
lab@Procyon # set protocols bgp group ibgp neighbor 172.30.5.7
```

```
[edit]
```

```
lab@Procyon # set protocols bgp group ibgp authentication-key
workbook
```

```
[edit]
```

```
lab@Procyon # set protocols bgp group cbgp type external
```

```
[edit]
```

```
lab@Procyon # set protocols bgp group cbgp local-address
172.30.5.8
```

```
[edit]
```

```
lab@Procyon # set protocols bgp group cbgp neighbor 172.30.5.1
```

```
[edit]
```

```
lab@Procyon # set protocols bgp group cbgp peer-as 65000
```

```
[edit]
```

```
lab@Procyon # set protocols bgp group cbgp authentication-key
workbook
```

```
[edit]
```

```
lab@Procyon # set protocols bgp group cbgp multihop
```

```
[edit]
```

```
lab@Procyon # commit
commit complete
```


NOTE: The configuration for IBGP and CBGP on the other routers in this network is similar to the config on R1, R2 and R8.

Verify the IBGP and CBGP sessions on the selected routers.

```
[edit]
lab@Sun# 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          0
Peer      AS    InPkt  OutPkt OutQ Flaps    Last Up/Dwn
State|#Active/Received/Accepted/Damped...
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@Sun# 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          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@Sun# run show bgp neighbor 172.30.5.2 | no-more
Peer: 172.30.5.2+179 AS 65000 Local: 172.30.5.1+56672 AS 65000
  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 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
  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 65000)
  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 20 Sent 19 Checked 75
  Input messages: Total 4 Updates 1 Refreshes 0 Octets
80
  Output messages: Total 5 Updates 0 Refreshes 0 Octets
158
  Output Queue[0]: 0
```

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@Sirius# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table
inet.0
Tot Paths  Act Paths Suppressed  History Damp State  Pending
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
```

```
[edit]
lab@Procyon# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table
inet.0
Tot Paths  Act Paths Suppressed  History Damp State  Pending
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 Two: EBGW Configuration.

The second task requires the configuration of EBGW sessions with the Internet Exchanges (IX), Peers (P) and Customer (C) BGP peers.

1) Configure EBGW on R1.

According to Figure 7, R1 and R2 have to build EBGW 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@Sun# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@Sun# set interfaces ge-0/0/5.300 vlan-id 300 family inet
address 192.168.1.1/24

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

Now you can configure the EBGW peer over the new interface.

```
[edit]
lab@Sun# set protocols bgp group IX peer-as 1620

[edit]
lab@Sun# set protocols bgp group IX neighbor 192.168.1.3

[edit]
lab@Sun# set protocols bgp group IX neighbor 192.168.1.4

[edit]
lab@Sun# set protocols bgp group IX log-updown

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

Group IX includes both EBGW peers, IX-1 and IX-2 because they share the same parameters. **log-update** accomplishes the requirement that EBGW state changes must be logged as syslog message.

```
[edit]
lab@Sun# 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@Sirius# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@Sirius# set interfaces ge-0/0/5.300 vlan-id 300 family inet
address 192.168.1.2/24

[edit]
lab@Sirius# set protocols bgp group IX peer-as 1620

[edit]
lab@Sirius# set protocols bgp group IX neighbor 192.168.1.3

[edit]
lab@Sirius# set protocols bgp group IX neighbor 192.168.1.4

[edit]
lab@Sirius# set protocols bgp group IX log-updown

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

Display the BGP summary in R2.

```
[edit]
lab@Sirius# 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@Sun# 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 -> {192.168.1.3}
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
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@Sirius# 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:
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
        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 EBGW 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@Canopus# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@Canopus# set interfaces ge-0/0/5.301 vlan-id 301 family inet
address 192.168.0.1/30

[edit]
lab@Canopus# set interfaces ge-0/0/5.302 vlan-id 302 family inet
address 192.168.0.5/30

[edit]
lab@Canopus# commit
commit complete

```

Now the EBGW neighbors can be configured. Two separate BGP groups are used for both peers.

```

[edit]
lab@Canopus# set protocols bgp group P2-1 peer-as 53732.2005

[edit]
lab@Canopus# set protocols bgp group P2-1 neighbor 192.168.0.2

[edit]
lab@Canopus# set protocols bgp group P2-1 log-updown

[edit]
lab@Canopus# set protocols bgp group P3-1 peer-as 43208.365

```

```
[edit]
lab@Canopus# set protocols bgp group P3-1 neighbor 192.168.0.6

[edit]
lab@Canopus# set protocols bgp group P3-1 log-updown

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

Verify the BGP session table for the new EBGP neighbors.

```
[edit]
lab@Canopus# 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.

```
[edit]
lab@A-Centauri# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@A-Centauri# set interfaces ge-0/0/5.303 vlan-id 303 family
inet address 192.168.0.9/30

[edit]
lab@A-Centauri# set interfaces ge-0/0/5.304 vlan-id 304 family
inet address 192.168.0.13/30

[edit]
lab@A-Centauri# 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@A-Centauri# set protocols bgp group C3 peer-as 64514

[edit]
lab@A-Centauri# set protocols bgp group C3 neighbor 192.168.0.10

[edit]
lab@A-Centauri# set protocols bgp group C3 neighbor 192.168.0.14

[edit]
lab@A-Centauri# set protocols bgp group C3 multipath

[edit]
lab@A-Centauri# set protocols bgp group C3 log-updown

[edit]
lab@A-Centauri# 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@A-Centauri# 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...
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@A-Centauri# 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@A-Centauri# 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    0 192.168.0.14 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 have to alter the default behavior. This is accomplished with a load-balancing policy with a special terminating action.

```

[edit]
lab@A-Centauri# set policy-options policy-statement load_balance
term 1 then load-balance per-packet

```

```

[edit]
lab@A-Centauri# set routing-options forwarding-table export
load_balance

```

```

[edit]
lab@A-Centauri# 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@A-Centauri# 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      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
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@Vega# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@Vega# set interfaces ge-0/0/0.305 vlan-id 305 family inet
address 192.168.0.17/30

[edit]
lab@Vega# set interfaces ge-0/0/5.306 vlan-id 306 family inet
address 192.168.0.21/30

[edit]
lab@Vega# set interfaces ge-0/0/5.307 vlan-id 307 family inet
address 192.168.0.25/30

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

Configure the RIP protocol to learn the Loopback IP address.

```
[edit]
lab@Vega# set protocols rip group peer neighbor ge-0/0/5.306
```

```
[edit]
lab@Vega# 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.

If you remember from the previous chapter, JUNOS only accepts routes from RIP neighbors, but does not advertise anything by default. So you have to configure and apply export policy to export R6 Loopback IP address to both neighbors.

```
[edit]
lab@Vega# set policy-options policy-statement loopback-to-rip term
1 from protocol direct
```

```
[edit]
lab@Vega# set policy-options policy-statement loopback-to-rip term
1 from route-filter 172.30.5.6/32 exact
```

```
[edit]
lab@Vega# set policy-options policy-statement loopback-to-rip term
1 then accept
```

```
[edit]
lab@Vega# set protocols rip group peer export loopback-to-rip
```

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

Verify if the Loopback IP address of C2-1 is received by R6.

```
[edit]
lab@Vega# run show rip neighbor
```

In Neighbor Met	Local State	Source Address	Destination Address	Send Mode	Receive Mode
ge-0/0/5.306 1	Up	192.168.0.21	224.0.0.9	mcast	both
ge-0/0/5.307 1	Up	192.168.0.25	224.0.0.9	mcast	both

```
[edit]
lab@Vega# 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)

```

C2-1 neighbor is reachable through both interfaces and the Loopback IP address is present in the routing table.

```

[edit]
lab@Vega# 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@Vega# 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@Vega# set protocols bgp group C2-1 local-address 172.30.5.6

[edit]
lab@Vega# set protocols bgp group C2-1 neighbor 172.31.31.1

[edit]
lab@Vega# set protocols bgp group C2-1 peer-as 64513

[edit]
lab@Vega# set protocols bgp group C2-1 multihop

[edit]
lab@Vega# set protocols bgp group C2-1 log-updown

[edit]
lab@Vega# 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.

```
[edit]
lab@Vega# 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 messages.

```
[edit]
lab@Vega# set protocols bgp group C1-1 neighbor 192.168.0.18

[edit]
lab@Vega# set protocols bgp group C1-1 peer-as 64512

[edit]
lab@Vega# set protocols bgp group C1-1 family inet unicast prefix-
limit maximum 20

[edit]
lab@Vega# set protocols bgp group C1-1 family inet unicast prefix-
limit teardown idle-timeout 3

[edit]
lab@Vega# set protocols bgp group C1-1 log-updown

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

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@Vega# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table
inet.0
Tot Paths  Act Paths  Suppressed  History Damp State  Pending
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.

Again first the interfaces should be configured for IP connectivity.

```
[edit]
lab@Rigel# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@Rigel# set interfaces ge-0/0/5.308 vlan-id 308 family inet
address 192.168.0.29/30

[edit]
lab@Rigel# set interfaces ge-0/0/5.309 vlan-id 309 family inet
address 192.168.0.33/30

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

Configure the EBGP neighbor with C1-1.

```
[edit]
lab@Rigel# set protocols bgp group C1-1 peer-as 64512

[edit]
lab@Rigel# set protocols bgp group C1-1 neighbor 192.168.0.34

[edit]
lab@Rigel# set protocols bgp group C1-1 family inet unicast
prefix-limit maximum 20

[edit]
lab@Rigel# set protocols bgp group C1-1 family inet unicast
prefix-limit teardown idle-timeout 3
```

```
[edit]
lab@Rigel# set protocols bgp group C1-1 log-updown
```

Configure the EBGP neighbor with P1-2.

```
[edit]
lab@Rigel# set protocols bgp group P1-2 peer-as 1679.12483

[edit]
lab@Rigel# set protocols bgp group P1-2 neighbor 192.168.0.30

[edit]
lab@Rigel# set protocols bgp group P1-2 log-updown

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

Verify the BGP session table for the EBGP neighbors.

```
[edit]
lab@Rigel# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.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
```

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@Procyon# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@Procyon# 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@Procyon# set protocols bgp group P1-1 peer-as 1679.12483

[edit]
lab@Procyon# set protocols bgp group P1-1 neighbor 192.168.0.38

[edit]
lab@Procyon# set protocols bgp group P1-1 log-updown

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

Verify the BGP session table.

```
[edit]
lab@Procyon# 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
```

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.

```
[edit]
lab@Canopus# set interfaces ge-0/0/5.301 family inet6

[edit]
lab@Canopus# 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@Canopus# set protocols bgp traceoptions file bgp.log

[edit]
lab@Canopus# set protocols bgp traceoptions flag packets detail

[edit]
lab@Canopus# 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@Canopus# run show log bgp.log | match bgp_listen
```

```
Aug 29 20:15:28.738997 bgp_listen_accept: Connection attempt from
unconfigured neighbor: fe80::223:9c01:2d8b:6c81+56998
```

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@Canopus# set protocols bgp group P2-1-ipv6 peer-as 53732.2005

[edit]
lab@Canopus# set protocols bgp group P2-1-ipv6 local-interface ge-
0/0/5.301

[edit]
lab@Canopus# set protocols bgp group P2-1-ipv6 neighbor
fe80::223:9c01:2d8b:6c81

[edit]
lab@Canopus# set protocols bgp group P2-1-ipv6 log-updown

[edit]
lab@Canopus# 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@Canopus# 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
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 particular BGP group.

```

[edit]
lab@Canopus# run show route receive-protocol bgp
fe80::223:9c01:2d8b:6c81

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, 50 routes (48 active, 0 holddown, 0
hidden)
  Prefix                Nexthop          MED      Lclpref    AS
path
  fd01:aaaa:bbbb:1::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:1::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:2::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:3::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:4::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:5::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:6::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:7::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:8::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:9::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:a::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:b::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:c::/80
*                fe80::223:9c01:2d8b:6c81
3521382357 I
  fd01:aaaa:bbbb:1:d::/80

```

```

*                               fe80::223:9c01:2d8b:6c81
3521382357 I
    fd01:aaaa:bbbb:1:e::/80
*                               fe80::223:9c01:2d8b:6c81
3521382357 I
    fd01:aaaa:bbbb:1:f::/80
*                               fe80::223:9c01:2d8b:6c81
3521382357 I

```

16 IPv6 prefixes are received from R3.

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@A-Centauri# set interfaces ge-0/0/5.303 family inet6 address
::192.168.0.9/126

```

```

[edit]
lab@A-Centauri# set interfaces ge-0/0/5.304 family inet6 address
::192.168.0.13/126

```

```

[edit]
lab@A-Centauri# commit
commit complete

```

Now you have to enable the appropriate family for the existing IPv4 session.

```

[edit]
lab@A-Centauri# set protocols bgp group C3 family inet6 unicast

[edit]
lab@A-Centauri# commit
commit complete

```

Verify the BGP session table.

```

[edit]
lab@A-Centauri# run show bgp summary
Groups: 3 Peers: 4 Down peers: 0
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@A-Centauri# 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]
lab@A-Centauri# 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. So you have explicitly enabled it again to have both families active over a single BGP session.

```

[edit]
lab@A-Centauri# set protocols bgp group C3 family inet unicast

[edit]
lab@A-Centauri# commit
commit complete

```

Verify if the IPv4 routes are now received together with IPv6 prefixes.

```

[edit]
lab@A-Centauri# 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

```

<http://www.inetzero.com> - Copyright 2017 iNET ZERO. All rights reserved

For personal non commercial use only – do not distribute

<https://t.me/learningnets>

```
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@Rigel# set interfaces ge-0/0/5.308 family inet6 address
fc09:c0:ffee::1/126

[edit]
lab@Rigel# set protocols bgp group P1-2-ipv6 peer-as 1679.12483

[edit]
lab@Rigel# set protocols bgp group P1-2-ipv6 neighbor
fc09:c0:ffee::2

[edit]
lab@Rigel# set protocols bgp group P1-2-ipv6 log-updown

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

After committing the configuration, verify the BGP summary table.

```
[edit]
lab@Rigel# 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@Procyon# set interfaces ge-0/0/5.310 family inet6 address
fc09:c0:ffee::5/126

[edit]
lab@Procyon# set protocols bgp group P1-1-ipv6 peer-as 1679.12483

[edit]
lab@Procyon# set protocols bgp group P1-1-ipv6 neighbor
fc09:c0:ffee::6

[edit]
lab@Procyon# set protocols bgp group P1-1-ipv6 log-updown

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

Verify the BGP session table.

```
[edit]
lab@Procyon# 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 period of time. This number is called figure of merit and by default is 0. Each time the prefix is withdraw the figure is increased with 1000, etc, etc

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.
- **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 is 60 minutes.

a. R3.

On R3, for both C3 customer peers you have to enable default damping parameters.

```
[edit]
lab@A-Centauri# set protocols bgp group C3 damping

[edit]
lab@A-Centauri# 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@Vega# set protocols bgp group C1-1 damping
```

damping 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@Vega# set policy-options damping aggressive half-life 20

[edit]
lab@Vega# set policy-options damping aggressive reuse 500

[edit]
lab@Vega# set policy-options damping aggressive suppress 2500

[edit]
lab@Vega# 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.

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@Vega# set protocols bgp group C1-1 import damp-aggressive

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

Policy is applied to the C1-1 peer only.

c. R7.

R7 has session with C1-1 as well. So the same damping and routing policies have to be configured.

```
[edit]
lab@Rigel# set protocols bgp group C1-1 damping

[edit]
lab@Rigel# set policy-options damping aggressive half-life 20

[edit]
lab@Rigel# set policy-options damping aggressive reuse 500

[edit]
lab@Rigel# set policy-options damping aggressive suppress 2500

[edit]
lab@Rigel# set policy-options policy-statement damp-aggressive
term 1 then damping aggressive

[edit]
lab@Rigel# set protocols bgp group C1-1 import damp-aggressive

[edit]
lab@Rigel# 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.

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 options 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.

```
[edit]
lab@Sun# 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.

```
[edit]
lab@Sun# 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@Sun# 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@Sun# set protocols isis interface ge-0/0/5.300 level 1 passive

[edit]
lab@Sun# commit
commit complete

```

Verify the IS-IS interface table.

```

[edit]
lab@Sun# 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   0    0x1 Passive              Disabled
10/10
lo0.0          0    0x1 Passive                Passive
0/0
lo0.32768      0    0x1 Passive                Passive
0/0

```

Verify the Level 1 LSP interface route is advertised to IS-IS.

```
[edit]
lab@Sun# run show isis database level 1 Sun.00-00 extensive
IS-IS level 1 link-state database:

Sun.00-00 Sequence: 0x2f, Checksum: 0xf4d7, Lifetime: 1154 secs
  IPV4 Unicast IS neighbor: Sirius.00      Metric:      5
    Two-way fragment: Sirius.00-00, Two-way first fragment:
Sirius.00-00
  IPV4 Unicast IS neighbor: Arcturus.00    Metric:      10
    Two-way fragment: Arcturus.00-00, Two-way first fragment:
Arcturus.00-00
  IPV4 Unicast IS neighbor: Procyon.00     Metric:      10
    Two-way fragment: Procyon.00-00, Two-way first fragment:
Procyon.00-00
  IPV6 Unicast IS neighbor: Sirius.00      Metric:      5
    Two-way fragment: Sirius.00-00, Two-way first fragment:
Sirius.00-00
  IPV6 Unicast IS neighbor: Procyon.00     Metric:      10
    Two-way fragment: Procyon.00-00, Two-way first fragment:
Procyon.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@Sirius# set protocols isis interface ge-0/0/5.300 level 1
passive
```

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

c. R3.

R3 has two EBGP sessions with P2-1 and P3-1 and receives combined more than 200 prefixes.

```
[edit]
lab@Canopus# 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@Canopus# run show route advertising-protocol bgp 172.30.5.4

inet.0: 878 destinations, 919 routes (298 active, 0 holddown, 619 hidden)
  Prefix                Nexthop          MED  Lclpref  AS path
* 0.0.0.0/5             192.168.0.2      100           3521382357 34427 7478 37839 39648
I
* 5.127.0.0/17          192.168.0.6      100           2831679853 9726 36659 30705 25538
37414 49276 ?
* 6.98.154.0/25         192.168.0.6      100           2831679853 26697 36659 39905 I
* 8.0.0.0/5             192.168.0.6      100           2831679853 9726 36659 42498 I
* 10.128.0.0/11         192.168.0.6      100           2831679853 26697 4341 43012 28104
39181 51157 ?
* 12.173.96.0/19        192.168.0.2      100           3521382357 4637 7478 16572 33786
?
* 15.36.0.0/15          192.168.0.2      100           3521382357 4637 57359 16881 43174
64323 I
* 17.96.0.0/13          192.168.0.2      100           3521382357 4637 22896 I
* 17.205.88.0/21        192.168.0.6      100           2831679853 9726 19412 16490 6122
I
* 18.100.0.0/17         192.168.0.6      100           2831679853 9726 19412 40487 63999
13005 28592 5369 9061 I
* 20.59.93.0/26         192.168.0.2      100           3521382357 46848 24130 9166 22927
23834 20922 27823 I
* 21.0.0.0/8            192.168.0.6      100           2831679853 26697 56251 10277
21151 I
```

You can see that R3 is advertising the routes to R4 with their original next-hop.

```
[edit]
lab@Canopus# set policy-options policy-statement nhs term 1 from
protocol bgp
```

```
[edit]
lab@Canopus# set policy-options policy-statement nhs term 1 from
route-type external
```

```
[edit]
```

```
lab@Canopus# set policy-options policy-statement nhs term 1 then
next-hop self
```

```
[edit]
```

```
lab@Canopus# set protocols bgp group ibgp export nhs
```

```
[edit]
```

```
lab@Canopus# set protocols bgp group cbgp export nhs
```

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@Canopus# run show route advertising-protocol bgp 172.30.5.4

inet.0: 878 destinations, 919 routes (298 active, 0 holddown, 619 hidden)
  Prefix            Nexthop        MED  Lclpref  AS path
* 0.0.0.0/5         Self           100          3521382357 34427 7478 37839 39648
I
* 5.127.0.0/17      Self           100          2831679853 9726 36659 30705 25538
37414 49276 ?
* 6.98.154.0/25     Self           100          2831679853 26697 36659 39905 I
* 8.0.0.0/5         Self           100          2831679853 9726 36659 42498 I
* 10.128.0.0/11     Self           100          2831679853 26697 4341 43012 28104
39181 51157 ?
* 12.173.96.0/19    Self           100          3521382357 4637 7478 16572 33786
?
* 15.36.0.0/15      Self           100          3521382357 4637 57359 16881 43174
64323 I
* 17.96.0.0/13      Self           100          3521382357 4637 22896 I
* 17.205.88.0/21    Self           100          2831679853 9726 19412 16490 6122
I
* 18.100.0.0/17     Self           100          2831679853 9726 19412 40487 63999
13005 28592 5369 9061 I
* 20.59.93.0/26     Self           100          3521382357 46848 24130 9166 22927
23834 20922 27823 I
* 21.0.0.0/8        Self           100          2831679853 26697 56251 10277
21151 I
```

d. R4.

R4 has no EBGP sessions, but it is a good candidate to explore the next-hop issue.

Looking at the BGP session table, you can see that R4 is receiving prefixes from R3 and R5.

```
[edit]
lab@Arcturus# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table
inet.0
      Tot Paths  Act Paths Suppressed    History Damp State    Pending
      258        258          0          0          0          0
Peer      AS InPkt  OutPkt  OutQ  Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.5.3 65001   875    464    0    0   3:23:20 211/211/211/0
0/0/0/0
172.30.5.5 65002   905    888    0    0   3:21:38 47/47/47/0
0/0/0/0

[edit]
lab@Arcturus# run show route terse

inet.0: 329 destinations, 329 routes (329 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        A 130                    Reject
* 0.0.0.0/5        B 170      100          to Reject     3521382357 34427 7478
37839 39648 I
* 5.127.0.0/17     B 170      100          to Reject     2831679853 9726 36659
30705 25538 37414 49276 ?
* 6.98.154.0/25    B 170      100          to Reject     2831679853 26697 36659
39905 I
* 8.0.0.0/5        B 170      100          to Reject     2831679853 9726 36659
42498 I
* 10.10.1.0/24     D 0                    >ge-0/0/0.0
* 10.10.1.4/32     L 0                    Local
* 10.10.10.0/24    S 5                    >10.10.1.254
* 10.128.0.0/11    B 170      100          to Reject     2831679853 26697 4341
43012 28104 39181 51157 ?
* 12.173.96.0/19  B 170      100          to Reject     3521382357 4637 7478
16572 33786 ?
* 15.36.0.0/15     B 170      100          to Reject     3521382357 4637 57359
16881 43174 64323 I
```

NOTE: This output is taken before applying next-hop self policy on R3.

You can see that R4 has no hidden prefixes, but all prefixes are resolved through a next-hop of **Reject**. This is because of the Aggregated default route configured on R4 in Chapter Two.

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@Arcturus# run show route terse

inet.0: 329 destinations, 329 routes (329 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        A 130                Reject
* 0.0.0.0/5        B 170      100      >172.30.0.21  3521382357 34427 7478
37839 39648 I
* 5.127.0.0/17     B 170      100      >172.30.0.21  2831679853 9726 36659
30705 25538 37414 49276 ?
* 6.98.154.0/25    B 170      100      >172.30.0.21  2831679853 26697
36659 39905 I
* 8.0.0.0/5        B 170      100      >172.30.0.21  2831679853 9726 36659
42498 I
* 10.10.1.0/24     D 0
* 10.10.1.4/32     L 0
* 10.10.10.0/24    S 5
* 10.128.0.0/11    B 170      100      >172.30.0.21  2831679853 26697 4341
43012 28104 39181 51157 ?
* 12.173.96.0/19   B 170      100      >172.30.0.21  3521382357 4637 7478
16572 33786 ?
* 15.36.0.0/15     B 170      100      >172.30.0.21  3521382357 4637 57359
16881 43174 64323 I
```

e. R5.

R5 has two neighbors, both from AS 64514. Verify BGP session table on R5.

```
[edit]
lab@A-Centauri# 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
```

Only one route from R6 is not active.

```
[edit]
lab@A-Centauri# run show route hidden extensive

inet.0: 332 destinations, 348 routes (331 active, 0 holddown, 1
hidden)
172.31.31.0/24 (1 entry, 0 announced)
    BGP      Preference: 170/-101
              Next hop type: Unusable
              Address: 0x113bc8c
              Next-hop reference count: 1
              State: <Hidden Int Ext>
              Local AS: 65002 Peer AS: 65002
              Age: 2:59:32
              Task: BGP_65002.172.30.5.6+62081
              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: 0 -
```

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.

Apply the next-hop self policy on R5 IBGP and CBGP session.

```
[edit]
lab@A-Centauri# set policy-options policy-statement nhs term 1
from protocol bgp

[edit]
lab@A-Centauri# set policy-options policy-statement nhs term 1
from route-type external

[edit]
lab@A-Centauri# set policy-options policy-statement nhs term 1
then next-hop self

[edit]
lab@A-Centauri# set protocols bgp group ibgp export nhs

[edit]
lab@A-Centauri# set protocols bgp group cbgp export nhs

[edit]
lab@A-Centauri# commit
commit complete
```

f. R6.

R6 has two neighbors. It has a lot of hidden prefixes.

```
[edit]
lab@Vega# run show route hidden terse

inet.0: 881 destinations, 905 routes (316 active, 0 holddown, 589 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf    Metric 1    Metric 2    Next hop      AS path
0.0.0.0/4          B 170      100          Unusable      (65003 65000) 1620
61671 I
0.0.0.0/5          B 170      100          Unusable      (65003 65000) 1620
61671 27075 I
1.64.0.0/10        B 170      100          Unusable      (65003 65000) 1620
61671 I
1.84.160.0/20      B 170      100          Unusable      (65003 65000) 1620
33112 I
1.96.0.0/11        B 170      100          Unusable      (65003 65000) 1620
33112 63164 40776 51777 I
1.161.192.0/21     B 170      100          Unusable      (65003 65000) 1620
33112 30404 32138 45045 I
1.166.89.0/24      B 170      100          Unusable      (65003) 110047427
56422 47123 34908 57085 3090 19128 21300 I
1.176.0.0/12       B 170      100          Unusable      (65003 65000) 1620
33112 49129 16320 52954 I
```

NOTE: The AS path of the prefixes contains AS numbers in brackets. Those are sub-AS added crossing CBGP sessions. All the prefixes are coming from sub-AS 65000. The Sub AS-numbers are removed when the prefixes are advertised to external peers.

After adding the interface on R1 and R2 as passive to IS-IS, the next-hops are resolved.

Adding the next-hop self policy on R6.

```
[edit]
lab@Vega# set policy-options policy-statement nhs term 1 from
protocol bgp

[edit]
lab@Vega# set policy-options policy-statement nhs term 1 from
route-type external

[edit]
lab@Vega# set policy-options policy-statement nhs term 1 then
next-hop self

lab@Vega# set protocols bgp group ibgp export nhs

[edit]
lab@Vega# set protocols bgp group cbgp export nhs

[edit]
lab@Vega# commit
commit complete

g. R7.
```

R7 has EBGP session to one customer and one peer router in different ASs.


```
[edit]
lab@Rigel# run show bgp summary
Groups: 5 Peers: 5 Down peers: 1
Table
  Tot Paths  Act Paths Suppressed    History  Damp State    Pending
inet.0
      612      595          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      6      501      0      0      13 15/32/32/0
0/0/0/0
172.30.5.8      65003      0      0      0      0      1:28 Active
192.168.0.30  110047427    495      7      0      0      6 564/564/564/0
0/0/0/0
192.168.0.34      64512      4      501      0      0      10 16/16/16/0
0/0/0/0
fc09:c0:ffee::2 110047427      3      3      0      0          2 Establ
inet6.0: 16/16/16/0
```

Apply the next-hop self policy on R7 IBGP and CBGP session.

```
[edit]
lab@Rigel# set policy-options policy-statement nhs term 1 from
protocol bgp
```

```
[edit]
lab@Rigel# set policy-options policy-statement nhs term 1 from
route-type external
```

```
[edit]
lab@Rigel# set policy-options policy-statement nhs term 1 then
next-hop self
```

```
[edit]
lab@Rigel# set protocols bgp group ibgp export nhs
```

```
[edit]
lab@Rigel# set protocols bgp group cbgp export nhs
```

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

h. R8.

R8 has single EBGP session with P1-1 router.

```
[edit]
lab@Procyon# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 2028 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...
172.30.5.1 65000 1928 1207 0 0 3:30:53 655/657/657/0
0/0/0/0
172.30.5.7 65003 1596 1001 0 1 41:44 45/807/807/0
0/0/0/0
192.168.0.38 110047427 946 555 0 0 3:31:04 108/564/564/0
0/0/0/0
fc09:c0:ffee::6 110047427 142 145 0 0 1:04:51 Establ
inet6.0: 16/16/16/0
```

Apply the next-hop self policy on R8 IBGP and CBGP sessions.

```
[edit]
lab@Procyon# set policy-options policy-statement nhs from protocol
bgp
```

```
[edit]
lab@Procyon# set policy-options policy-statement nhs from route-
type external
```

```
[edit]
lab@Procyon# set policy-options policy-options policy-statement
nhs then next-hop self
```

```
[edit]
lab@Procyon# set protocols bgp group ibgp export nhs
```

```
[edit]
lab@Procyon# set protocols bgp group cbgp export nhs
```

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

11) Verify the routing table for hidden routes.

a. R1.

```
[edit]
lab@Sun# run show route hidden terse

inet.0: 827 destinations, 1955 routes (827 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.

```
[edit]
lab@Sirius# run show route hidden

inet.0: 827 destinations, 1963 routes (827 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 on R2 as well.

c. R3.

```
[edit]
lab@Canopus# 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, 50 routes (48 active, 0 holddown, 0 hidden)
```

d. R4.

```
[edit]
lab@Arcturus# run show route hidden terse

inet.0: 329 destinations, 329 routes (329 active, 0 holddown, 0 hidden)
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
inet6.0: 33 destinations, 35 routes (33 active, 0 holddown, 0 hidden)
```

e. R5.

```
[edit]
lab@A-Centauri# 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]
lab@Vega# 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, 32 routes (31 active, 0 holddown, 0 hidden)
```

g. R7.

```
[edit]
lab@Rigel# run show route hidden

inet.0: 878 destinations, 2114 routes (878 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)
```

h. R8.

```
[edit]
lab@Procyon# run show route hidden

inet.0: 876 destinations, 2097 routes (876 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)
```

All routers do not have hidden routes anymore after applying the next-hop self policies and adding the external interface addresses as ISIS 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.

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.

a. R1.

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@Sun# set policy-options community C1 members 54591:64512

[edit]
lab@Sun# set policy-options community C2 members 54591:64513

[edit]
lab@Sun# set policy-options community C3 members 54591:64514

[edit]
lab@Sun# set policy-options community IX members 54591:1620

[edit]
lab@Sun# set policy-options community P1 members 54591:1679

[edit]
lab@Sun# set policy-options community P2 members 54591:53732

[edit]
lab@Sun# set policy-options community P3 members 54591:43208
```

```
[edit]
lab@Sun# set policy-options community rtbh members 6451.:666

[edit]
lab@Sun# 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@Sun# set policy-options policy-statement default-filter term 1
from route-filter 0/0 through 0/32

[edit]
lab@Sun# 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. The **through** match type is very specific.

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 and CBGP sessions on each router.

```
[edit]
lab@Sun# set policy-options policy-statement rtbh term 1 from
community rtbh

[edit]
lab@Sun# 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.

```
[edit]
lab@Sun# set routing-options aggregate route 172.30.0.0/16

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

NOTE: Here again you can use **load merge terminal relative** to add this configuration to other devices.

Now you can continue with Configure other routing policies on R1 to meet the task requirements..

You have to 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.

```
[edit]
lab@Sun# set policy-options policy-statement IX-export term 1 from
protocol bgp
```

```
[edit]
lab@Sun# set policy-options policy-statement IX-export term 1 from
community P1
```

```
[edit]
lab@Sun# set policy-options policy-statement IX-export term 1 then
reject
```

```
[edit]
lab@Sun# set policy-options policy-statement IX-export term 2 from
protocol aggregate
```

```
[edit]
lab@Sun# set policy-options policy-statement IX-export term 2 from
route-filter 172.30.0.0/16 exact
```

```
[edit]
lab@Sun# set policy-options policy-statement IX-export term 2 then
accept
```

An import policy has to 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@Sun# set policy-options policy-statement IX-filter term 1 from
route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
lab@Sun# set policy-options policy-statement IX-filter term 1 then
community set IX
```

```
[edit]
lab@Sun# set policy-options policy-statement IX-filter term 1 then
accept
```

```
[edit]
```

```
lab@Sun# 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@Sun# set protocols bgp group IX import default-filter
```

```
[edit]
lab@Sun# set protocols bgp group IX import IX-filter
```

Policies default-filter and IX-filter are applied in import direction.

```
[edit]
lab@Sun# 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@Sun# set protocols bgp group ibgp import rtbh
```

```
[edit]
lab@Sun# set protocols bgp group cbgp import rtbh
```

```
[edit]
lab@Sun# 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@Sirius# set policy-options policy-statement IX-export term 1
from protocol bgp
```

```
[edit]
lab@Sirius# set policy-options policy-statement IX-export term 1
from community P1
```

```
[edit]
lab@Sirius# set policy-options policy-statement IX-export term 1
then reject
```



```

[edit]
lab@Sirius# set policy-options policy-statement IX-export term 2
from protocol aggregate

[edit]
lab@Sirius# set policy-options policy-statement IX-export term 2
from route-filter 172.30.0.0/16 exact

[edit]
lab@Sirius# set policy-options policy-statement IX-export term 2
then as-path-prepend "54591 54591 54591"

[edit]
lab@Sirius# set policy-options policy-statement IX-export term 2
then accept

[edit]
lab@Sirius# set policy-options policy-statement IX-export term 3
from protocol bgp

[edit]
lab@Sirius# set policy-options policy-statement IX-export term 3
then as-path-prepend "54591 54591 54591"

[edit]
lab@Sirius# 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]
lab@Sirius# set policy-options policy-statement IX-filter term 1
from route-filter 0/0 prefix-length-range /8-/24

[edit]
lab@Sirius# set policy-options policy-statement IX-filter term 1
then community set IX

[edit]
lab@Sirius# set policy-options policy-statement IX-filter term 1
then accept

[edit]
lab@Sirius# set policy-options policy-statement IX-filter term 2
then reject

```

Apply the policies on R2.

```
[edit]
lab@Sirius# set protocols bgp group IX import default-filter

[edit]
lab@Sirius# set protocols bgp group IX import IX-filter

[edit]
lab@Sirius# set protocols bgp group IX export IX-export

[edit]
lab@Sirius# set protocols bgp group ibgp import rtbh

[edit]
lab@Sirius# set protocols bgp group cbgp import rtbh

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

c. R3.

R3 has two external sessions with P2 and P3.

Create policies for session with P2.

```
[edit]
lab@Canopus# set policy-options policy-statement P2-export term 1
from protocol aggregate

[edit]
lab@Canopus# set policy-options policy-statement P2-export term 1
from route-filter 172.30.0.0/16 exact

[edit]
lab@Canopus# set policy-options policy-statement P2-export term 1
then accept
```

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@Canopus# set policy-options policy-statement P2-filter term 1
from route-filter 0.0.0.0/0 prefix-length-range /8-/24

[edit]
lab@Canopus# set policy-options policy-statement P2-filter term 1
then local-preference 200

[edit]
```

```
lab@Canopus# set policy-options policy-statement P2-filter term 1
then community set P2
```

```
[edit]
```

```
lab@Canopus# set policy-options policy-statement P2-filter term 1
then accept
```

```
[edit]
```

```
lab@Canopus# 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@Canopus# set policy-options policy-statement P3-export term 1
from protocol aggregate
```

```
[edit]
```

```
lab@Canopus# set policy-options policy-statement P3-export term 1
from route-filter 172.30.0.0/16 exact
```

```
[edit]
```

```
lab@Canopus# 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]
```

```
lab@Canopus# 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 an other AS and passing through AS 43208.365 are not matched.

Create the import policy for session with P3.

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 1
from as-path P3-local-routes
```

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 1
route-filter 0/0 prefix-length-range /32-/32
```

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 1
then local-preference 200
```

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 1
then community set P3
```

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 1
then accept
```

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 2
from route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 2
then local-preference 200
```

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 2
then community set P3
```

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 2
then accept
```

```
[edit]
```

```
lab@Canopus# set policy-option policy-statement P3-filter term 3
then reject
```

Apply the policies to P2 and P3 sessions.

```
[edit]
```

```
lab@Canopus# set protocols bgp group P2-1 export P2-export
```

```
[edit]
```

```
lab@Canopus# set protocols bgp group P2-1 import default-filter
```

```
[edit]
```

```
lab@Canopus# set protocols bgp group P2-1 import P2-filter
```

```
[edit]
```

```
lab@Canopus# set protocols bgp group P3-1 export P3-export
```

```
[edit]
```

```
lab@Canopus# set protocols bgp group P3-1 import default-filter
```

```
[edit]
```

```
lab@Canopus# set protocols bgp group P3-1 import P3-filter
```

```
[edit]
lab@Canopus# set protocols bgp group ibgp import rtbh

[edit]
lab@Canopus# set protocols bgp group cbgp import rtbh

[edit]
lab@Canopus# 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@Arcturus # set protocols bgp group ibgp import rtbh

[edit]
lab@Arcturus # set protocols bgp group cbgp import rtbh

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

e. R5.

R5 has two sessions with customers from AS 64514. Here you have to be careful because these BGP sessions carry IPv6 and IPv4 prefixes. Routing policies should deal with IPv4 traffic also.

For this reason, you have to 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@A-Centauri# set policy-options policy-statement as-internal
term 1 from protocol aggregate

[edit]
lab@A-Centauri# set policy-options policy-statement as-internal
term 1 from route-filter 172.30.0.0/16 exact

[edit]
lab@A-Centauri# 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 one additional community to each routes having EBGp sessions with customer neighbors.

```
[edit]
```

```

lab@A-Centauri# set policy-options community C3-low-pref members
64514:90
Creating import policy.

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
1 from family inet6

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
1 then accept

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
2 from community C3-low-pref

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
2 from prefix-filter 0/0 prefix-length-range /8-/24

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
2 then local-preference 90

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
2 then community add C3

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
2 then accept

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
3 from route-filter 0/0 prefix-length-range /8-/24

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
3 then local-preference 300

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
3 then community add C3

[edit]
lab@A-Centauri# set policy-options policy-statement C3-filter term
3 then accept

[edit]
lab@A-Centauri# 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@A-Centauri# set protocols bgp group C3 export as-internal

[edit]
lab@A-Centauri# set protocols bgp group C3 import default-filter

[edit]
lab@A-Centauri# set protocols bgp group C3 import C3-filter

[edit]
lab@A-Centauri# set protocols bgp group ibgp import rtbh

[edit]
lab@A-Centauri# set protocols bgp group cbgp import rtbh

[edit]
lab@A-Centauri# commit
commit complete
```

f. R6.

R6 has EBGp sessions with two customers from different ASs.

Here you have to create two new communities for matching less preferred routes, one for C2 and one C1.

```
[edit]
lab@Vega# set policy-options community C1-low-pref members
64512:90

[edit]
lab@Vega# 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.

```
[edit]
lab@Vega# set routing-options aggregate route 0/0
```

Create export policies for the sessions with C1-1 neighbor.

```
[edit]
lab@Vega# set policy-options policy-statement as-internal term 1
from protocol aggregate
```

```
[edit]
lab@Vega# set policy-options policy-statement as-internal term 1
from route-filter 172.30.0.0/16 exact
```

```
[edit]
lab@Vega# set policy-options policy-statement as-internal term 1
then metric 10
```

```
[edit]
lab@Vega# 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 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@Vega# set policy-options policy-statement med-10 term 1 from
protocol bgp
```

```
[edit]
lab@Vega# set policy-options policy-statement med-10 term 1 then
metric 10
```

```
[edit]
lab@Vega# set policy-options policy-statement med-10 term 1 then
accept
```

Configure import policy for C1 session.

```
[edit]
lab@Vega# set policy-option policy-statement C1-filter term 1 from
community C1-low-pref
```

```
[edit]
lab@Vega# set policy-option policy-statement C1-filter term 1
route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
lab@Vega# set policy-option policy-statement C1-filter term 1 then
local-preference 90
```

```
[edit]
lab@Vega# set policy-option policy-statement C1-filter term 1 then
community add C1
```

```
[edit]
```



```
lab@Vega# set policy-option policy-statement C1-filter term 1 then
accept
```

```
[edit]
```

```
lab@Vega# set policy-option policy-statement C1-filter term 2 from
route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
```

```
lab@Vega# set policy-option policy-statement C1-filter term 2 then
local-preference 400
```

```
[edit]
```

```
lab@Vega# set policy-option policy-statement C1-filter term 2 then
community add C1
```

```
[edit]
```

```
lab@Vega# set policy-option policy-statement C1-filter term 2 then
accept
```

```
[edit]
```

```
lab@Vega# 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@Vega# set policy-options policy-statement default-only term 1
from protocol aggregate
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement default-only term 1
from route-filter 0/0 exact
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement default-only term 1
then accept
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement default-only term 2
then reject
```

Unlike the export policies on the other routers here you have to configure Term 2 to reject the re-advertisement of other BGP routes, because you have to advertise only a default route.

Configure import policy for C2-1.

```
[edit]
```

```
lab@Vega# set policy-options policy-statement C2-filter term 1
from community C2-low-pref
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement C2-filter term 1
route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement C2-filter term 1
then local-preference 90
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement C2-filter term 1
then community add C2
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement C2-filter term 1
then accept
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement C2-filter term 2
from route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement C2-filter term 2
then local-preference 300
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement C2-filter term 2
then community add C2
```

```
[edit]
```

```
lab@Vega# set policy-options policy-statement C2-filter term 2
then accept
```

```
[edit]
```

```
lab@Vega# 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@Vega# set protocols bgp group C1-1 export as-internal
```

```
[edit]
```

```
lab@Vega# set protocols bgp group C1-1 export med-10
```

```
[edit]
```

```
lab@Vega# set protocols bgp group C1-1 import default-filter
```

```
[edit]
```

```
lab@Vega# set protocols bgp group C1-1 import C1-filter
```

```
[edit]
```

```
lab@Vega# set protocols bgp group C2-1 export default-only
```

```
[edit]
lab@Vega# set protocols bgp group C2-1 import default-filter

[edit]
lab@Vega# set protocols bgp group C2-1 import C2-filter

[edit]
lab@Vega# set protocols bgp group ibgp import rtbh

[edit]
lab@Vega# set protocols bgp group cbgp import rtbh

[edit]
lab@Vega# 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@Rigel# set policy-options community C1-low-pref members
64512:90
```

Configure the export policies for the session with C1.

```
[edit]
lab@Rigel# set policy-options policy-statement as-internal term 1
from protocol aggregate

[edit]
lab@Rigel# set policy-options policy-statement as-internal term 1
from route-filter 172.30.0.0/16 exact

[edit]
lab@Rigel# set policy-options policy-statement as-internal as-
internal term 1 then metric 20

[edit]
lab@Rigel# 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@Rigel# set policy-options policy-statement med-20 term 1 from
protocol bgp

[edit]
lab@Rigel# set policy-options policy-statement med-20 term 1 then
metric 20
```

```
[edit]
lab@Rigel# set policy-options policy-statement med-20 term 1 then
accept
```

Configure import policy for C1 neighbor.

```
[edit]
lab@Rigel# set policy-options policy-statement C1-filter term 1
from community C1-low-pref
```

```
[edit]
lab@Rigel# set policy-options policy-statement C1-filter term 1
from route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
lab@Rigel# set policy-options policy-statement C1-filter term 1
then local-preference 90
```

```
[edit]
lab@Rigel# set policy-options policy-statement C1-filter term 1
then community add C1
```

```
[edit]
lab@Rigel# set policy-options policy-statement C1-filter term 1
then accept
```

```
[edit]
lab@Rigel# set policy-options policy-statement C1-filter term 2
from route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
lab@Rigel# set policy-options policy-statement C1-filter term 2
then local-preference 300
```

```
[edit]
lab@Rigel# set policy-options policy-statement C1-filter term 2
then community add C1
```

```
[edit]
lab@Rigel# set policy-options policy-statement C1-filter term 2
then accept
```

```
[edit]
lab@Rigel# 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@Rigel# set routing-options aggregate route 172.30.128.0/17

[edit]
lab@Rigel# set policy-options community no-export members no-
export

[edit]
lab@Rigel# set policy-options policy-statement P1-export term 1
from community IX

[edit]
lab@Rigel# set policy-options policy-statement P1-export term 1
then reject

[edit]
lab@Rigel# set policy-options policy-statement P1-export term 2
from protocol aggregate

[edit]
lab@Rigel# set policy-options policy-statement P1-export term 2
from route-filter 172.30.128.0/17 exact

[edit]
lab@Rigel# set policy-options policy-statement P1-export term 2
then community set no-export

[edit]
lab@Rigel# set policy-options policy-statement P1-export term 2
then accept

[edit]
lab@Rigel# set policy-options policy-statement P1-export term 3
from protocol aggregate

[edit]
lab@Rigel# set policy-options policy-statement P1-export term 3
from route-filter 172.16.0.0/16 exact

[edit]
lab@Rigel# 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@Rigel# set policy-options as-path P1 110047427
```

Creating AS path that matches routes originated by P1 AS.

```
[edit]
lab@Rigel# set policy-options policy-statement P1-filter term 1
from as-path P1
```

```
[edit]
lab@Rigel# set policy-options policy-statement P1-filter term 1
from route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
lab@Rigel# set policy-options policy-statement P1-filter term 1
then local-preference 200
```

```
[edit]
lab@Rigel# set policy-options policy-statement P1-filter term 1
then community set P1
```

```
[edit]
lab@Rigel# set policy-options policy-statement P1-filter term 1
then accept
```

```
[edit]
lab@Rigel# set policy-options policy-statement P1-filter term 2
then reject
```

Applying the policies to the BGP sessions.

```
[edit]
lab@Rigel# set protocols bgp group C1-1 import default-filter
```

```
[edit]
lab@Rigel# set protocols bgp group C1-1 import C1-filter
```

```
[edit]
lab@Rigel# set protocols bgp group C1-1 export as-internal
```

```
[edit]
lab@Rigel# set protocols bgp group C1-1 export med-20
```

```
[edit]
lab@Rigel# set protocols bgp group P1-2 import default-filter
```

```
[edit]
lab@Rigel# set protocols bgp group P1-2 import P1-filter
```

```
[edit]
lab@Rigel# set protocols bgp group P1-2 export P1-export
```

```
[edit]
lab@Rigel# set protocols bgp group ibgp import rtbh

[edit]
lab@Rigel# set protocols bgp group cbgp import rtbh

[edit]
lab@Rigel# 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 the your local routes with the no-export community.

```
[edit]
lab@Procyon# set routing-options aggregate route 172.30.0.0/17

[edit]
lab@Procyon# set policy-options community no-export members no-
export

[edit]
lab@Procyon# set policy-options policy-statement P1-export term 1
from protocol bgp

[edit]
lab@Procyon# set policy-options policy-statement P1-export term 1
from community IX

[edit]
lab@Procyon# set policy-options policy-statement P1-export term 1
then reject

[edit]
lab@Procyon# set policy-options policy-statement P1-export term 2
from protocol aggregate

[edit]
lab@Procyon# set policy-options policy-statement P1-export term 2
from route-filter 172.30.0.0/17 exact

[edit]
lab@Procyon# set policy-options policy-statement P1-export term 2
then community set no-export

[edit]
lab@Procyon# set policy-options policy-statement P1-export term 2
then accept

[edit]
lab@Procyon# set policy-options policy-statement P1-export term 3
from protocol aggregate
```

```
[edit]
lab@Procyon# set policy-options policy-statement P1-export term 3
from route-filter 172.30.0.0/16 exact
```

```
[edit]
lab@Procyon# 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 than the same routes received from R7.

```
[edit]
lab@Procyon# set policy-options as-path P1 110047427
```

```
[edit]
lab@Procyon# set policy-options policy-statement P1-filter term 1
from as-path P1
```

```
[edit]
lab@Procyon# set policy-options policy-statement P1-filter term 1
from route-filter 0/0 prefix-length-range /8-/24
```

```
[edit]
lab@Procyon# set policy-options policy-statement P1-filter term 1
then local-preference 210
```

```
[edit]
lab@Procyon# set policy-options policy-statement P1-filter term 1
then community set P1
```

```
[edit]
lab@Procyon# set policy-options policy-statement P1-filter term 1
then accept
```

```
[edit]
lab@Procyon# set policy-options policy-statement P1-filter term 2
then reject
```

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@Sun# 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@Sun# run show route protocol bgp terse | match "(/[0-7] )|(/2[5-9]
)|(/3[0-2] )"
* 172.17.47.2/32      B 170    100      >172.30.0.2      (65001)
2831679853 I
* 172.17.47.3/32      B 170    100      >172.30.0.2      (65001)
2831679853 I
```

Only both routes received from P3-1 are present with mask longer than /32. They are allow by the task.

- Verify the routing table for 0.0.0.0 with any mask.

```
[edit]
lab@Sun# 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@Sun# run show route advertised-protocol bgp 172.30.5.8 aspath-regex
"1620 .*" community-name IX

inet.0: 830 destinations, 1796 routes (668 active, 0 holddown, 324
hidden)
  Prefix                Nexthop          MED    Lclpref    AS path
* 1.64.0.0/10           192.168.1.3      100      100        1620 61671 I
* 1.84.160.0/20         192.168.1.3      100      100        1620 33112 I
* 1.96.0.0/11           192.168.1.3      100      100        1620 33112 63164 40776
51777 I
* 1.161.192.0/21        192.168.1.3      100      100        1620 33112 30404 32138
45045 I
* 1.166.89.0/24         192.168.1.3      100      100        1620 110047427 56422
47123 34908 57085 3090 19128 21300 I
* 1.176.0.0/12         192.168.1.3      100      100        1620 33112 49129 16320
52954 I
* 3.49.128.0/19         192.168.1.3      100      100        1620 110047427 56422
7923 42335 17359 9293 38097 38522 61520 I
* 3.189.0.0/16         192.168.1.3      100      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@Sun# run show route advertised-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		
* 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@Sun# 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@Sun# 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
                        > 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@Sun# 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@Sun# 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

```

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@Sun# run show route advertising-protocol bgp 192.168.1.3
community-name P1

inet.0: 830 destinations, 1796 routes (668 active, 0 holddown, 324
hidden)

[edit]
lab@Sun# run show route advertising-protocol bgp 192.168.1.4
community-name P1

inet.0: 830 destinations, 1796 routes (668 active, 0 holddown, 324
hidden)

```

b. R2.

- Verify if unresolved routes are present in the routing table.

```

[edit]
lab@Sirius# 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@Sirius# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32    B 170    100    >172.30.0.14    (65001)
2831679853 I
* 172.17.47.3/32    B 170    100    >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@Sirius# 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@Sirius# run show route advertised-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      100      1620 61671 I
* 1.84.160.0/20         192.168.1.4      100      100      1620 33112 I
* 1.96.0.0/11           192.168.1.4      100      100      1620 33112 63164
40776 51777 I
* 1.161.192.0/21        192.168.1.4      100      100      1620 33112 30404
32138 45045 I
* 1.166.89.0/24         192.168.1.4      100      100      1620 110047427
56422 47123 34908 57085 3090 19128 21300 I
* 1.176.0.0/12          192.168.1.4      100      100      1620 33112 49129
16320 52954 I
* 3.49.128.0/19         192.168.1.4      100      100      1620 110047427
56422 7923 42335 17359 9293 38097 38522 61520 I
* 3.189.0.0/16          192.168.1.4      100      100      1620 33112 41164
7136 6370 35389 39823 16965 I
```

- Verify if customer routes are advertised to IX peers.

```
[edit]
lab@Sirius# run show route advertised-protocol bgp 192.168.1.4
aspath-regex ".* [64512-64514]"

inet.0: 830 destinations, 1852 routes (668 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
* 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@Sirius# run show route advertising-protocol bgp 192.168.1.4
172.30/16

inet.0: 830 destinations, 1852 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@Sirius# run show route protocol bgp aspath-regex ".* [64512-
64514]"

inet.0: 830 destinations, 1852 routes (668 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@Sirius# run show route protocol bgp community-name rtbh

inet.0: 830 destinations, 1852 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@Sirius# run show route 172.17.0.0/24

inet.0: 830 destinations, 1852 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
                    > to 192.168.1.4 via ge-0/0/5.300
```

- Verify if routes from P1 are not send to IX peers.

```
[edit]
lab@Sirius# run show route adverting-protocol bgp 192.168.1.3
community-name P1

inet.0: 830 destinations, 1852 routes (668 active, 0 holddown, 324
hidden)

[edit]
lab@Sirius# run show route advertising-protocol bgp 192.168.1.4
community-name P1

inet.0: 830 destinations, 1852 routes (668 active, 0 holddown, 324
hidden)
```

c. R3.

- Verify if unresolved routes are present in the routing table.

```
[edit]
lab@Canopus# 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@Canopus# 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
```

- Verify the routing table for route/ prefix 0.0.0.0 with any mask.

```
[edit]
lab@Canopus# 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@Canopus# run show route adverting-protocol bgp 172.30.5.4
community-name P2

inet.0: 731 destinations, 734 routes (673 active, 0 holddown, 60
hidden)
  Prefix                Nexthop      MED      Lclpref    AS path
* 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@Canopus# run show route advertising-protocol bgp 172.30.5.4
community-name P3

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 ?
```

```
* 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 routes with a mask shorter than /8 and longer than /24 are present in routing table.

```
[edit]
lab@Canopus# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32      B 170      100      >172.30.0.14      (65001)
2831679853 I
* 172.17.47.3/32      B 170      100      >172.30.0.14      (65001)
2831679853 I
```

- Verify the routing table for route 0.0.0.0 with any mask.

```
[edit]
lab@Sirius# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if the customer routes are advertised to P2 and P3 peers.

```
[edit]
lab@Canopus# run show route advertised-protocol bgp 192.168.0.2
aspath-regex "64512|64513|64514"

inet.0: 731 destinations, 734 routes (673 active, 0 holddown, 60
hidden)
  Prefix                Nexthop                MED      Lc1pref      AS path
* 172.31.0.0/24        Self                    0         0             64512 I
* 172.31.1.0/24        Self                    0         0             64512 I
* 172.31.2.0/24        Self                    0         0             64512 I
* 172.31.3.0/24        Self                    0         0             64512 I
* 172.31.4.0/24        Self                    0         0             64512 I
* 172.31.5.0/24        Self                    0         0             64512 I
* 172.31.6.0/24        Self                    0         0             64512 I
* 172.31.7.0/24        Self                    0         0             64512 I
* 172.31.8.0/24        Self                    0         0             64512 I
* 172.31.9.0/24        Self                    0         0             64512 I
* 172.31.10.0/24       Self                    0         0             64512 I
* 172.31.11.0/24       Self                    0         0             64512 I
* 172.31.12.0/24       Self                    0         0             64512 I
* 172.31.13.0/24       Self                    0         0             64512 I
* 172.31.14.0/24       Self                    0         0             64512 I
* 172.31.15.0/24       Self                    0         0             64512 I
* 172.31.16.0/24       Self                    0         0             64513 I
* 172.31.17.0/24       Self                    0         0             64513 I
* 172.31.18.0/24       Self                    0         0             64513 I
* 172.31.19.0/24       Self                    0         0             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@Canopus# run show route advertised-protocol bgp 192.168.0.6
aspath-regex "64512|64513|64514"

inet.0: 731 destinations, 734 routes (673 active, 0 holddown, 60
hidden)
  Prefix          Nexthop          MED      Lclpref    AS path
* 172.31.0.0/24   Self              0
* 172.31.1.0/24   Self              0
* 172.31.2.0/24   Self              0
* 172.31.3.0/24   Self              0
* 172.31.4.0/24   Self              0
* 172.31.5.0/24   Self              0
* 172.31.6.0/24   Self              0
* 172.31.7.0/24   Self              0
* 172.31.8.0/24   Self              0
* 172.31.9.0/24   Self              0
* 172.31.10.0/24  Self              0
* 172.31.11.0/24  Self              0
* 172.31.12.0/24  Self              0
* 172.31.13.0/24  Self              0
* 172.31.14.0/24  Self              0
* 172.31.15.0/24  Self              0
* 172.31.16.0/24  Self              0

```

- Verify if the local summary prefix is sent to EBGp peers.

```
[edit]
lab@Canopus# run show route advertising-protocol bgp 192.168.0.2
172.30/16

inet.0: 731 destinations, 734 routes (673 active, 0 holddown, 60
hidden)
  Prefix                Nexthop                MED      Lclpref
AS path
* 172.30.0.0/16         Self
I

[edit]
lab@Canopus# run show route advertising-protocol bgp 192.168.0.6
172.30/16

inet.0: 731 destinations, 734 routes (673 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@Canopus# run show route protocol bgp aspath-regex ".* [64512-
64514]"

inet.0: 731 destinations, 734 routes (673 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@Canopus# run show route protocol bgp community-name rtbh

inet.0: 731 destinations, 734 routes (673 active, 0 holddown, 60
hidden)
+ = 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@Arcturus# 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@Arcturus# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32  B 170    100          >172.30.0.21
2831679853 I
* 172.17.47.3/32  B 170    100          >172.30.0.21
2831679853 I
```

- Verify the routing table for 0.0.0.0 with any mask.

```
[edit]
lab@Arcturus# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if customer routes are received by R4.


```
[edit]
lab@Arcturus# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 673 destinations, 674 routes (673 active, 0 holddown, 0 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
                    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@Arcturus# run show route protocol bgp community-name rtbh

inet.0: 673 destinations, 674 routes (673 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.1.0/24      *[BGP/170] 00:57:38, 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] 00:57:38, localpref 400, from
172.30.5.5
                    AS path: (65002) 64512 I
                    > to 172.30.0.30 via ge-0/0/4.145
```

e. R5.

- Verify if unresolved routes are present in the routing table.

```
[edit]
lab@A-Centauri# 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@A-Centauri# 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    (65001)
2831679853 I
* 172.17.47.3/32    B 170    100    >172.30.0.34    (65001)
2831679853 I
```

- Verify the routing table for 0.0.0.0 with any mask.

```
[edit]
lab@A-Centauri# run show route protocol bgp terse | match "
0.0.0.0"
```

- Verify if customer routes are advertised to the EBGp peers.

```
[edit]
lab@A-Centauri# run show route advertising-protocol bgp
192.168.0.10 aspath-regex "64512|64513|64514"

inet.0: 675 destinations, 691 routes (675 active, 0 holddown, 0
hidden)
  Prefix                Nexthop                MED      Lclpref    AS path
* 172.31.0.0/24         Self                    0         0          64512 I
* 172.31.1.0/24         Self                    0         0          64512 I
* 172.31.2.0/24         Self                    0         0          64512 I
* 172.31.3.0/24         Self                    0         0          64512 I
* 172.31.4.0/24         Self                    0         0          64512 I
* 172.31.5.0/24         Self                    0         0          64512 I
* 172.31.6.0/24         Self                    0         0          64512 I
* 172.31.7.0/24         Self                    0         0          64512 I
* 172.31.8.0/24         Self                    0         0          64512 I
* 172.31.9.0/24         Self                    0         0          64512 I
* 172.31.10.0/24        Self                    0         0          64512 I
* 172.31.11.0/24        Self                    0         0          64512 I
* 172.31.12.0/24        Self                    0         0          64512 I
* 172.31.13.0/24        Self                    0         0          64512 I
* 172.31.14.0/24        Self                    0         0          64512 I
* 172.31.15.0/24        Self                    0         0          64512 I
* 172.31.16.0/24        Self                    0         0          64513 I
* 172.31.17.0/24        Self                    0         0          64513 I
* 172.31.18.0/24        Self                    0         0          64513 I
* 172.31.19.0/24        Self                    0         0          64513 I
* 172.31.20.0/24        Self                    0         0          64513 I
* 172.31.21.0/24        Self                    0         0          64513 I
* 172.31.22.0/24        Self                    0         0          64513 I
* 172.31.23.0/24        Self                    0         0          64513 I
* 172.31.24.0/24        Self                    0         0          64513 I
* 172.31.25.0/24        Self                    0         0          64513 I
* 172.31.26.0/24        Self                    0         0          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
```

- Verify if the local summary prefix is sent to EBGp peers.

```
[edit]
lab@A-Centauri# run show route advertising-protocol bgp
192.168.0.10 172.30/16

inet.0: 675 destinations, 691 routes (675 active, 0 holddown, 0
hidden)
  Prefix                Nexthop                MED      Lclpref
AS path
* 172.30.0.0/16         Self
I
```

```
[edit]
lab@A-Centauri# run show route advertising-protocol bgp
192.168.0.14 172.30/16

inet.0: 675 destinations, 691 routes (675 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.

```
[edit]
lab@A-Centauri# run show route protocol bgp aspath-regex ".*
[64512-64514]"

inet.0: 675 destinations, 691 routes (675 active, 0 holddown, 0
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

```

- Verify the remote triggered black hole routes with community **rtbh**.

```

[edit]
lab@A-Centauri# run show route protocol bgp community-name rtbh

inet.0: 675 destinations, 691 routes (675 active, 0 holddown, 0
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@A-Centauri# run show route aspath-regex 64514

inet.0: 675 destinations, 691 routes (675 active, 0 holddown, 0
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@A-Centauri# 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      192.168.0.10                ulst 262142 17
ge-0/0/5.303         ucst    618    3
192.168.0.14        ucst    616    3
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
```

f. R6.

- Verify if unresolved routes are present in the routing table.

```
[edit]
lab@Vega# 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@Vega# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32 B 170 100 >172.30.0.25 (65001)
2831679853 I
* 172.17.47.3/32 B 170 100 >172.30.0.25 (65001)
2831679853 I
```

- Verify the routing table for a route 0.0.0.0 with any mask.

```
[edit]
lab@Vega# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if customer routes are advertised to the EBGp peers.

```
[edit]
lab@Vega# run show route advertising-protocol bgp 172.31.31.1
aspath-regex "64512|64513|64514"

inet.0: 676 destinations, 693 routes (676 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
* 172.31.32.0/24     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

```

[edit]

```
lab@Vega# run show route advertising-protocol bgp 192.168.0.18
aspath-regex "64512|64513|64514"
```

```
inet.0: 676 destinations, 693 routes (676 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			
* 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@Vega# run show route advertising-protocol bgp 192.168.0.18
172.30/16

inet.0: 676 destinations, 693 routes (676 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@Vega# 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                      {101 235 339 448 534
647 794 859 ... 63887} ?

```

NOTE: AS path holds all AS numbers of the contributing routes.

- Verify the local-preference of the customer routes received by R6.

```
[edit]
lab@Vega# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 676 destinations, 693 routes (676 active, 0 holddown, 0 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@Vega# run show route protocol bgp community-name rtbh

inet.0: 676 destinations, 693 routes (676 active, 0 holddown, 0 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@Vega# run show route aspath-regex 64513

inet.0: 676 destinations, 693 routes (676 active, 0 holddown, 0 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
                   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@Vega# run show route 172.31.1/24

inet.0: 676 destinations, 693 routes (676 active, 0 holddown, 0
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@Rigel# 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@Rigel# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32   B 170      100      >172.30.0.17      (65002
65001) 2831679853 I
* 172.17.47.3/32   B 170      100      >172.30.0.17      (65002
65001) 2831679853 I

```

- Verify the routing table for a route 0.0.0.0 with any mask.

```

[edit]
lab@Rigel# run show route protocol bgp terse | match " 0.0.0.0"

```

- Verify if customer routes are advertised to the EBGp peers.

```
[edit]
lab@Rigel# run show route advertising-protocol bgp 192.168.0.30
aspath-regex "64512|64513|64514"

inet.0: 834 destinations, 1779 routes (672 active, 0 holddown, 548
hidden)
  Prefix                Nexthop                MED      Lclpref    AS path
* 172.31.0.0/24         Self                    20        0          64512 I
* 172.31.1.0/24         Self                    20        0          64512 I
* 172.31.2.0/24         Self                    20        0          64512 I
* 172.31.3.0/24         Self                    20        0          64512 I
* 172.31.4.0/24         Self                    20        0          64512 I
* 172.31.5.0/24         Self                    20        0          64512 I
* 172.31.6.0/24         Self                    20        0          64512 I
* 172.31.7.0/24         Self                    20        0          64512 I
* 172.31.8.0/24         Self                    20        0          64512 I
* 172.31.9.0/24         Self                    20        0          64512 I
* 172.31.10.0/24        Self                    20        0          64512 I
* 172.31.11.0/24        Self                    20        0          64512 I
* 172.31.12.0/24        Self                    20        0          64512 I
* 172.31.13.0/24        Self                    20        0          64512 I
* 172.31.14.0/24        Self                    20        0          64512 I
* 172.31.15.0/24        Self                    20        0          64512 I
* 172.31.16.0/24        Self                    20        0          64513 I
* 172.31.17.0/24        Self                    20        0          64513 I
* 172.31.18.0/24        Self                    20        0          64513 I
* 172.31.19.0/24        Self                    20        0          64513 I
[edit]
lab@Rigel# run show route advertising-protocol bgp 192.168.0.34
aspath-regex "64512|64513|64514"

inet.0: 834 destinations, 1779 routes (672 active, 0 holddown, 548
hidden)
  Prefix                Nexthop                MED      Lclpref    AS path
* 172.31.16.0/24        Self                    20        0          64513 I
* 172.31.17.0/24        Self                    20        0          64513 I
* 172.31.18.0/24        Self                    20        0          64513 I
* 172.31.19.0/24        Self                    20        0          64513 I
* 172.31.20.0/24        Self                    20        0          64513 I
* 172.31.21.0/24        Self                    20        0          64513 I
* 172.31.22.0/24        Self                    20        0          64513 I
* 172.31.23.0/24        Self                    20        0          64513 I
* 172.31.24.0/24        Self                    20        0          64513 I
* 172.31.25.0/24        Self                    20        0          64513 I
* 172.31.26.0/24        Self                    20        0          64513 I
* 172.31.27.0/24        Self                    20        0          64513 I
* 172.31.28.0/24        Self                    20        0          64513 I
* 172.31.29.0/24        Self                    20        0          64513 I
* 172.31.30.0/24        Self                    20        0          64513 I
* 172.31.31.0/24        Self                    20        0          64513 I
* 172.31.32.0/24        Self                    20        0          64514 I
* 172.31.33.0/24        Self                    20        0          64514 I
* 172.31.34.0/24        Self                    20        0          64514 I
* 172.31.35.0/24        Self                    20        0          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@Rigel# run show route advertising-protocol bgp 192.168.0.30
172.30/16

inet.0: 834 destinations, 1779 routes (672 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@Rigel# run show route advertising-protocol bgp 192.168.0.34
172.30/16

inet.0: 834 destinations, 1779 routes (672 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@Rigel# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 834 destinations, 1779 routes (672 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@Rigel# run show route protocol bgp community-name rtbh

inet.0: 834 destinations, 1779 routes (672 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@Rigel# run show route 172.31.1/24

inet.0: 834 destinations, 1779 routes (672 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
                    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@Rigel# run show route receive-protocol bgp 192.168.0.30
aspath-regex "110047427 .+"

inet.0: 834 destinations, 1779 routes (672 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 the routes to P1 are advertised with community no-export.

```
[edit]
lab@Rigel# run show route advertising-protocol bgp 192.168.0.30
172.30/16 extensive

inet.0: 834 destinations, 1779 routes (672 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@Procyon# 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@Procyon# run show route protocol bgp | match "(/[0-7] )|(/2[5-9] )|(/3[0-2] )"
* 172.17.47.2/32  B 170    100      >172.30.0.9      (65000 65001)
2831679853 I
* 172.17.47.3/32  B 170    100      >172.30.0.9      (65000 65001)
2831679853 I
```

- Verify the routing table for route 0.0.0.0 with any mask.

```
[edit]
lab@Procyon# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if customer routes are advertised to the EBGp peers.

```
[edit]
lab@Procyon# run show route advertising-protocol bgp 192.168.0.38
aspath-regex "64512|64513|64514"

inet.0: 832 destinations, 1745 routes (670 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		
* 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		
* 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

```

- Verify if the local summary prefix is sent to EBGp peers.

```

[edit]
lab@Procyon# run show advertising-protocol bgp 192.168.0.38
172.30/16

inet.0: 832 destinations, 1745 routes (670 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

```

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@Procyon# run show route protocol bgp aspath-regex ".* [64512-64514]"

inet.0: 832 destinations, 1745 routes (670 active, 0 holddown, 548 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
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@Procyon# run show route protocol bgp community-name rtbh

inet.0: 832 destinations, 1745 routes (670 active, 0 holddown, 548 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@Procyon# run show route 172.17.0.0/24

inet.0: 832 destinations, 1745 routes (670 active, 0 holddown, 548 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@Procyon# run show route receive-protocol bgp 192.168.0.38
aspath-regex "110047427 .+"

inet.0: 832 destinations, 1745 routes (670 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 the routes to P1 are advertised with community no-export.

```
[edit]
lab@Procyon# run show route advertising-protocol bgp 192.168.0.38
172.30/16 extensive

inet.0: 832 destinations, 1745 routes (670 active, 0 holddown, 548
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 have to redesign the network by introducing Route Reflection instead of confederations. You have to 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 has the ability to 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 as 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.
- *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 is able to 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 route reflector are capable of caring a lot prefixes in the RIB and having less power in the forwarding engine.

- 1) Delete IBGP configuration on all routers including the confederation.

```
[edit]
lab@Sun# delete routing-options confederation

[edit]
lab@Sun# delete protocols bgp group ibgp

[edit]
lab@Sun# delete protocols bgp group cbgp

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

- 2) Configure additional interfaces on R1 and R2.

For the Route Reflector we will be using a dedicated J-series router. It is physically connected to R1 and R2 as shown in the beginning of the Task.

a. Configure R1.

```
[edit]
lab@Sun# set interfaces ge-0/0/4.206 vlan-id 206

[edit]
lab@Sun# set interfaces ge-0/0/4.206 family inet address
172.30.0.65/30

[edit]
lab@Sun# set interfaces ge-0/0/4.206 family iso

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

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. Configure R2.

```
[edit]
lab@Sirius# set interfaces ge-0/0/4.207 vlan-id 207

[edit]
lab@Sirius# set interfaces ge-0/0/4.207 family inet address
172.30.0.69/30

[edit]
lab@Sirius# set interfaces ge-0/0/4.207 family iso

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

c. Verify the connectivity.

```
[edit]
lab@Sun# run ping 172.30.0.66
PING 172.30.0.66 (172.30.0.66): 56 data bytes
64 bytes from 172.30.0.66: icmp_seq=1 ttl=64 time=30.538 ms
64 bytes from 172.30.0.66: icmp_seq=2 ttl=64 time=1.505 ms
64 bytes from 172.30.0.66: icmp_seq=3 ttl=64 time=1.312 ms
64 bytes from 172.30.0.66: icmp_seq=4 ttl=64 time=1.234 ms
64 bytes from 172.30.0.66: icmp_seq=5 ttl=64 time=1.372 ms
64 bytes from 172.30.0.66: icmp_seq=6 ttl=64 time=1.289 ms
64 bytes from 172.30.0.66: icmp_seq=7 ttl=64 time=1.116 ms
```

```
[edit]
lab@Sirius# run ping 172.30.0.70
```

```
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
^C
--- 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.

All routers should be configured with a global AS instead of the one configured for the confederation.

```
[edit]
lab@Sun# set routing-options autonomous-system 54591
```

4) Configure IBGP configuration.

Only one IBGP neighbor, the route reflector should be configured on each routers 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.

```
[edit]
lab@Sun# set protocols bgp group ibgp type internal

[edit]
lab@Sun# set protocols bgp group ibgp local-address 172.30.5.1

[edit]
lab@Sun# set protocols bgp group ibgp import rtbh

[edit]
lab@Sun# set protocols bgp group ibgp authentication-key workbook

[edit]
lab@Sun# set protocols bgp group ibgp bfd-liveness-detection
minimum-interval 500

[edit]
lab@Sun# set protocols bgp group ibgp neighbor 172.30.5.41

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

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 but R1 and R3.

The policy configuration is available from the previous confederation exercise. You only have to re-apply.

```
[edit]
lab@Canopus# set protocols bgp group ibgp export nhs
```

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

6) Configure the router reflector.

a. Enabling family iso and Configure NET address.

```
[edit]
root@route-reflector# set interfaces ge-0/0/1.206 family iso
```

```
[edit]
root@route-reflector# set interfaces ge-0/0/1.207 family iso
```

```
[edit]
root@route-reflector# set interface lo0.0 family iso address
49.0001.1720.3000.5041.00
```

b. Configure IS-IS protocol.

```
[edit]
root@route-reflector# set protocols isis level 2 disable
```

```
[edit]
root@route-reflector# set protocols isis level 1 authentication-
type md5
```

```
[edit]
root@route-reflector# set protocols isis level 1 authentication-
key workbook
```

```
[edit]
root@route-reflector# set protocols isis interface all point-to-
point
```

```
[edit]
root@route-reflector# set protocols isis interface all bfd-
liveness-detection multiplier 3
```

```
[edit]
root@route-reflector# set protocols isis interface all bfd-
liveness-detection minimum-interval 150
```

c. Configure the autonomous system

```
[edit]
root@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]
root@route-reflector# set protocols bgp group cluster-1 type
internal

[edit]
root@route-reflector# set protocols bgp group cluster-1 local-
address 172.30.5.41

[edit]
root@route-reflector# set protocols bgp group cluster-1
authentication-key workbook

[edit]
root@route-reflector# set protocols bgp group cluster-1 cluster
0.0.0.1

[edit]
root@route-reflector# set protocols bgp group cluster-1 bfd-
liveness-detection minimum-interval 500

[edit]
root@route-reflector# set protocols bgp group cluster-1 neighbor
172.30.5.1

[edit]
root@route-reflector# set protocols bgp group cluster-1 neighbor
172.30.5.6

[edit]
root@route-reflector# set protocols bgp group cluster-1 neighbor
172.30.5.7

[edit]
root@route-reflector# set protocols bgp group cluster-1 neighbor
172.30.5.8
Configure the second cluster group with cluster 0.0.0.2

[edit]
root@route-reflector# set protocols bgp group cluster-2 type
internal

[edit]
root@route-reflector# set protocols bgp group cluster-2 local-
address 172.30.5.41

[edit]
root@route-reflector# set protocols bgp group cluster-2
authentication-key workbook

[edit]
```

```

root@route-reflector# set protocols bgp group cluster-2 cluster
0.0.0.2

[edit]
root@route-reflector# set protocols bgp group cluster-2 bfd-
liveness-detection minimum-interval 500

[edit]
root@route-reflector# set protocols bgp group cluster-2 neighbor
172.30.5.2

[edit]
root@route-reflector# set protocols bgp group cluster-2 neighbor
172.30.5.3

[edit]
root@route-reflector# set protocols bgp group cluster-2 neighbor
172.30.5.4

[edit]
root@route-reflector# set protocols bgp group cluster-2 neighbor
172.30.5.5

[edit]
root@route-reflector# commit
commit complete

```

e. Verify the operation of both RR clusters.

```

[edit]
root@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      984        600         0           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

10) [edit]

11)	root@route-reflector# run show bfd session				
12)	Transmit				
13)	Address	State	Interface	Detect Time	Interval
	Multiplier				
14)	172.30.0.65	Up	ge-0/0/1.206	0.450	0.150
3					
15)	172.30.0.69	Up	ge-0/0/1.207	0.450	0.150
3					
16)	172.30.5.1	Up		1.500	0.500
3					
17)	172.30.5.4	Up		1.500	0.500
3					
18)	172.30.5.8	Up		1.500	0.500
3					
19)	172.30.5.7	Up		1.500	0.500
3					
20)	172.30.5.5	Up		1.500	0.500
3					
21)	172.30.5.6	Up		1.500	0.500
3					
22)	172.30.5.2	Up		1.500	0.500
3					
23)	172.30.5.3	Up		1.500	0.500
3					

All 8 client routers sessions are established.

```
[edit]
root@route-reflector# run show bgp group brief
Group Type: Internal      AS: 54591      Local AS: 54591
Name: cluster-1          Index: 0        Flags: <>
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.

```
[edit]
lab@Sun# run show bgp neighbor 172.30.5.41
Peer: 172.30.5.41+1385 AS 54591 Local: 172.30.5.1+179 AS 54591
Type: Internal      State: Established      Flags: <ImportEval Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Import: [ rtbh ]
Options: <Preference LocalAddress AuthKey LogUpDown Refresh>
Options: <BfdEnabled>
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: enabled, 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: 10001
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:                215
  Received prefixes:              215
  Accepted prefixes:              215
  Suppressed due to damping:      0
  Advertised prefixes:            383
  Last traffic (seconds): Received 22   Sent 23   Checked 43
  Input messages:  Total 341   Updates 335   Refreshes 0
Octets 33018
  Output messages: Total 378   Updates 371   Refreshes 0
Octets 27599
  Output Queue[0]: 0

```

NOTE: The iNET ZERO 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@Sun# 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

```

RR is receiving the route and reflects it to the clients.

```
[edit]
root@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
```

```

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 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: 172.30.5.2
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		
	Nexthop: 172.30.0.69 via ge-	
0/0/1.207		

Here 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.

Verify the route received on R8.

```
[edit]
lab@Procyon# 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
```



```

0/0/4.118                                     Next hop: 172.30.0.9 via ge-
                                           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@Sun# 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@Sun# 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@Sun# run show route advertising-protocol bgp 172.30.5.41

```

```
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        100        1620
61671 I
* 1.84.160.0/20  192.168.1.3          100        100        1620
33112 I
* 1.96.0.0/11    192.168.1.3          100        100        1620
33112 63164 40776 51777 I
* 1.161.192.0/21 192.168.1.3          100        100        1620
33112 30404 32138 45045 I
--- (more) ---
```

Routes from R1 is advertised with the nexthop of the IX routers. Because they are advertised in IS-IS all routers are able to resolve the routes without a next-hop self policy.

- Verify for routes with mask shorter than /8 and longer than /24

```
[edit]
lab@Sun# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32 B 170 100 >172.30.0.2
2831679853 I
* 172.17.47.3/32 B 170 100 >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@Sun# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if the correct community is set and send to RR.

```
lab@Sun# 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        100        1620 61671 I
* 1.84.160.0/20  192.168.1.3          100        100        1620 33112 I
* 1.96.0.0/11    192.168.1.3          100        100        1620 33112
63164 40776 51777 I
* 1.161.192.0/21 192.168.1.3          100        100        1620 33112
30404 32138 45045 I
* 1.166.89.0/24  192.168.1.3          100        100        1620 110047427
56422 47123 34908 57085 3090 19128 21300 I
* 1.176.0.0/12   192.168.1.3          100        100        1620 33112
49129 16320 52954 I
```

- Verify if customer routes are advertised to the IX peers.

```
[edit]
lab@Sun# run show route adverting-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
--- (more) ---
```

- Verify if the local summary route is advertised to IX peers.

```
[edit]
lab@Sun# 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@Sun# 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.

```

[edit]
lab@Sun# 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@Sun# 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@Sun# 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@Sun# 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.

b. R2.

- Verify the BGP session status with the RR.

```
[edit]
lab@Sirius# run show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table      Tot Paths  Act Paths Suppressed    History Damp State
Pending
inet.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@Sirius# 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@Sirius# 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      100          1620
61671 I
* 1.84.160.0/20 192.168.1.3  100      100          1620
33112 I
* 1.96.0.0/11 192.168.1.3  100      100          1620
33112 63164 40776 51777 I
* 1.161.192.0/21 192.168.1.3  100      100          1620
33112 30404 32138 45045 I
* 1.166.89.0/24 192.168.1.3  100      100          1620
110047427 56422 47123 34908 57085 3090 19128 21300 I
```

As with R1, R2 is not using next-hop self policy.

- Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@Sirius# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32    B 170    100          >172.30.0.14
2831679853 I
* 172.17.47.3/32    B 170    100          >172.30.0.14
2831679853 I
```

- Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@Sirius# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if the correct community is set and send to RR.

```
[edit]
lab@Sirius# run show route advertising-protocol 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       100        1620 61671 I
* 1.84.160.0/20         192.168.1.3      100       100        1620 33112 I
* 1.96.0.0/11           192.168.1.3      100       100        1620 33112
63164 40776 51777 I
* 1.161.192.0/21        192.168.1.3      100       100        1620 33112
30404 32138 45045 I
* 1.166.89.0/24         192.168.1.3      100       100        1620
110047427 56422 47123 34908 57085 3090 19128 21300 I
* 1.176.0.0/12          192.168.1.3      100       100        1620 33112
49129 16320 52954 I
```

- Verify if the customer routes are advertised to the IX peers.

```
[edit]
lab@Sirius# 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
* 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
```

Here you can see that the AS-path prepend makes the routes advertised by R1 more preferred than 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@Sirius# 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@Sirius# 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@Sirius# run show route protocol bgp community-name rtbh

inet.0: 834 destinations, 1800 routes (672 active, 0 holddown, 324
hidden)
+ = 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@Sirius# 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@Sirius# 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@Sirius# run show route advertising-protocol bgp 192.168.1.4
aspath-regex "110047427 .*"
```

c. R3.

- Verify the BGP session status with the RR.

```
[edit]
lab@Canopus# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table Tot Paths  Act Paths Suppressed  History Damp State
Pending
inet.0
          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@Canopus# 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@Canopus# 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]
lab@Canopus# 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
```

- Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@Canopus# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if the correct community is set and send to RR.

```
[edit]
lab@Canopus# 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)
  Prefix                Nexthop          MED      Lclpref    AS path
* 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@Canopus# run show route advertising-protocol bgp 172.30.5.41
community-name P3

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 ?
* 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			

- Verify if customer routes are advertised to the IX peers.

```
[edit]
lab@Canopus# 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
--- (more) ---
```

```
[edit]
lab@Canopus# 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
* 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
--- (more) ---
```

- Verify if the local summary route is advertised to P peers.

```
[edit]
lab@Canopus# 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@Canopus# 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
```

- Verify the local preference of the customer routes received by R3.

```
[edit]
lab@Canopus# 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.
-
-
-
-

```
[edit]
lab@Canopus# 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@Canopus# 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@Arcturus# 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@Arcturus# run show route resolution unresolved
Tree Index 1
Tree Index 2
Tree Index 3
Tree Index 4
```

- Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@Arcturus# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32    B 170    100    >172.30.0.21
2831679853 I
```

```
* 172.17.47.3/32    B 170    100    >172.30.0.21
2831679853 I
```

- Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@Arcturus# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify the local preference of the customer routes received by R4.

```
[edit]
lab@Arcturus# 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@Arcturus# 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

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@A-Centauri# 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@A-Centauri# 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@A-Centauri# 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      Lc1pref      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
--- (more) ---
```

- Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@A-Centauri# 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
```

- Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@A-Centauri# run show route protocol bgp terse | match "
0.0.0.0"
```

- Verify if the correct community is set and send to RR.

```
[edit]
lab@A-Centauri# 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
--- (more) ---
```

- Verify if customer routes are advertised to the customer AS.

```
[edit]
lab@A-Centauri# 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)
```


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				
--- (more) ---				

```
[edit]
lab@A-Centauri# 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
--- (more) ---
```

- Verify if the local summary route is advertised to Customer AS.

```
[edit]
lab@A-Centauri# run show route 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@A-Centauri# run show route 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.

```
[edit]
lab@A-Centauri# 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@A-Centauri# 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@A-Centauri# 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]
lab@A-Centauri# 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
```

f. R6.

- Verify the BGP session status with the RR.

```
[edit]
lab@Vega# 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@Vega# 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@Vega# 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
--- (more) ---
```

- Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@Vega# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32      B 170      100          >172.30.0.25
2831679853 I
* 172.17.47.3/32      B 170      100          >172.30.0.25
2831679853 I
```

- Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@Vega# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if the correct community is set and send to RR.

```
[edit]
lab@Vega# 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
--- (more) ---

```

```

[edit]
lab@Vega# 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
--- (more) ---

```

- Verify if customer routes are advertised to the EBGp peers.

```

[edit]
lab@Vega# 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@Vega# 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
--- (more) ---
```

- Verify if the local summary route is advertised to C1-1 peer.

```
[edit]
lab@Vega# 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@Vega# 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
                        AS path: 64512 I
                        > to 192.168.0.18 via ge-0/0/5.305
```



```

172.30.5.41          [BGP/170] 00:01:41, localpref 90, from
                    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@Vega# 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@Vega# 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
                    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@Vega# 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@Rigel# 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.

24)

```
[edit]
lab@Rigel# 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@Rigel# 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@Rigel# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9] )|(/3[0-2] )"

```

```
* 172.17.47.2/32    B 170      100          >172.30.0.17
2831679853 I
* 172.17.47.3/32    B 170      100          >172.30.0.17
2831679853 I
```

- Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@Rigel# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if the correct community is set and send to RR.

```
[edit]
lab@Rigel# 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@Rigel# 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.

- Verify if customer routes are advertised to the EBGp peers.

```
[edit]
lab@Rigel# 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
--- (more) ---
```

```
[edit]
lab@Rigel# 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
--- (more) ---
```

- Verify if the local summary route is advertised to EBGp peers.

```
[edit]
lab@Rigel# 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
```

R7 is advertising part of the internal routes to P1-2 together with the summary route of all internal routes.

```
[edit]
lab@Rigel# 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@Rigel# 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.

```
[edit]
lab@Rigel# 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@Rigel# 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@Rigel# 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@Rigel# run show route advertising-protocol bgp 192.168.0.30
172.30/16 extensive

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@Procyon# 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@Procyon# 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@Procyon# run show route advertising-protocol bgp 172.30.5.41
```



```
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
--- (more) ---
```

- Verify for routes with mask shorter than /8 and longer than /24.

```
[edit]
lab@Procyon# run show route protocol bgp terse | match "(/[0-7]
)|(/2[5-9])|(/3[0-2])"
* 172.17.47.2/32 B 170 100 >172.30.0.9
2831679853 I
* 172.17.47.3/32 B 170 100 >172.30.0.9
2831679853 I
```

- Verify for routes 0.0.0.0 with any mask.

```
[edit]
lab@Procyon# run show route protocol bgp terse | match " 0.0.0.0"
```

- Verify if the correct community is set and send to RR.

```
[edit]
lab@Procyon# 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
--- (more) ---
```

- Verify if customer routes are advertised to the P1-1 peer.

```
[edit]
lab@Procyon# run show route advertising-protocol 192.168.0.38
aspath-regex "64512|64513|64514"

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
--- (more) ---
```

- Verify if the local summary route is advertised to P1-1 peer.

```
[edit]
lab@Procyon# 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@Procyon# 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.

```

[edit]
lab@Procyon# 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@Procyon# 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@Procyon# 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
```

This was the last check for this Chapter. Now you should save the configurations of all routers as you need them for the next Chapter. Don't forget to save the configuration of the RR!

Chapter Four solutions: 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 plane 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 a 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 were called EXP (experimental).
- **Bottom of the stack** – 1 bit field known as 'S' bit is used to indicate 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 a 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 Router ID is not configured, the interface address is used as the Transport address. A 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.

```
[edit]
lab@Sun# set groups if-families interfaces ge-0/0/4 unit <*>
family mpls
```

```
[edit]
lab@Sun# 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** stanza. 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@Sun# set protocols mpls interface all
```

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

This step is performed on all routers in the network.

Verify if the interfaces are enabled for MPLS.

```
[edit]
lab@Sun# run show mpls interface
```

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-0/0/4.206	Up	<none>

Four interfaces on R1 are enabled for MPLS.

2) Configure LDP according Table 8.

Next you have to enable LDP as label distribution protocol according Figure 8. Some additional features like authentication should be enabled on the sessions. Configurations are the same on all routers.

```
[edit]
lab@Sun# set protocols ldp interface ae0.0
```

```
[edit]
lab@Sun# set protocols ldp interface ge-0/0/4.114
```

Adding both interfaces to LDP protocol.

```
[edit]
lab@Sun# set protocols ldp session 172.30.5.2 authentication-key
workbook
```

```
[edit]
```

```
lab@Sun# set protocols ldp session 172.30.5.4 authentication-key
workbook
```

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@Sun# 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]
lab@Sun# set protocols ldp explicit-null
```

```
[edit]
lab@Sun# 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.

After committing and adding the necessary LDP configuration to all routers you can verify the LDP sessions.

```
[edit]
lab@Sun# 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@Sun# 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@Sun# 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
  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.

```
[edit]
lab@Sirius# 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
}
```

Activate the **explicit-null** feature will instruct LDP on R2 to announce a label value of 0 for its Loopback IP address.

```
[edit]
lab@Sirius# set protocols ldp explicit-null

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

Verify the LDP database on R1 again.

```
[edit]
lab@Sun# 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.

```
[edit]
lab@Sun# 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
                   > 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@Sun# 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.

3) Configure IS-IS LDP synchronization.

This step requires to configure the LDP IGP synchronization feature. In this network ISIS 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@Sun# 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 150;
        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 ISIS interface type must be set to point-to-point.

NOTE: LDP synchronization is supported only on point-to-point interfaces.

```
[edit]
lab@Sun# set protocols isis interface ge-0/0/4.114 point-to-point
```

```
[edit]
lab@Sun# set protocols isis interface ge-0/0/4.114 bfd-liveness-
detection minimum-interval 150
```

```
[edit]
lab@Sun# set protocols isis interface ge-0/0/4.114 bfd-liveness-
detection multiplier 3
```

```
[edit]
lab@Sun# set protocols isis interface ge-0/0/4.114 ldp-
synchronization
```

```
[edit]
lab@Sun# set protocols isis interface ae0.0 point-to-point
```

```
[edit]
lab@Sun# set protocols isis interface ae0.0 bfd-liveness-detection
minimum-interval 150
```

```
[edit]
```

```
lab@Sun# set protocols isis interface ae0.0 bfd-liveness-detection
multiplier 3
```

```
[edit]
```

```
lab@Sun# set protocols isis interface ae0.0 ldp-synchronization
```

```
[edit]
```

```
lab@Sun# commit
commit complete
```

Display the IS-IS interface details.

```
[edit]
lab@Sun# run show isis interface ae0.0 extensive
IS-IS interface database:
ae0.0
  Index: 69, State: 0x6, Circuit id: 0x1, Circuit type: 1
  LSP interval: 100 ms, CSNP interval: 20 s, Loose Hello padding
  Adjacency advertisement: Advertise
  LDP sync state: in sync, for: 00:00:18, reason: LDP up during
  config
  config holdtime: infinity
  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 a LDP egress policy. Where you will match the routes in inet.0 for which you want to build a 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.

```
[edit]
lab@Sun# set policy-options policy-statement ldp-routes term 1
from route-filter 192.168.1.0/24 exact
```

```
[edit]
lab@Sun# set policy-options policy-statement ldp-routes term 1
from route-filter 172.30.5.1/32 exact
```

```
[edit]
lab@Sun# set policy-options policy-statement ldp-routes term 1
then accept
```

```
[edit]
lab@Sun# set protocols ldp egress-policy ldp-routes

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

You have to 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@Sun# 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
  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 have to 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 have 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@Arcturus# 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

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.

```
[edit]
lab@Sun# set protocols ldp deaggregate

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

Verify again the LDP database on R4.

```
[edit]
lab@Arcturus# 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
  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 a 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 constraints or can follow the IGP path. Mandatory and optional objects are carried by the Path message.
- *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 an included color
- 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)

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.

```
[edit]
lab@Sun# set protocols rsvp interface ge-0/0/4.114 authentication-
key workbook
```

```
[edit]
lab@Sun# set protocols rsvp interface ge-0/0/4.114 bandwidth 333m
```

```
[edit]
lab@Sun# set protocols rsvp interface ge-0/0/4.118 authentication-
key workbook
```

```
[edit]
lab@Sun# set protocols rsvp interface ge-0/0/4.118 bandwidth 333m
```

```
[edit]
lab@Sun# set protocols rsvp interface ae0.0 authentication-key
workbook
```

Verify the RSVP interface database.

```
[edit]
lab@Sun# run show rsvp interface
RSVP interface: 3 active
```

	Active	Subscr-	Static	Available	Reserved	
Highwater						
Interface	State	resv	ption	BW	BW	mark
ae0.0	Up	0	100%	2Gbps	2Gbps	0bps
ge-0/0/4.114	Up	0	100%	333Mbps	333Mbps	0bps
ge-0/0/4.118	Up	0	100%	333Mbps	333Mbps	0bps

You can see that the Static bandwidth on the ge- interfaces is changed to the configured value of 333Mbps. 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 HelloInterval and Protection.

```
[edit]
lab@Sun# 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 a LSP must be signaled/ configured before this show command will display output.

2) Configure administrative groups on all routers.

In Table 9 and Figure 15 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 not 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 have to define a name for the group and assign a affinity bit to it. Different groups can be assigned to interfaces. Usually colors are used for names of the groups.

For this task you have to configure green and red interfaces. The group configuration must be identical on all routers in the MPLS network.

```
[edit]
lab@Sun# set protocols mpls admin-groups green 0
```

```
[edit]
lab@Sun# set protocols mpls admin-groups red 1
```

Next you have to assign the groups to the interfaces according to Figure 15. Previously you already configured **protocols mpls - interfaces all**. However now you have to make the statement more specific as for each interface a different group might be needed.

```
[edit]
```

```

lab@Sun# set protocols mpls interface ge-0/0/4.114 admin-group
green

[edit]
lab@Sun# set protocols mpls interface ge-0/0/4.118 admin-group red

[edit]
lab@Sun# set protocols mpls interface ae0.0 admin-group [ green
red ]

[edit]
lab@Sun# commit
commit complete

```

Interface **ae0.0** is assigned with both groups.

Verify the administrative group on the interfaces.

```

[edit]
lab@Sun# 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 9 and Figure 15.

After Configure all routers you can check the traffic-engineering database (TED). Traffic engineering support is enabled by default for ISIS.

```

[edit]
lab@Sun# run show ted database extensive
TED database: 8 ISIS nodes 8 INET nodes
NodeID: Sun.00(172.30.5.1)
  Type: Rtr, Age: 37 secs, LinkIn: 3, LinkOut: 3
  Protocol: IS-IS(1)
    To: Procyon.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: Arcturus.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
    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: Sirius.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: Sirius.00(172.30.5.2)
  Type: Rtr, Age: 37 secs, LinkIn: 3, LinkOut: 3
  Protocol: IS-IS(1)
    To: Sun.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: Rigel.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
    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: Canopus.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 Configure administrative groups.

```
[edit]
lab@Sirius# run show isis database level 1 Sun.00-00 extensive
IS-IS level 1 link-state database:

Sun.00-00 Sequence: 0x1b0, Checksum: 0x7eda, Lifetime: 1188 secs
  IPV4 Unicast IS neighbor: Sirius.00      Metric:      5
    Two-way fragment: Sirius.00-00, Two-way first fragment:
Sirius.00-00
  IPV4 Unicast IS neighbor: Arcturus.00    Metric:      10
    Two-way fragment: Arcturus.00-00, Two-way first fragment:
Arcturus.00-00
  IPV4 Unicast IS neighbor: Procyon.00     Metric:      10
    Two-way fragment: Procyon.00-00, Two-way first fragment:
Procyon.00-00
  IPV4 Unicast IS neighbor: route-reflector.00 Metric:      10
    Two-way fragment: route-reflector.00-00, Two-way first
fragment: route-reflector.00-00
  IPV6 Unicast IS neighbor: Sirius.00      Metric:      5
    Two-way fragment: Sirius.00-00, Two-way first fragment:
Sirius.00-00
  IPV6 Unicast IS neighbor: Procyon.00    Metric:      10
    Two-way fragment: Procyon.00-00, Two-way first fragment:
Procyon.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: Sun.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: Sun.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: Sun
IS extended neighbor: Sirius.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
  Maximum reservable bandwidth: 2Gbps
  Maximum bandwidth: 2Gbps
  Administrative groups: 0x3 red green
IS extended neighbor: Arcturus.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: Procyon.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: Sirius.00, Metric: default 5
IPV6 UnicastIS neighbor: Procyon.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
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 have to configure the RSVP signaled LSPs according Table 10 and Figure 16. CSPF will be used to create the ERO. BFD must to be configured as well.

a. R1.

```

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon to
172.30.5.8

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon oam
bfd-liveness-detection minimum-interval 300

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Vega to
172.30.5.6

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Vega oam
bfd-liveness-detection minimum-interval 300

[edit]
lab@Sun# commit
commit complete

```

RSVP signaled LSP are configured with **label-switched-path** under **family mpls** stanza. 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@Sun# 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 Sun-to-
Vega
172.30.5.8  172.30.5.1  Up    0   1 FF    -          3 Sun-to-
Procyon
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@Sun# run show rsvp session extensive
Ingress RSVP: 2 sessions

172.30.5.6
  From: 172.30.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Sun-to-Vega, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  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: Sun-to-Procyon, 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. **Explct route** is showing the calculated / configured ERO. **Record route** is the path that the Path message traverses. **Label out** and **Label in** are shown as well.

Another command which is useful to receive information about the RSVP signaled LSP is the **show mpls lsp**.

```

[edit]
lab@Sun# 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	*		Sun-to-Vega
172.30.5.8	172.30.5.1	Up	0	*		Sun-to-Procyon

```

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@Sun# 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
Sun-to-Vega	
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 Sun-to-Procyon	

```

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 via 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@Sun# run show mpls lsp extensive
Ingress LSP: 2 sessions

172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
Vega
  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: Sun-to-
Procyon
  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.

However the BFD session are DOWN. Let's verify the BFD sessions table.

```

[edit]
lab@Sun# run show bfd session

```

Transmit Address	State	Interface	Time	Detect Interval	
Multiplier					
127.0.0.1	Down	ae0.0	0.000	1.000	3
127.0.0.1	Down	ge-0/0/4.118	0.000	1.000	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
172.30.5.41	UP	ge-0/0/4.206	0.450	0.150	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@Sun# 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 for Verify the data path of the LSPs. For proper operation of MPLS ping you have to 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 has to be configured on the Loopback interface on all routers. Secondly the filter protecting the RE must be changed to accept the MPLS ping packets.

```

[edit]

```

```
lab@Sun# set interfaces lo0.0 family inet address 127.0.0.1/32

[edit]
lab@Sun# set interfaces lo0.0 family inet address 172.30.5.1/32
primary
```

Making the address primary ensures that this IP address would be used for sourcing packets from the RE.

```
[edit]
lab@Procyon# set firewall family inet filter protect-re term lsp-
ping from protocol udp

[edit]
lab@Procyon# set firewall family inet filter protect-re term lsp-
ping from port 3503

[edit]
lab@Procyon# set firewall family inet filter protect-re term lsp-
ping then policer re-policer

[edit]
lab@Procyon# set firewall family inet filter protect-re term lsp-
ping then accept

[edit]
lab@Procyon# insert firewall family inet filter protect-re term
lsp-ping before term last

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

The new term has to 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.

```
[edit]
lab@Sun# run show mpls lsp name Sun-to-Procyon extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
Procyon
  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)
  172.30.0.2 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@Sun# run show bfd session

```

Address Multiplier	State	Interface	Detect Time	Transmit Interval	
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@Sun# run ping mpls rsvp Sun-to-Procyon
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss

```

b. R2.

```

[edit]

```



```

lab@Sirius# set protocols mpls label-switched-path Sirius-to-A-
Centauri to 172.30.5.5

[edit]
lab@Sirius# set protocols mpls label-switched-path Sirius-to-A-
Centauri oam bfd-liveness-detection minimum-interval 300

[edit]
lab@Sirius# set protocols mpls label-switched-path Sirius-to-Rigel
to 172.30.5.7

[edit]
lab@Sirius# set protocols mpls label-switched-path Sirius-to-Rigel
oam bfd-liveness-detection minimum-interval 300

[edit]
lab@Sirius# commit
commit complete

```

Verify the RSVP session for the configured LSPs.

```

[edit]
lab@Sirius# 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 Sirius-
to-A-Centauri
172.30.5.7   172.30.5.2   Up    0  1 FF      -          3 Sirius-
to-Rigel
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 Sun-to-
Vega
Total 1 displayed, Up 1, Down 0

```

Both session are UP and signaled successfully. On 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.

```

lab@Sirius# 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  *           Sirius-
to-A-Centauri
172.30.5.7   172.30.5.2   Up    0  *           Sirius-
to-Rigel
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  Sun-to-
Vega
Total 1 displayed, Up 1, Down 0
```

```
[edit]
lab@Sirius# 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
Sirius-to-A-Centauri
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 Sirius-to-Rigel
```

c. R3.

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Vega to 172.30.5.6
```

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Vega oam bfd-liveness-detection minimum-interval 300
```

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-1 to 172.30.5
```

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-1 oam bfd-liveness-detection minimum-interval 300
```

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-2 to 172.30.5.8
```

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-2 oam bfd-liveness-detection minimum-interval 300
```

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

R3 has two LSPs to R8. Listing the RSVP session and MPLS LSP table.

```
[edit]
lab@Canopus# 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
Canopus-to-Vega
172.30.5.8   172.30.5.3  Up       0   1 FF      -      299840
Canopus-to-Procyon-1
172.30.5.8   172.30.5.3  Up       0   1 FF      -      299856
Canopus-to-Procyon-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@Canopus# 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 *          Canopus-to-Vega
172.30.5.8   172.30.5.3   Up    0 *          Canopus-to-
Procyon-1
172.30.5.8   172.30.5.3   Up    0 *          Canopus-to-
Procyon-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.

```
[edit]
lab@Canopus# 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 Canopus-to-Vega
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 Canopus-to-Procyon-1
                  to 172.30.0.13 via ge-0/0/4.123, label-
switched-path Canopus-to-Procyon-2
```

d. R4.

```
[edit]
lab@Arcturus# set protocols label-switched-path Arcturus-to-Rigel-
1 to 172.30.5.7
```

```
[edit]
lab@Arcturus# set protocols label-switched-path Arcturus-to-Rigel-
1 oam bfd-liveness-detection minimum-interval 300
```

```
[edit]
lab@Arcturus# set protocols label-switched-path Arcturus-to-Rigel-
2 to 172.30.5.7
```

```
[edit]
lab@Arcturus# set protocols label-switched-path Arcturus-to-Rigel-
2 oam bfd-liveness-detection minimum-interval 300
```

```
[edit]
```

```
lab@Arcturus# set protocols label-switched-path Arcturus-to-A-Centauri to 172.30.5.5
```

```
[edit]
```

```
lab@Arcturus# set protocols label-switched-path Arcturus-to-A-Centauri oam bfd-liveness-detection minimum-interval 300
```

```
[edit]
```

```
lab@Arcturus# commit
commit complete
```

R4 has two LSPs to R7 Loopback IP address.

```
[edit]
lab@Arcturus# run show rsvp session
Ingress RSVP: 3 sessions
To          From          State   Rt  Style Labelin Labelout LSPname
172.30.5.5   172.30.5.4   Up      0   1 FF      -       3
Arcturus-to-A-Centauri
172.30.5.7   172.30.5.4   Up      0   1 FF      -    299872
Arcturus-to-Rigel-1
172.30.5.7   172.30.5.4   Up      0   1 FF      -    299888
Arcturus-to-Rigel-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@Arcturus# 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 *           Arcturus-
to-A-Centauri
172.30.5.7   172.30.5.4   Up    0 *           Arcturus-
to-Rigel-1
172.30.5.7   172.30.5.4   Up    0 *           Arcturus-
to-Rigel-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@Arcturus# 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 Arcturus-to-A-Centauri
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 Arcturus-to-Rigel-1
                   to 172.30.0.30 via ge-0/0/4.145, label-
switched-path Arcturus-to-Rigel-2

```

e. R5.

```

[edit]
lab@A-Centauri# set protocols mpls label-switched-path A-Centauri-
to-Arcturus to 172.30.5.4

```

```

[edit]
lab@A-Centauri# set protocols mpls label-switched-path A-Centauri-
to-Arcturus oam bfd-liveness-detection minimum-interval 300

```

```

[edit]
lab@A-Centauri# set protocols mpls label-switched-path A-Centauri-
to-Sirius to 172.30.5.2

```

```

[edit]
lab@A-Centauri# set protocols mpls label-switched-path A-Centauri-
to-Sirius oam bfd-liveness-detection minimum-interval 300

```

```

[edit]
lab@A-Centauri# commit
commit complete

```

Here you can see that there are Ingress and Egress LSPs. LSPs are unidirectional, this means that a LSP should be configured in both directions for bidirectional communication.

```

[edit]
lab@A-Centauri# 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 A-
Centauri-to-Sirius
172.30.5.4   172.30.5.5   Up      0   1 FF    -        3 A-
Centauri-to-Arcturus
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        -
Sirius-to-A-Centauri
172.30.5.5   172.30.5.4   Up      0   1 FF    3        -
Arcturus-to-A-Centauri
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
Arcturus-to-Rigel-2						
Total 1 displayed, Up 1, Down 0						

R5 is advertising an Implicit Null label 3.

```
[edit]
lab@A-Centauri# 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 *          A-
Centauri-to-Sirius
172.30.5.4  172.30.5.5  Up    0 *          A-
Centauri-to-Arcturus
Total 2 displayed, Up 2, Down 0

Egress LSP: 2 sessions
To          From          State Rt Style Labelin Labelout
LSPname
172.30.5.5  172.30.5.2  Up    0  1 FF      3      -
Sirius-to-A-Centauri
172.30.5.5  172.30.5.4  Up    0  1 FF      3      -
Arcturus-to-A-Centauri
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
Arcturus-to-Rigel-2
Total 1 displayed, Up 1, Down 0
```

f. R6.

```
[edit]
lab@Vega# set protocols mpls label-switched-path Vega-to-Sun to
172.30.5.1

[edit]
lab@Vega# set protocols mpls label-switched-path Vega-to-Sun oam
bfd-liveness-detection minimum-interval 300

[edit]
lab@Vega# set protocols mpls label-switched-path Vega-to-Canopus
to 172.30.5.3

[edit]
lab@Vega# set protocols mpls label-switched-path Vega-to-Canopus
oam bfd-liveness-detection minimum-interval 300

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

Verify the RSVP sessions and MPLS LSP table.

```
[edit]
lab@Vega# 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  Vega-
to-Sun
172.30.5.3  172.30.5.6  Up       0   1 FF    -          3  Vega-
to-Canopus
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         -        Canopus-to-Vega
172.30.5.6  172.30.5.1  Up       0   1 FF    3         -        Sun-to-
Vega
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    299856    3        Arcturus-to-Rigel-2
Total 1 displayed, Up 1, Down 0

[edit]
lab@Vega# 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   *      Vega-to-Sun
172.30.5.3  172.30.5.6  Up       0   *      Vega-to-
Canopus
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         -        Canopus-to-Vega
172.30.5.6  172.30.5.1  Up       0   1 FF    3         -        Sun-to-
Vega
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        Arcturus-to-Rigel-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@Vega# 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 Vega-to-Sun
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 Vega-to-Canopus
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.

```
[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-Sirius
to 172.30.5.2
```

```
[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-Sirius
oam bfd-liveness-detection minimum-interval 300
```

```
[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-
Arcturus-1 to 172.30.5.4
```

```
[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-
Arcturus-1 oam bfd-liveness-detection minimum-interval 300
```

```
[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-
Arcturus-2 to 172.30.5.4
```

```
[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-
Arcturus-2 oam bfd-liveness-detection minimum-interval 300
```

```
[edit]
lab@Rigel# 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@Rigel# 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 Rigel-
to-Sirius
172.30.5.4 172.30.5.7 Up 0 1 FF - 299872 Rigel-
to-Arcturus-1
172.30.5.4 172.30.5.7 Up 0 1 FF - 299888 Rigel-
to-Arcturus-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 - Sirius-
to-Rigel
172.30.5.7 172.30.5.4 Up 0 1 FF 3 -
Arcturus-to-Rigel-1
172.30.5.7 172.30.5.4 Up 0 1 FF 3 -
Arcturus-to-Rigel-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 Sun-to-
Vega
Total 1 displayed, Up 1, Down 0

```

```

[edit]
lab@Rigel# 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 * Rigel-to-
Sirius
172.30.5.4 172.30.5.7 Up 0 * Rigel-to-
Arcturus-1
172.30.5.4 172.30.5.7 Up 0 * Rigel-to-
Arcturus-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 - Sirius-
to-Rigel
172.30.5.7 172.30.5.4 Up 0 1 FF 3 -
Arcturus-to-Rigel-1
172.30.5.7 172.30.5.4 Up 0 1 FF 3 -
Arcturus-to-Rigel-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 Sun-to-
Vega
Total 1 displayed, Up 1, Down 0

```

You can see that the LSP to R2 is configured for bidirectional connectivity.

```
[edit]
```

```
lab@Rigel# 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 Rigel-to-Sirius
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 Rigel-to-Arcturus-1
                  > to 172.30.0.41 via ge-0/0/4.167, label-
switched-path Rigel-to-Arcturus-2
```

h. R8.

```
[edit]
lab@Procyon# set protocols mpls label-switched-path Procyon-to-
Canopus-1 to 172.30.5.3

[edit]
lab@Procyon# set protocols mpls label-switched-path Procyon-to-
Canopus-1 oam bfd-liveness-detection minimum-interval 300

[edit]
lab@Procyon# set protocols mpls label-switched-path Procyon-to-
Canopus-2 to 172.30.5.3

[edit]
lab@Procyon# set protocols mpls label-switched-path Procyon-to-
Canopus-2oam bfd-liveness-detection minimum-interval 300

[edit]
lab@Procyon# set protocols mpls label-switched-path Procyon-to-Sun
to 172.30.5.1

[edit]
lab@Procyon# set protocols mpls label-switched-path Procyon-to-Sun
oam bfd-liveness-detection minimum-interval 300

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

Verify the RSVP session table and MPLS LSP table.

```
[edit]
lab@Procyon# 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
Procyon-to-Sun
172.30.5.3  172.30.5.8  Up      0   1 FF    -        299952
Procyon-to-Canopus-1
172.30.5.3  172.30.5.8  Up      0   1 FF    -        299904
Procyon-to-Canopus-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        -
Canopus-to-Procyon-1
172.30.5.8  172.30.5.3  Up      0   1 FF      3        -
Canopus-to-Procyon-2
172.30.5.8  172.30.5.1  Up      0   1 FF      3        - Sun-to-
Procyon
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 Vega-
to-Sun
172.30.5.2  172.30.5.5  Up      0   1 FF  299840  299888 A-
Centauri-to-Sirius
172.30.5.5  172.30.5.2  Up      0   1 FF  299824      3 Sirius-
to-A-Centauri
Total 3 displayed, Up 3, Down 0
```

```
lab@Procyon# run show mpls lsp
Ingress LSP: 3 sessions
To          From          State   Rt  P      ActivePath      LSPname
172.30.5.1  172.30.5.8  Up      0  *      Procyon-
to-Sun
172.30.5.3  172.30.5.8  Up      0  *      Procyon-
to-Canopus-1
172.30.5.3  172.30.5.8  Up      0  *      Procyon-
to-Canopus-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        -
Canopus-to-Procyon-1
172.30.5.8  172.30.5.3  Up      0   1 FF      3        -
Canopus-to-Procyon-2
172.30.5.8  172.30.5.1  Up      0   1 FF      3        - Sun-to-
Procyon
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 Vega-
to-Sun
172.30.5.2  172.30.5.5  Up      0   1 FF  299840  299888 A-
Centauri-to-Sirius
172.30.5.5  172.30.5.2  Up      0   1 FF  299824      3 Sirius-
to-A-Centauri
Total 3 displayed, Up 3, Down 0
```

Again on R8 all LSPs are configured to ensure bidirectional paths to R1 and R3.

```
[edit]
lab@Procyon# 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 Procyon-to-Sun
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 Procyon-to-Canopus-1
                  > to 172.30.0.9 via ge-0/0/4.118, label-
switched-path Procyon-to-Canopus-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 a LSP to R8 should include links which have the green color assigned.

```
[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon
admin-group include-any green

[edit]
lab@Sun# 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@Sun# run show mpls lsp name Sun-to-Procyon extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
Procyon
  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
  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@Sirius# set protocols mpls label-switched-path Sirius-to-Rigel
admin-group include-any red

[edit]
lab@Sirius# commit
commit complete

```

Verify the LSP parameters.

```

[edit]
lab@Sirius# run show mpls lsp name Sirius-to-Rigel extensive

```

```

Ingress LSP: 2 sessions

172.30.5.7
  From: 172.30.5.2, State: Up, ActiveRoute: 0, LSPname: Sirius-to-
Rigel
  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
  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@Canopus# set protocols mpls label-switched-path Canopus-to-
Vega admin-group include-any red

```

```

[edit]
lab@Canopus# commit
commit complete

```

d. R4.

```

[edit]
lab@Arcturus# set protocols mpls label-switched-path Arcturus -to-
A-Centauri admin-group include-any green

```

```

[edit]
lab@Arcturus# commit

```

```
commit complete
```

e. R5.

```
[edit]
lab@A-Centauri# set protocols mpls label-switched-path A-Centauri-
to-Arcturus admin-group include-any green
```

```
[edit]
lab@A-Centauri# commit
commit complete
```

f. R6.

```
[edit]
lab@Vega# set protocols mpls label-switched-path Vega-to-Canopus
admin-group include-any red
```

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

g. R7.

```
[edit]
lab@Rigel# set protocols mpls label-switch-path Ridel-to-Sirius
admin-group include-any green
```

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

h. R8.

```
[edit]
lab@Procyon# set protocols mpls label-switch-path Procyon-to-Sun
admin-group include-any green
```

```
[edit]
lab@Procyon# 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 inherit from the **label-switched-path** configuration.


```

[edit]
lab@Canopus# set protocols mpls path path-1 172.30.5.2

[edit]
lab@Canopus# set protocols mpls path path-1 172.30.5.1

[edit]
lab@Canopus# set protocols mpls path path-1 172.30.5.8

[edit]
lab@Canopus# set protocols mpls path path-2 172.30.5.6

[edit]
lab@Canopus# set protocols mpls path path-2 172.30.5.7

[edit]
lab@Canopus# set protocols mpls path path-2 172.30.5.8

[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-1 primary path-1

[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-2 primary path-2

[edit]
lab@Canopus# 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@Canopus# run show mpls lsp name Canopus-to-Procyon-1 extensive
Ingress LSP: 3 sessions

172.30.5.8
  From: 172.30.5.3, State: Up, ActiveRoute: 0, LSPname: Canopus-
to-Procyon-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)
  172.30.0.13 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@Canopus# run show mpls lsp name Canopus-to-Procyon-2 extensive
Ingress LSP: 3 sessions

172.30.5.8
From: 172.30.5.3, State: Up, ActiveRoute: 0, LSPname: Canopus-
to-Procyon-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.

```

[edit]
lab@Procyon# set protocols mpls path path-1 172.30.5.1

```

```

[edit]
lab@Procyon# set protocols mpls path path-1 172.30.5.2

```

```

[edit]
lab@Procyon# set protocols mpls path path-1 172.30.5.3

```

```
[edit]
lab@Procyon# set protocols mpls path path-2 172.30.5.5

[edit]
lab@Procyon# set protocols mpls path path-2 172.30.5.4

[edit]
lab@Procyon# set protocols mpls path path-2 172.30.5.3

[edit]
lab@Procyon# set protocols mpls label-switched-path Procyon-to-Canopus-1 primary path-1

[edit]
lab@Procyon# set protocols mpls label-switched-path Procyon-to-Canopus-2 primary path-2

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

Verify the route of the LSPs.

```
[edit]
lab@Procyon# run show mpls lsp name Procyon-to-Canopus-1 extensive
Ingress LSP: 3 sessions

172.30.5.3
  From: 172.30.5.8, State: Up, ActiveRoute: 0, LSPname: Procyon-
to-Canopus-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
    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@Procyon# run show mpls lsp name Procyon-to-Canopus-2 extensive
Ingress LSP: 3 sessions

172.30.5.3
  From: 172.30.5.8, State: Up, ActiveRoute: 0, LSPname: Procyon-
to-Canopus-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 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.

```

[edit]
lab@Arcturus# set protocols mpls path path-1 172.30.5.3

[edit]
lab@Arcturus# set protocols mpls path path-2 172.30.5.5

[edit]
lab@Arcturus# set protocols mpls label-switch-path Arcturus-to-
Rigel-1 admin-group include-any red

[edit]

```

```
lab@Arcturus# set protocols mpls label-switch-path Arcturus-to-Rigel-1 primary path-1
```

```
[edit]
```

```
lab@Arcturus# set protocols mpls label-switch-path Arcturus-to-Rigel-2 admin-group include-any red
```

```
[edit]
```

```
lab@Arcturus# set protocols mpls label-switch-path Arcturus-to-Rigel-2 primary path-2
```

```
[edit]
```

```
lab@Arcturus# commit
```

```
commit complete
```

Verify the LSPs paths.

```
[edit]
lab@Arcturus# run show mpls lsp name Arcturus-to-Rigel-1 extensive
Ingress LSP: 3 sessions

172.30.5.7
  From: 172.30.5.4, State: Up, ActiveRoute: 0, LSPname: Arcturus-to-Rigel-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
  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@Arcturus# run show mpls lsp name Arcturus-to-Rigel-2 extensive
Ingress LSP: 3 sessions
```

```

172.30.5.7
  From: 172.30.5.4, State: Up, ActiveRoute: 0, LSPName: Arcturus-
to-Rigel-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.

```

[edit]
lab@Rigel# set protocols mpls path path-1 172.30.5.2

[edit]
lab@Rigel# set protocols mpls path path-2 172.30.5.8

[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-
Arcturus-1 admin-group include-any green

[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-
Arcturus-1 primary path-1

[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-
Arcturus-2 admin-group include-any green

[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-
Arcturus-2 primary path-2

[edit]
lab@Rigel# commit
commit complete

```

Verify the changed path of both LSPs.

```
[edit]
lab@Rigel# run show mpls lsp name Rigel-to-Arcturus-1 extensive
Ingress LSP: 3 sessions

172.30.5.4
  From: 172.30.5.7, State: Up, ActiveRoute: 0, LSPname: Rigel-to-
Arcturus-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
```

```

[edit]
lab@Rigel# run show mpls lsp name Rigel-to-Arcturus-2 extensive
Ingress LSP: 3 sessions

172.30.5.4
  From: 172.30.5.7, State: Up, ActiveRoute: 0, LSPName: Rigel-to-
Arcturus-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 makes a decision 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 R1 LSPs with bandwidth reservation of 60Mbps. Configuration is simple and straightforward.

```

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Vega
bandwidth 60m

```

```

[edit]
lab@Sun# commit
commit complete

```

Verify the LSP after the commit.


```
[edit]
lab@Sun# run show mpls lsp name Sun-to-Vega extensive
Ingress LSP: 2 sessions

172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
Vega
  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 are able to 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 resigaled using the make-before-break feature. This ensures that traffic is not dropped during resignaling of the LSP.

Configure autobandwidth on the LSP from R1 to R8.

```
[edit]
lab@Sun# set protocols mpls statistics file Auto-Bandwidth

[edit]
lab@Sun# set protocols mpls statistics auto-bandwidth

[edit]
```

```
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon
auto-bandwidth adjust-interval 172800
```

```
[edit]
```

```
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon
auto-bandwidth minimum-bandwidth 30m
```

```
[edit]
```

```
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon
auto-bandwidth maximum-bandwidth 120m
```

```
[edit]
```

```
lab@Sun# 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.

Verify the result of the autobandwidth LSP configuration.

```
[edit]
lab@Sun# run show mpls lsp name Sun-to-Procyon extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
  Procyon
  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 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 on the LSP from R1 to R8.

```
[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon
priority 6 6
```

```
[edit]
lab@Sun# commit
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@Sun# run show mpls lsp name Sun-to-Procyon extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
Procyon
  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.

```

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Vega
priority 7 7

```

```

[edit]
lab@Sun# commit
commit complete

```

```

[edit]
lab@Sun# run show mpls lsp name Sun-to-Vega extensive
Ingress LSP: 2 sessions

172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
Vega
  ActivePath: (primary)
  LSPTtype: 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
    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 a LSP which is preempted is torn down before a new path is calculated. For example if a 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.

```

[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-2 soft-preemption

```

```

[edit]
lab@Canopus# commit
commit complete

```

Verify if soft-preemption is signaled.

```
[edit]
lab@Canopus# run show rsvp session name Canopus-to-Procyon-2
extensive
Ingress RSVP: 8 sessions

172.30.5.8
  From: 172.30.5.3, LSPstate: Up, ActiveRoute: 0
  LSPname: Canopus-to-Procyon-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 a 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.

Another requirement for the LSPs is to enable the make-before-break 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 on R3.

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-1 optimize-timer 28800
```

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-1 adaptive
```

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-2 optimize-timer 28800
```

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-2 adaptive
```

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

After committing, both LSPs should be recomputed and the optimization period is taken into account.

```
[edit]
lab@Canopus# run show mpls lsp name Canopus-to-Procyon-1 extensive
Ingress LSP: 3 sessions

172.30.5.8
  From: 172.30.5.3, State: Up, ActiveRoute: 0, LSPname: Canopus-
to-Procyon-1
  ActivePath: path-1 (primary)
  LSPTtype: 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) (Sun.00/172.30.5.1)-
>172.30.0.10(172.30.5.8:71) (Procyon.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: Sun Sep 22 10:34:08 2013
Total 1 displayed, Up 1, Down 0

```

Reoptimization should take place in 27344 seconds.

```

[edit]
lab@Canopus# 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 Canopus-
to-Vega
172.30.5.6  172.30.5.3  Up     0   1 FF      -    300640 Canopus-
to-Vega
172.30.5.8  172.30.5.3  Up     0   1 SE      -    301216 Canopus-
to-Procyon-1
172.30.5.8  172.30.5.3  Up     0   1 SE      -    300464 Canopus-
to-Procyon-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 knob** 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 have to 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@A-Centauri# 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
            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

[edit]

```

```

lab@A-Centauri# set protocols mpls label-switched-path A-Centauri-
to-Sirius install 192.168.1.0/24

```

```

[edit]
lab@A-Centauri# commit
commit complete

```

Verify if the route is added to inet.3

```

[edit]
lab@A-Centauri# 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 A-Centauri-to-Sirius

```

Verify if the next-hop of prefix advertised from IX peers is now resolved through the LSP.

```

[edit]
lab@A-Centauri# 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: 672
              Next hop: 172.30.0.29 via ge-0/0/4.145 weight 0x1,
selected
              Label-switched-path A-Centauri-to-Sirius
              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.

```

[edit]
lab@Vega# set protocols mpls label-switched-path Vega-to-Sun
install 192.168.1.0/24

[edit]
lab@Vega# commit
commit complete

```

Verify if prefix 192.168.1.0/24 is present in inet.3

```

[edit]
lab@Vega# 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 Vega-to-Sun

```

- 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 routers loopback interfaces. LDP tunneling over the RSVP LSP ensures that LDP LSPs are signaled between each and every router.

First you have to enable LDP on Loopback interface on all routers.

```
[edit]
lab@Sun# set protocols ldp interface lo0.0
```

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

Next you have to 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.

```
[edit]
lab@Sun# set protocols label-switched-path Sun-to-Procyon ldp-
tunneling
```

```
[edit]
lab@Sun# 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@Canopus# 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 a 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@Canopus# 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 Canopus-to-Vega
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-
path Canopus-to-Vega
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 Canopus-to-Vega
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-
path Bypass->172.30.0.22
                  [LDP/9] 00:05:11, metric 10
                  > to 172.30.0.22 via ge-0/0/4.134, label-
switched-path Canopus-to-Vega
                  to 172.30.0.13 via ge-0/0/4.123, label-
switched-path Canopus-to-Vega
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 Canopus-to-Vega
                  to 172.30.0.13 via ge-0/0/4.123, label-switched-
path Canopus-to-Vega
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 Canopus-to-Procyon-1
                  to 172.30.0.26 via ge-0/0/4.136, label-switched-
path Canopus-to-Procyon-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] 00:00:41, 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@Canopus# 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 Canopus-to-Vega
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 Canopus-to-Vega
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 Canopus-to-Vega
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 Canopus-to-Vega
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 Canopus-to-Vega
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 Canopus-to-Vega
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 Canopus-to-Vega
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 Canopus-to-Vega
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 Canopus-to-Vega
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 Canopus-to-Procyon-1
> to 172.30.0.26 via ge-0/0/4.136, label-switched-
path Canopus-to-Procyon-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 have to 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 have to 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@Procyon# 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 Procyon-to-Canopus-1
                    to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
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 Procyon-to-Canopus-1
                    > to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-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 Procyon-to-Canopus-1
                    to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
26.0.0.0/8         *[BGP/170] 00:18:57, localpref 200, from
172.30.5.41
    --- (more) ---

[edit]
lab@Procyon# 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 Procyon-to-Canopus-1
                  > to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-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 Procyon-to-Canopus-1
                  > to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-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 Procyon-to-Canopus-1
                  to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-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@Procyon# set policy-options policy-statement lsp-map term 1
from protocol bgp

[edit]
lab@Procyon# set policy-options policy-statement lsp-map term 1
from community P2

[edit]
lab@Procyon# set policy-options policy-statement lsp-map term 1
then install-nexthop lsp Procyon-to-Canopus-1

[edit]
lab@Procyon# set policy-options policy-statement lsp-map term 2
from protocol bgp

[edit]
lab@Procyon# set policy-options policy-statement lsp-map term 2
from community P3

[edit]
lab@Procyon# set policy-options policy-statement lsp-map term 2
then install-nexthop lsp Procyon-to-Canopus-2

```

Apply the policy to the forwarding-table.


```
[edit]
lab@Procyon# set routing-options forwarding-table export lsp-map

[edit]
lab@Procyon# 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@Procyon# 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 Procyon-to-Canopus-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 Procyon-to-Canopus-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 Procyon-to-Canopus-1
26.0.0.0/8      *[BGP/170] 00:01:59, localpref 200, from
172.30.5.41
                AS path: 3521382357 46848 I
                to 172.30.0.9 via ge-0/0/4.118, label-
switched-path Procyon-to-Canopus-1
--- (more) ---
```

```
[edit]
lab@Procyon# 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 Procyon-to-Canopus-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 Procyon-to-Canopus-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 Procyon-to-Canopus-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 Procyon-to-Canopus-2

```

16) Configure a per flow load balancing policy on R4 and R7.

This Step requires 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 Step).

Display the forwarding table for R7 Loopback IP address before applying the policy.

```

[edit]
lab@Arcturus# 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
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@Arcturus# 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 Arcturus-to-Rigel-1
    to 172.30.0.30 via ge-0/0/4.145, label-
switched-path Arcturus-to-Rigel-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.

```

[edit]
lab@Arcturus# set policy-options policy-statement load-balance
term 1 then load-balance per-packet

```

```

[edit]
lab@Arcturus# set routing-options forwarding-table export load-
balance

```

```

[edit]
lab@Arcturus# 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@Arcturus# 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          ulst 262144    4
                    ucst    586        24
ge-0/0/4.114
                    172.30.0.30        ucst    591        11
ge-0/0/4.145

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.

On R7 traffic to R4 should be load balanced for bidirectional communication. The same policy should be applied to the forwarding-table.

```
[edit]
lab@Rigel# set policy-options policy-statement load-balance term 1
then load-balance per-packet
```

```
[edit]
lab@Rigel# set routing-options forwarding-table export load-
balance
```

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

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.

Issue a traceroute with the default settings.

Don't forget 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@Sun# set protocols mpls traffic-engineering mpls-forwarding

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

That knob tells JUNOS to copy all inet.3 routes to inet.0.

```
[edit]
lab@Sun# run show route 172.30.5.5

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
Sun-to-Procyon

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
Sun-to-Procyon
```

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.

```
[edit]
lab@Sun# 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
```

It looks like a normal traceroute output. Every hop on the paths is shown together with the label values used. However when you configure the **no-decrement-ttl** feature, and issue again the same traceroute command the result is much different.

```
[edit]
lab@Sun# 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
```

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 egress PE.

Task 3: RSVP Protection

By default when network failures occurs 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 secondary paths. Paths are configured to use CSPF routes without strict ERO configuration.

```
[edit]
lab@Sun# set protocols mpls path primary-1

[edit]
lab@Sun# set protocols mpls path primary-2

[edit]
lab@Sun# set protocols mpls path secondary-1

[edit]
lab@Sun# set protocols mpls path secondary-2

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Vega
primary primary-1

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Vega
secondary secondary-1

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon
primary primary-2

[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon
secondary secondary-2

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

Verify the LSP to R8 after the changes are committed.

```
[edit]
lab@Sun# run show mpls lsp name Sun-to-Procyon extensive
```

```

Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
  Procyon
  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: Sun Sep 22 13:15:05 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 2 sessions
Total 0 displayed, Up 0, Down 0

```

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.

You have to configure the other LSPs, (except LSPs K, L, O and P) the same way.

- 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 aits 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 LSP from R1 to R6.

```
[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Vega
secondary secondary-1 standby
```

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

After committing the changes you verify the LSP to R6.

```
[edit]
lab@Sun# run show mpls lsp name Sun-to-Vega extensive
Ingress LSP: 2 sessions

172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
Vega
  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
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: Sun 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 particular 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.

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@Sun# run show rsvp session name Sun-to-Vega
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 Sun-to-
Vega
172.30.5.6  172.30.5.1  Up     0   1 FF      -    300560 Sun-to-
Vega
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.

```
[edit]
lab@Sun# set protocols mpls label-switched-path Sun-to-Vega
adaptive
```

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

Verify the reservation after the change.

```
[edit]
lab@Sun# run show rsvp session name Sun-to-Vega
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 Sun-to-
Vega
172.30.5.6  172.30.5.1  Up      0   1 SE      -    301296 Sun-to-
Vega
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 Configure a primary and a secondary path.

Configure revert-timer to 0 on R7 for the LSP towards R2

```
[edit]
lab@Rigel# set protocols mpls label-switched-path Rigel-to-Sirius
revert-timer 0
```

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

Verify the result of the changed revert-timer.

```
lab@Rigel# run show mpls lsp extensive name Rigel-to-Sirius
Ingress LSP: 3 sessions

172.30.5.2
  From: 172.30.5.7, State: Up, ActiveRoute: 0, LSPname: Rigel-to-Sirius
  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
  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) (Canopus.00/172.30.5.3)-
>172.30.0.13(172.30.5.2:73) (Sirius.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) (Sun.00/172.30.5.1) -
>172.30.0.2(172.30.5.2:69) (Sirius.00/172.30.5.2)
 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: Sun 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 are 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.

```

[edit]
lab@Sun# set protocol mpls label-switched-path Sun-to-Vega fast-
reroute hop-limit 5

```

```

[edit]
lab@Sun# set protocol mpls label-switched-path Sun-to-Vega fast-
reroute no-include-any

```

```

[edit]
lab@Sun# 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.

```
[edit]
lab@Sun# run show mpls lsp name Sun-to-Vega extensive
Ingress LSP: 6 sessions

172.30.5.6
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-Vega
  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]
  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 have to check a transit LSRs for that LSP. For both primary and secondary paths fast reroute detours are signaled. Let's for example take a 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@Canopus# run show mpls lsp transit name Sun-to-Vega extensive
Transit LSP: 28 sessions, 4 detours

172.30.5.6
  From: 172.30.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Sun-to-Vega, 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

172.30.5.6
  From: 172.30.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Sun-to-Vega, LSPpath: Secondary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 SE, Label in: 300304, 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 3 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) 22 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: 300160
Total 2 displayed, Up 2, Down 0

```

On R3 a detour is created through R4 to protect the direct link between R3 and R6.

```

[edit]
lab@Arcturus# run show mpls lsp transit name Sun-to-Vega extensive
Transit LSP: 29 sessions, 4 detours

172.30.5.6
  From: 172.30.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Sun-to-Vega, 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
    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: Sun-to-Vega, 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
  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 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.

First you have to enable link-protection on all RSVP interfaces.

- a. Enable link-protection on all routers' RSVP interfaces.

[edit]

```
lab@Sun# set protocols rsvp interface ge-0/0/4.114 link-protection
```

```
[edit]
```

```
lab@Sun# set protocols rsvp interface ge-0/0/4.118 link-protection
```

```
[edit]
```

```
lab@Sun# set protocols rsvp interface ae0.0 link-protection
```

```
[edit]
```

```
lab@Sun# commit
```

```
commit complete
```

b. Configure link protection for LSPs A, B, E, F, Q, R, S and T.

```
[edit]
```

```
lab@Sun# set protocols mpls label-switched-path Sun-to-Procyon  
link-protection
```

```
[edit]
```

```
lab@Sun# commit
```

```
commit complete
```

Verify the LSP state after enabling link-protection.

```
lab@Sun# run show mpls lsp name Sun-to-Procyon extensive
Ingress LSP: 2 sessions

172.30.5.8
  From: 172.30.5.1, State: Up, ActiveRoute: 0, LSPname: Sun-to-
  Procyon
    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: Sun 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

```

Verify the RSVP interface which is enabled for link-protection.

```

[edit]
lab@Sun# 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.

```
[edit]
lab@Canopus# 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      Canopus-
to-Vega
172.30.5.8  172.30.5.3  Up    0 *    path-1         Canopus-
to-Procyon-1
172.30.5.8  172.30.5.3  Up    0 *    primary-1      Canopus-
to-Procyon-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      - Vega-
to-Canopus
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      -
Procyon-to-Canopus-1
172.30.5.3  172.30.5.8  Up    0  1 SE      3      -
Procyon-to-Canopus-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 Vega-
to-Sun
172.30.5.1  172.30.5.6  Up    0  1 SE  302352  302592 Vega-
to-Sun
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 A-
Centauri-to-Sirius
172.30.5.2  172.30.5.5  Up    0  1 SE  302272      3 A-
Centauri-to-Sirius
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 Rigel-
to-Arcturus-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 Sirius-
to-A-Centauri
```

```

172.30.5.5 172.30.5.2 Up 0 1 SE 302288 302560 Sirius-
to-A-Centauri
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 Sun-to-
Vega
172.30.5.6 172.30.5.1 Up 0 1 SE 302400 3 Sun-to-
Vega
172.30.5.7 172.30.5.4 Up 0 1 SE 302448 302672
Arcturus-to-Rigel-1
Total 17 displayed, Up 17, Down 0

```

NOTE: The output above is taken after node-link-protection is enabled as well as link-protection.

Let's list the RSVP sessions.

```

[edit]
lab@Canopus# 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
Canopus-to-Vega
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
Canopus-to-Procyon-2
172.30.5.8  172.30.5.3  Up      0   1 SE    -        300352
Canopus-to-Procyon-1
Total 7 displayed, Up 7, Down 0

Egress RSVP: 8 sessions
To          From      State   Rt  Style  Labelin  Labelout  LSPname
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        -
Procyon-to-Canopus-1
172.30.5.3  172.30.5.8  Up      0   1 SE    3        -
Procyon-to-Canopus-1
172.30.5.3  172.30.5.8  Up      0   1 SE    3        -
Procyon-to-Canopus-2
172.30.5.3  172.30.5.6  Up      0   1 SE    3        - Vega-
to-Canopus
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
Procyon-to-Sun
172.30.5.1 172.30.5.6 Up        0  1 SE  300384  300016 Vega-
to-Sun
172.30.5.1 172.30.5.6 Up        0  1 SE  300416  300032 Vega-
to-Sun
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 A-
Centauri-to-Sirius
172.30.5.2 172.30.5.5 Up        0  1 SE  300656      3 A-
Centauri-to-Sirius
172.30.5.2 172.30.5.7 Up        0  1 SE  300544      3 Rigel-
to-Sirius
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 A-
Centauri-to-Arcturus
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 Rigel-
to-Arcturus-2
172.30.5.5 172.30.5.2 Up        0  1 SE  300096  299824 Sirius-
to-A-Centauri
172.30.5.5 172.30.5.2 Up        0  1 SE  300112  299840 Sirius-
to-A-Centauri
172.30.5.5 172.30.5.4 Up        0  1 SE  300320  299952
Arcturus-to-A-Centauri
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 Sun-to-
Vega
172.30.5.6 172.30.5.1 Up        0  1 SE  300304      3 Sun-to-
Vega
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 Sirius-
to-Rigel
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	
Arcturus-to-Rigel-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	Sun-to-Procyon
172.30.5.8	172.30.5.1	Up	0	1	SE	300704	300400	Sun-to-Procyon
Total 30 displayed, Up 30, Down 0								

c. Configure node link protection for LSPs I, J, M and N.

Node link protection automatically protects the link but also the downstream node.

Configure node link protection on R3 for the LSP to R8.

```
[edit]
lab@Canopus# set protocols mpls label-switched-path Canopus-to-Procyon-1 node-link-protection
```

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

Verify node-link protection configured on the LSP.

```
[edit]
lab@Canopus# run show mpls lsp name Canopus-to-Procyon-1 extensive
Ingress LSP: 3 sessions

172.30.5.8
  From: 172.30.5.3, State: Up, ActiveRoute: 0, LSPname: Canopus-
to-Procyon-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):
    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.

```

[edit]
lab@Sun# set policy-options policy-statement load-balance term 1
then load-balance per-packet

```

```
[edit]
lab@Sun# set routing-options forwarding-table export load-balance

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

The load balancing policy should 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 behaviour 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.

```
[edit]
lab@Sun# set protocols mpls ipv6-tunneling

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

This statement instructs JUNOS to copy the routes from `inet.3` to `inet6.3` as v4 mapped IPv6 routes.

```
[edit]
lab@Sun# 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
Sun-to-Procyon
    to 172.30.0.6 via ge-0/0/4.114, label-
switched-path Sun-to-Procyon
::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
Sun-to-Vega
    to 172.30.0.6 via ge-0/0/4.114, label-
switched-path Sun-to-Vega
    to 172.30.0.10 via ge-0/0/4.118, label-
switched-path Sun-to-Vega
    to 172.30.0.10 via ge-0/0/4.118, label-
switched-path Sun-to-Vega
    to 172.30.0.2 via ae0.0, label-switched-path
Sun-to-Vega
    [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
Sun-to-Procyon
    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
Sun-to-Procyon
    to 172.30.0.6 via ge-0/0/4.114, label-
switched-path Sun-to-Procyon

```

IPv6 tunneling should be configured on all PE routers.

- 2) Configure IPv6 labeled-unicast family on all routers.

IPv6 labeled unicast should be enabled on all PE routers. This enables the exchange of IPv6 labeled routes.

```

[edit]
lab@Sun# set protocols bgp group ibgp family inet6 labeled-unicast
explicit-null

```

NOTE: The **explicit-null** knob 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@Sun# set protocols bgp group ibgp family inet unicast

[edit]
lab@Sun# 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.

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 family
inet6 labeled-unicast explicit-null

[edit]
root@route-reflector# set protocols bgp group cluster-1 family
inet unicast

[edit]
root@route-reflector# set protocols bgp group cluster-2 family
inet6 labeled-unicast explicit-null

[edit]
root@route-reflector# set protocols bgp group cluster-2 family
inet unicast

[edit]
root@route-reflector# commit
[edit protocols bgp group cluster-1 cluster]
commit complete
```

RR should start to receive routes in table **inet6.0**

```
[edit]
root@route-reflector# run show route table inet6

inet6.0: 16 destinations, 16 routes (0 active, 0 holddown, 64
hidden)
```

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]
root@route-reflector# run show route table inet6 hidden

inet6.0: 32 destinations, 48 routes (0 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
        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
--- (more) ---

```

As you can see the routes are present, but are Unusable.

```

[edit]
root@route-reflector# set routing-options rib inet6.3 static route
0:0:0:0:0:ffff::/96 receive

[edit]
root@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]
root@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
                    *[BGP/170] 00:18:02, localpref 100, from
172.30.5.7
                    AS path: 23456 I
                    Indirect
--- (more) ---

```

```

[edit]
root@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
        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
  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@Canopus# 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
  *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 Canopus-to-Vega
            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
            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@A-Centauri# 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                 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@A-Centauri# set policy-options policy-statement ipv6-next-hop-
1 term 1 from family inet6
```

```
[edit]
lab@A-Centauri# set policy-options policy-statement ipv6-next-hop-
1 term 1 from protocol bgp
```

```
[edit]
```

```
lab@A-Centauri# set policy-options policy-statement ipv6-next-hop-1 term 1 then next-hop ::192.168.0.9
```

```
[edit]
```

```
lab@A-Centauri# set policy-options policy-statement ipv6-next-hop-2 term 1 from family inet6
```

```
[edit]
```

```
lab@A-Centauri# set policy-options policy-statement ipv6-next-hop-2 term 1 from protocol bgp
```

```
[edit]
```

```
lab@A-Centauri# set policy-options policy-statement ipv6-next-hop-2 term 1 then next-hop ::192.168.0.13
```

```
[edit]
```

```
lab@A-Centauri# set protocols bgp group C3 neighbor 192.168.0.10 export ipv6-next-hop-1
```

```
[edit]
```

```
lab@A-Centauri# set protocol bgp group C3 neighbor 192.168.0.14 export ipv6-next-hop-2
```

```
[edit]
```

```
lab@A-Centauri# commit
commit complete
```

Verify the IPv6 routes advertised from R5.

```
[edit]
lab@A-Centauri# 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
--- (more) ---
```

```
[edit]
lab@A-Centauri# 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
--- (more) ---
```

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.

This was the last step of Chapter 4. Do not forget to save the configuration on all routers for next chapter!

Chapter Five solutions: 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 18.

```
[edit]
lab@Sun# set interfaces ge-0/0/5.311 vlan-id 311

[edit]
lab@Sun# set interfaces ge-0/0/5.311 family inet address
192.168.0.41/30

[edit]
lab@Sun# set interfaces ge-0/0/5.311 description "CE2-1 connection
1"
```



```

[edit]
lab@Sun# set interfaces ge-0/0/5.312 vlan-id 312

[edit]
lab@Sun# set interfaces ge-0/0/5.312 family inet address
192.168.0.45/30

[edit]
lab@Sun# set interfaces ge-0/0/5.312 description "CE2-1 connection
2"

[edit]
lab@Sun# set interfaces ge-0/0/5.313 vlan-id 313

[edit]
lab@Sun# set interfaces ge-0/0/5.313 family inet address
192.168.0.49/30

[edit]
lab@Sun# set interfaces ge-0/0/5.313 description "CE2-1 connection
3"

[edit]
lab@Sun# 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.

```

[edit]
lab@Sun# set protocols bgp group ibgp family inet-vpn unicast

[edit]
lab@Sun# commit
commit complete

```

- 3) Configure the VPNv4 address family on the Route Reflector.

```

[edit]
root@route-reflector# set protocols bgp group cluster-1 family
inet-vpn unicast

[edit]
root@route-reflector# set protocols bgp group cluster-2 family
inet-vpn unicast

[edit]
root@route-reflector# commit
[edit protocols bgp group cluster-1 cluster]
commit complete

```

4) Configure customer C1 VPN

Customer C1 has sites attached to R3, R4, R7 and R8.

a. R3

```
[edit]
lab@Canopus# set interfaces lo0.1 family inet address
172.30.5.17/32

[edit]
lab@Canopus# set routing-options route-distinguisher-id 172.30.5.3

[edit]
lab@Canopus# set routing-instances C1 instance-type vrf

[edit]
lab@Canopus# set routing-instances C1 interface lo0.1

[edit]
lab@Canopus# set routing-instances C1 interface ge-0/0/5.318

[edit]
lab@Canopus# set routing-instances C1 vrf-target target:54591:100

[edit]
lab@Canopus# set routing-instances C1 protocols ospf area 0
interface all

[edit]
lab@Canopus# 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 automatic 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.

```
[edit]
lab@Canopus# 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 have to specify the **instance** with the command to see the output specifically for a particular routing-instance

b. R4.

```
[edit]
lab@Arcturus# set interfaces lo0.1 family inet address
172.30.5.21/32

[edit]
lab@Arcturus# set routing-options route-distinguisher-id
172.30.5.4

[edit]
lab@Arcturus# set routing-instances C1 instance-type vrf

[edit]
lab@Arcturus# set routing-instances C1 interface lo0.1

[edit]
lab@Arcturus# set routing-instances C1 interface ge-0/0/5.319

[edit]
lab@Arcturus# set routing-instances C1 vrf-target target:54591:100

[edit]
lab@Arcturus# set routing-instances C1 protocols ospf area 0
interface all

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

The routing instance on R4 is configured following the same principle as on R3. The import and export route target is 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** stanza.

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 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@Vega# set interfaces lo0.1 family inet address 172.30.5.29/32

[edit]
lab@Vega# set routing-options route-distinguisher-id 172.30.5.6

[edit]
lab@Vega# set routing-instances C1 interface lo0.1

[edit]
lab@Vega# set routing-instances C1 interface ge-0/0/5.322

[edit]
lab@Vega# set routing-instances C1 instance-type vrf

[edit]
lab@Vega# set routing-instances C1 vrf-target target:54591:100

[edit]
lab@Vega# set routing-instances C1 protocols ospf area 0 interface
all

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

```
[edit]
lab@Vega# 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@Procyon# set interfaces lo0.1 family inet address
172.30.5.37/32

[edit]
lab@Procyon# set routing-options route-distinguisher-id 172.30.5.8

[edit]
lab@Procyon# set routing-instances C1 instance-type vrf

[edit]
lab@Procyon# set routing-instances C1 interface lo0.1

[edit]
lab@Procyon# set routing-instances C1 interface ge-0/0/5.324

[edit]
lab@Procyon# set routing-instances C1 vrf-target target:54591:100
```

```
[edit]
lab@Procyon# set routing-instances C1 protocols ospf area 0
interface all
```

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

```
[edit]
lab@Procyon# 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.13vpn.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.

```
[edit]
root@route-reflector# set routing-options rib inet.3 static route
172.30.5.0/24 receive
```

```
[edit]
root@route-reflector# commit
[edit protocols bgp group cluster-1 cluster]

commit complete
```

Verify if the routes specified in the **bgp.l3vpn.0** table are resolved.

```
[edit]
root@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
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.

Each routing-instance has its own routing table where all routes for a particular VPN is stored.

```

[edit]
lab@Canopus# run show route table C1.inet.0
C1.inet.0: 18 destinations, 34 routes (18 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.17/32      *[Direct/0] 01:05:57
                   > via lo0.1
                   [BGP/170] 00:08:06, MED 3, localpref 100, from
172.30.5.41
                   AS path: I
                   > to 172.30.0.22 via ge-0/0/4.134, Push 303280
172.30.5.21/32     *[OSPF/10] 00:08:04, metric 3
                   > to 192.168.0.70 via ge-0/0/5.318
                   [BGP/170] 00:08:11, localpref 100, from
172.30.5.41
                   AS path: I
                   > to 172.30.0.22 via ge-0/0/4.134, Push 303168
172.31.48.0/30     *[OSPF/10] 00:21:14, metric 6
                   > to 192.168.0.70 via ge-0/0/5.318
                   [BGP/170] 00:08:06, MED 7, localpref 100, from
172.30.5.41
                   AS path: I
                   > to 172.30.0.22 via ge-0/0/4.134, Push 303280
172.31.48.4/30     *[OSPF/10] 00:21:14, metric 2
                   > to 192.168.0.70 via ge-0/0/5.318
                   [BGP/170] 00:08:06, MED 2, localpref 100, from
172.30.5.41
                   AS path: I
                   > to 172.30.0.22 via ge-0/0/4.134, Push 303280
172.31.48.8/30     *[OSPF/10] 00:21:14, metric 2
                   > to 192.168.0.70 via ge-0/0/5.318
                   [BGP/170] 00:08:06, MED 3, localpref 100, from
172.30.5.41
                   AS path: I
                   > to 172.30.0.22 via ge-0/0/4.134, Push 303280

```

Verify the OSPF database on R3.


```
[edit]
lab@Canopus# 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    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 a 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@Arcturus# 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@Vega# 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
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
Network 192.168.0.86      172.31.63.4      0x80000001    61    0x22
0x5f13  32
```

Verify the OSPF database on R8.

```
[edit]
lab@Procyon# 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
Network 192.168.0.86      172.31.63.4      0x80000001    276   0x22
0x5f13 32
Network 192.168.0.94      172.31.63.1      0x80000001    18    0x22
0x73f4 32
```

- a. Configure a Sham Link on R3.

```
[edit]
lab@Canopus# set routing-instances C1 protocols ospf sham-link
local 172.30.5.17

[edit]
lab@Canopus# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.21 metric 100

[edit]
lab@Canopus# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.29

[edit]
lab@Canopus# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.37

[edit]
lab@Canopus# 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.

- b. Configure the sham link on R4.

```
[edit]
lab@Arcturus# set routing-instances C1 protocols ospf sham-link
local 172.30.5.21

[edit]
lab@Arcturus# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.17 metric 100

[edit]
lab@Arcturus# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.29

[edit]
lab@Arcturus# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.37

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

c. Configure the sham link on R6.

```
[edit]
lab@Vega# set routing-instances C1 protocols ospf sham-link local
172.30.5.29

[edit]
lab@Vega# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.17

[edit]
lab@Vega# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.21

[edit]
lab@Vega# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.37

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

On R6 it is not needed to increase the metric.

d. Configure the sham link on R8.

```
[edit]
lab@Procyon# set routing-instances C1 protocols ospf sham-link
local 172.30.5.37

[edit]
lab@Procyon# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.17
```

```
[edit]
lab@Procyon# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.21
```

```
[edit]
lab@Procyon# set routing-instances C1 protocols ospf area 0 sham-
link-remote 172.30.5.29
```

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

Verify the operation of the sham link.

```
[edit]
lab@Canopus# 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@Canopus# 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.

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 multi-homed 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@Sun# set interfaces lo0.1 family inet address 172.30.5.9/32

[edit]
lab@Sun# set interfaces lo0.2 family inet address 172.30.5.10/32
```

Configure the automatic route distinguisher.

```
[edit]
lab@Sun# 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 also!

```
[edit]
lab@Sun# set routing-options autonomous-system 54951 loops 3

[edit]
root@route-reflector# 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@Sun# set routing-instances C2-hub instance-type vrf

[edit]
lab@Sun# set routing-instances C2-hub interface lo0.1

[edit]
lab@Sun# set routing-instances C2-hub interface ge-0/0/5.311

[edit]
lab@Sun# set routing-instances C2-hub vrf-import C2-hub-import

[edit]
lab@Sun# 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@Sun# set routing-instances C2-hub protocols bgp group ce peer-
as 64600

[edit]
lab@Sun# 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@Sun# set routing-instances C2-spoke instance-type vrf

[edit]
lab@Sun# set routing-instances C2-spoke interface lo0.2

[edit]
lab@Sun# set routing-instances C2-spoke interface ge-0/0/5.312

[edit]
lab@Sun# set routing-instances C2-spoke vrf-import C2-spoke-import

[edit]
lab@Sun# set routing-instances C2-spoke vrf-export C2-spoke-export
```

Again different policies are used for importing and exporting route-targets.

```
[edit]
lab@Sun# set routing-instances C2-spoke protocols bgp group ce
peer-as 64600

[edit]
lab@Sun# set routing-instances C2-spoke protocols bgp group ce as-
override

[edit]
lab@Sun# 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** knob. This ensures that the CE router on the other end will accept the routes and not drop them because of the AS path loop.

```
[edit]
lab@Sun# set routing-instances C2-spoke protocols bgp group ce
peer-as 64600

[edit]
lab@Sun# set routing-instances C2-spoke protocols bgp group ce as-
override

[edit]
lab@Sun# 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@Sun# set policy-options policy-statement C2-hub-export term 1
from protocol [ bgp direct ]

[edit]
```



```
lab@Sun# set policy-options policy-statement C2-hub-export term 1
then community set CE2-hub
```

```
[edit]
```

```
lab@Sun# 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 don't include the direct routes they will be rejected.

```
[edit]
```

```
lab@Sun# set policy-options policy-statement C2-hub-import term 1
then reject
```

The import policy rejects everything.

```
[edit]
```

```
lab@Sun# set policy-options policy-statement C2-spoke-import term
1 from protocol bgp
```

```
[edit]
```

```
lab@Sun# set policy-options policy-statement C2-spoke-import term
1 from community CE2-spoke
```

```
[edit]
```

```
lab@Sun# 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@Sun# set policy-options policy-statement C2-spoke-export then
reject
```

Nothing is exported from the spoke routing instance.

At last the communities must be defined.

```
[edit]
```

```
lab@Sun# set policy-options community CE2-hub members
target:54591:200
```

```
[edit]
```

```
lab@Sun# set policy-options community CE2-spoke members
target:54591:201
```

```
[edit]
```

```
lab@Sun# commit
commit complete
```

b. 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@Sirius# set interfaces lo0.1 family inet address
172.30.5.13/32

[edit]
lab@Sirius# set interfaces lo0.2 family inet address
172.30.5.14/32

[edit]
lab@Sirius# set routing-options route-distinguisher-id 172.30.5.2

[edit]
lab@Sirius# set routing-options autonomous-system 54591 loops 3
```

Configure the hub routing instance.

```
[edit]
lab@Sirius# set routing-instances C2-hub instance-type vrf

[edit]
lab@Sirius# set routing-instances C2-hub interface lo0.1

[edit]
lab@Sirius# set routing-instances C2-hub interface ge-0/0/5.314

[edit]
lab@Sirius# set routing-instances C2-hub vrf-import C2-hub-import

[edit]
lab@Sirius# set routing-instances C2-hub vrf-export C2-hub-export

[edit]
lab@Sirius# set routing-instances C2-hub vrf-table-label

[edit]
lab@Sirius# set routing-instances C2-hub protocols bgp group ce
peer-as 64600

[edit]
lab@Sirius# set routing-instances C2-hub protocols bgp group ce
neighbor 192.168.0.54
```

Configure the spoke routing instance.

```
[edit]
lab@Sirius# set routing-instances C2-spoke instance-type vrf

[edit]
lab@Sirius# set routing-instances C2-spoke interface lo0.2

[edit]
lab@Sirius# set routing-instances C2-spoke interface ge-0/0/5.315
```

```
[edit]
lab@Sirius# set routing-instances C2-spoke vrf-import C2-spoke-
import

[edit]
lab@Sirius# set routing-instances C2-spoke vrf-export C2-spoke-
export

[edit]
lab@Sirius# set routing-instances C2-spoke protocols bgp group ce
peer-as 64600

[edit]
lab@Sirius# set routing-instances C2-spoke protocols bgp group ce
neighbor 192.168.0.58

[edit]
lab@Sirius# set routing-instances C2-spoke protocols bgp group ce
as-override
```

Configure the import and export route target policies.

```
[edit]
lab@Sirius# set policy-options policy-statement C2-hub-export term
1 from protocol [ bgp direct ]

[edit]
lab@Sirius# set policy-options policy-statement C2-hub-export term
1 then community set CE2-hub

[edit]
lab@Sirius# set policy-options policy-statement C2-hub-import then
reject

[edit]
lab@Sirius# set policy-options policy-statement C2-spoke-export
term 1 then reject

[edit]
lab@Sirius# set policy-options policy-statement C2-spoke-import
term 1 from protocol bgp

[edit]
lab@Sirius# set policy-options policy-statement term 1 from
community CE2-spoke

[edit]
lab@Sirius# set policy-options policy-statement C2-spoke-import
term 1 then accept

[edit]
lab@Sirius# set policy-options community C2-hub members
target:54591:200
```

```
[edit]
lab@Sirius# set policy-options community C2-spoke members
target:54591:201
```

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

c. Configure C2 VPN on R4.

R4 is a spoke site, which is multi-homed on R5 also.

```
[edit]
lab@Arcturus# set interfaces lo0.2 family inet address
172.30.5.22/32
```

```
[edit]
lab@Arcturus# 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@Arcturus# set routing-instances C2-spoke instance-type vrf
```

```
[edit]
lab@Arcturus# set routing-instances C2-spoke interface lo0.2
```

```
[edit]
lab@Arcturus# set routing-instances C2-spoke interface ge-
0/0/5.320
```

```
[edit]
lab@Arcturus# set routing-instances C2-spoke vrf-import C2-spoke-
import
```

```
[edit]
lab@Arcturus# set routing-instances C2-spoke vrf-export C2-spoke-
export
```

```
[edit]
lab@Arcturus# set routing-instances C2-spoke protocols bgp group
ce peer-as 64600
```

```
[edit]
lab@Arcturus# set routing-instances C2-spoke protocols bgp group
ce as-override
```

```
[edit]
lab@Arcturus# 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@Arcturus# set policy-options policy-statement C2-spoke-import
term 1 from protocol bgp
```

```
[edit]
```

```
lab@Arcturus# set policy-options policy-statement C2-spoke-import
term 1 from community CE2-hub
```

```
[edit]
```

```
lab@Arcturus# 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@Arcturus# set policy-options policy-statement C2-spoke-export
term 1 from protocol [ bgp direct ]
```

```
[edit]
```

```
lab@Arcturus# set policy-options policy-statement C2-spoke-export
term 1 then community set CE2-spoke
```

```
[edit]
```

```
lab@Arcturus# set policy-options policy-statement C2-spoke-export
term 1 then accept
```

```
[edit]
```

```
lab@Arcturus# 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@Arcturus# set policy-options community CE2-hub members
target:54591:200
```

```
[edit]
```

```
lab@Arcturus# set policy-options community CE2-spoke members
target:54591:201
```

```
[edit]
```

```
lab@Arcturus# commit
commit complete
```

- d. 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@A-Centauri# set interfaces lo0.1 family inet address
172.30.5.25/32
```

```
[edit]
```

```

lab@A-Centauri# set routing-options route-distinguisher-id
172.30.5.5

[edit]
lab@A-Centauri# set routing-options autonomous-system 54591 loops
3

[edit]
lab@A-Centauri# set routing-instances C2-spoke instance-type vrf

[edit]
lab@A-Centauri# set routing-instances C2-spoke interface lo0.1

[edit]
lab@A-Centauri# set routing-instances C2-spoke interface ge-
0/0/5.321

[edit]
lab@A-Centauri# set routing-instances C2-spoke vrf-import C2-
spoke-import

[edit]
lab@A-Centauri# set routing-instances C2-spoke vrf-export C2-
spoke-export

[edit]
lab@A-Centauri# set routing-instances C2-spoke protocols bgp group
ce peer-as 64600

[edit]
lab@A-Centauri# set routing-instances C2-spoke protocols bgp group
ce as-override

[edit]
lab@A-Centauri# set routing-instances C2-spoke rotocols bgp group
ce neighbor 192.168.0.82

```

Configure the export and import policies.

```

[edit]
lab@A-Centauri# set policy-options policy-statement C2-spoke-
import term 1 from protocol bgp

[edit]
lab@A-Centauri# set policy-options policy-statement C2-spoke-
import term 1 from community CE2-hub

[edit]
lab@A-Centauri# set policy-options policy-statement C2-spoke-
import term 1 then accept

[edit]
lab@A-Centauri# set policy-options policy-statement C2-spoke-
export term 1 from protocol [ bgp direct ]

```

```

[edit]
lab@A-Centauri# set policy-options policy-statement C2-spoke-
export term 1 then community set CE2-spoke

[edit]
lab@A-Centauri# set policy-options policy-statement C2-spoke-
export term 1 then accept

[edit]
lab@A-Centauri# set policy-options community CE2-hub members
target:54591:200

[edit]
lab@A-Centauri# set policy-options community CE2-spoke members
target:54591:201

[edit]
lab@A-Centauri# commit
commit complete

```

- e. Configure C2 VPN on R7.

Customer C2 has one or more spoke sites connected to R7.

```

[edit]
lab@Rigel# set interfaces lo0.1 family inet address 172.30.5.33/32

[edit]
lab@Rigel# set routing-options autonomous-system 54591 loops 3

[edit]
lab@Rigel# set routing-instances C2-spoke instance-type vrf

[edit]
lab@Rigel# set routing-instances C2-spoke interface lo0.1

[edit]
lab@Rigel# set routing-instances C2-spoke interface ge-0/0/5.323

[edit]
lab@Rigel# set routing-instances C2-spoke vrf-import C2-spoke-
import

[edit]
lab@Rigel# set routing-instances C2-spoke vrf-export C2-spoke-
export

[edit]
lab@Rigel# set routing-instances C2-spoke protocols bgp group ce
peer-as 64600

[edit]

```

```
lab@Rigel# set routing-instances C2-spoke protocols bgp group ce
as-override
```

```
[edit]
```

```
lab@Rigel# set routing-instances C2-spoke protocols bgp group ce
neighbor 192.168.0.90
```

Configure the export and import policies.

```
[edit]
```

```
lab@Rigel# set policy-options policy-statement C2-spoke-import
term 1 from protocol bgp
```

```
[edit]
```

```
lab@Rigel# set policy-options policy-statement C2-spoke-import
term 1 from community CE2-hub
```

```
[edit]
```

```
lab@Rigel# set policy-options policy-statement C2-spoke-import
term 1 then accept
```

```
[edit]
```

```
lab@Rigel# set policy-options policy-statement C2-spoke-export
term 1 from protocol [ bgp direct ]
```

```
[edit]
```

```
lab@Rigel# set policy-options policy-statement C2-spoke-export
term 1 then community set CE2-spoke
```

```
[edit]
```

```
lab@Rigel# set policy-options policy-statement C2-spoke-export
term 1 then accept
```

```
[edit]
```

```
lab@Rigel# set policy-options community CE2-hub members
target:54591:200
```

```
[edit]
```

```
lab@Rigel# set policy-options community CE2-spoke members
target:54591:201
```

```
[edit]
```

```
lab@Rigel# commit
commit complete
```

Verify the hub and spoke L3VPN configuration.

```
[edit]
```

```
lab@Sun# 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.13vpn.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.13vpn.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 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@Sun# run show route advertising-protocol bgp 192.168.0.42

```

Nothing is advertised from R1 to the hub CE EBGp peer.

```

[edit]
lab@Sun# 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** knob. This is also the reason why you need to enable each router to accept an AS path with its' own AS.

The spoke routing instance advertises the routes learned from the other spokes to the CE router.

```

[edit]
lab@Sun# 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@Sun# 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.

```
[edit]
lab@Sun# 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
Sun-to-Procyon
to 172.30.0.2 via ae0.0, label-switched-path
Sun-to-Procyon

```

R2 is also connected to the hub site for redundancy.

```

[edit]
lab@Sirius# run show bgp summary
Groups: 4 Peers: 5 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp
State    Pending
inet.0
0          0          1128          402          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    182377    143083    0      385
32 Active
192.168.0.54    64600      836      850      0      0
6:25:46 Establ
C2-hub.inet.0: 20/20/20/0
192.168.0.58    64600      826     1206      0      0
6:25:30 Establ
C2-spoke.inet.0: 0/0/0/0
192.168.1.3     1620     1294     48109      0      0
6:25:49 Establ
inet.0: 402/564/402/0
192.168.1.4     1620     1294     48108      0      0
6:25:47 Establ
inet.0: 0/564/402/0

```

Both EBGP sessions are established and as on R1 only the hub session is receiving routes.

```

[edit]
lab@Sirius# 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@Sirius# run show route advertising-protocol bgp 192.168.0.54

```

No routes are advertised to that peer from hub routing instance.

```

[edit]
lab@Sirius# 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)

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@Sirius# 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.

```
[edit]
lab@Arcturus# 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
* 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@Arcturus# 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.

On R7 the two prefixes are received from the spoke. All other prefixes from the hub are send to the CE.

```
[edit]
lab@Rigel# 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
* 172.30.5.22/32        Self
* 172.30.5.25/32        Self
* 172.31.64.0/24        Self
* 172.31.65.0/24        Self
* 172.31.66.0/24        Self
* 172.31.67.0/24        Self
* 172.31.68.0/24        Self
* 172.31.69.0/24        Self
* 172.31.70.0/24        Self
* 172.31.71.0/24        Self
* 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        192.168.0.90
* 172.31.77.0/24        192.168.0.90
* 192.168.0.40/30       Self
* 192.168.0.76/30       Self
* 192.168.0.80/30       Self
```



```
[edit]
lab@Rigel# run show route receive-protocol bgp 192.168.0.90

inet.0: 841 destinations, 1242 routes (673 active, 0 holddown, 554
hidden)

inet.3: 18 destinations, 25 routes (8 active, 0 holddown, 15
hidden)

C2-spoke.inet.0: 23 destinations, 27 routes (23 active, 0
holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref    AS
path
* 172.31.76.0/24        192.168.0.90                64600
I
* 172.31.77.0/24        192.168.0.90                64600
I

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

mpls.0: 45 destinations, 45 routes (45 active, 0 holddown, 0
hidden)

bgp.l3vpn.0: 22 destinations, 22 routes (22 active, 0 holddown, 0
hidden)

inet6.0: 82 destinations, 84 routes (82 active, 0 holddown, 0
hidden)

inet6.3: 8 destinations, 10 routes (8 active, 0 holddown, 0
hidden)
```

Verify the VPN connectivity from R4 to R7.

```
[edit]
lab@Arcturus# 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
```

```
[edit]
lab@Arcturus# 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, 40 byte
packets
 1  172.30.5.9  2.927 ms  8.475 ms  2.839 ms
 2  192.168.0.42 5.090 ms  5.168 ms  4.940 ms
 3  192.168.0.45 7.782 ms  7.615 ms  8.111 ms
 4  * * *
```

```

5 * * *
6 * * 192.168.0.73 7.862 ms !N
7 192.168.0.73 8.224 ms !N 13.988 ms !N 5.843 ms !N

```

Ping and traceroute from R4 to R5.

```

[edit]
lab@Arcturus# 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@Arcturus# 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 it 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** feature.

By default JUNOS assigns a different VPN label for each CE interface. When the penultimate router removes the transport label, the packet received on the egress LSR has only 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.

```

[edit]
lab@Sun# set routing-instances C2-hub vrf-table-label

[edit]
lab@Sun# set routing-instances C2-spoke vrf-table-label

[edit]
lab@Sun# commit
commit complete

```

Verify the MPLS label advertisement after the change.

```
[edit]
lab@Sun# 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
```

```
[edit]
lab@Arcturus# run show route table mpls.0 label 16 extensive

mpls.0: 31 destinations, 31 routes (31 active, 0 holddown, 0
hidden)
16 (1 entry, 0 announced)
  *VPN Preference: 0
    Next hop:
    Next table: C2-hub.inet.0
    Label operation: Pop
    Address: 0x155fb70
    Next-hop reference count: 1
    State: <Active NotInstall Int Ext>
    Age: 1:47
    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@Sun# set policy-options community C2-S1-origin members
origin:54591:1

[edit]
lab@Sun# set policy-options policy-statement C2-hub-export term 1
then community add C2-S1-origin
```

Create the **origin** extended community type. It's structure is the same as the target community but instead of target, the origin keyword is used when creating the community.

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@Sun# set policy-options policy-statement C2-spoke-import term
reject-origin from protocol bgp

[edit]
lab@Sun# set policy-options policy-statement C2-spoke-import term
reject-origin from community C2-S1-origin

[edit]
lab@Sun# 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@Sun# insert policy-options policy-statement C2-spoke-import
term reject-origin before term 1

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

After the commit, the routes from R1 are sent with both a route target and origin community.

```
[edit]
lab@Sun# 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
    VPN Label: 16
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 100
    AS path: [54591] 64600 54591 I
    Communities: target:54591:200 origin:54591:1
```

On R2 the configuration should be the same as on R1, because the same customer CE is attached to them.

On R4 and R5 a different origin community should be used to represent the routes advertised from that side. This way they will be 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@Arcturus# set routing-options community C2-S2-origin members
origin:54591:2

[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-export
term 1 then community add C2-S2-origin

[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-import
term reject-origin from protocol bgp

[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-import
reject-origin from community C2-S2-origin

[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-import
term reject-origin then reject

[edit]
```

```
lab@Arcturus# insert policy-options policy-statement C2-spoke-  
import term reject-origin before term 1
```

```
[edit]  
lab@Arcturus# 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 configured VRFs. The Route reflector only reflects the routes with the required route target to the PEs.

Configure the route target family on all routers.

```
[edit]  
lab@Sun# set protocols bgp group ibgp family route-target  
  
[edit]  
lab@Sun# commit  
commit complete
```

Configure the route target family on the Route Reflector.

```
[edit]  
root@route-reflector# set protocols bgp group cluster-1 family  
route-target  
  
[edit]  
root@route-reflector# set protocols bgp group cluster-2 family  
route-target  
  
[edit]  
root@route-reflector# commit  
[edit protocols bgp group cluster-1 cluster]  
  
commit complete
```

Verify the content of **bgp.rtarget.0** table, which is used for storing **route target family** information.

```
[edit]
root@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
172.30.5.3      *[BGP/170] 00:01:51, localpref 100, from
                AS path: I
                Indirect
172.30.5.4      [BGP/170] 00:01:46, localpref 100, from
                AS path: I
                Indirect
172.30.5.6      [BGP/170] 00:01:39, localpref 100, from
                AS path: I
                Indirect
172.30.5.8      [BGP/170] 00:01:26, localpref 100, from
                AS path: I
                Indirect
54591:54591:200/96
172.30.5.4      *[BGP/170] 00:01:46, localpref 100, from
                AS path: I
                Indirect
172.30.5.5      [BGP/170] 00:01:43, localpref 100, from
                AS path: I
                Indirect
172.30.5.7      [BGP/170] 00:01:31, localpref 100, from
                AS path: I
                Indirect
```

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@Sun# 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] 00:02:45
                Local
```


But on R1 only the local route is present in the table. Besides that three more hidden routes are present.

```
[edit]
lab@Sun# 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
            Address: 0x113bc8c
            Next-hop reference count: 3
            State: <Hidden Int Ext>
            Local AS: 54591 Peer AS: 54591
            Age: 3:02
            Task: BGP_54591.172.30.5.41+179
            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: 0

--- (more) ---
```

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 the rib-groups to copy the RR Loopback IP address from `inet.0` to `inet.3`.

```
[edit]
lab@Sun# set policy-options policy-statement rr-loopback-to-inet3
term 1 from route-filter 172.30.5.41/32 exact
```

```
[edit]
lab@Sun# set policy-options policy-statement rr-loopback-to-inet3
term 1 to rib inet.3
```

```
[edit]
lab@Sun# set policy-options policy-statement rr-loopback-to-inet3
term 1 then accept
```

```
[edit]
lab@Sun# set policy-options policy-statement rr-loopback-to-inet3
term 2 to rib inet.3
```

```
[edit]
lab@Sun# 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]
lab@Sun# set routing-options rib-groups rr-inet0-inet3 import-rib
[ inet.0 inet.3 ]
```

```
[edit]
lab@Sun# set routing-options rib-groups rr-inet0-inet3 import-
policy rr-loopback-to-inet3
```

```
[edit]
lab@Sun# set protocols isis rib-group rr-inet0-inet3
```

```
[edit]
lab@Sun# 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@Sun# 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@Sun# 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] 00:20:42, localpref 100, from
172.30.5.41
        AS path: I
        > to 172.30.0.66 via ge-0/0/4.206
54591:54591:200/96
    *[BGP/170] 00:20:42, localpref 100, from
172.30.5.41
        AS path: I
        > to 172.30.0.66 via ge-0/0/4.206
54591:54591:201/96
    *[RTarget/5] 00:21:13
        Local
        [BGP/170] 00:20:59, localpref 100, from
172.30.5.41
        AS path: I
        > 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@Arcturus# set routing-options rib-groups C1-C2-vpn import-rib
[ C1.inet.0 C2-spoke.inet.0 ]
```

```
[edit]
lab@Arcturus# 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@Arcturus# 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** stanza copies routes in both directions. The order of the ribs in rib-groups definition does not matter.

```
[edit]
lab@Arcturus# set routing-instances C1 protocols ospf rib-group
C1-C2-vpn

[edit]
lab@Arcturus# set routing-instances C2-spoke routing-options
interface-routes rib-group inet C2-C1-vpn

[edit]
lab@Arcturus# set routing-instances C2-spoke protocols bgp family
inet unicast rib-group C2-C1-vpn

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

Verify the presence of the copied routes.

```
[edit]
lab@Arcturus# 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]
lab@Arcturus# 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** knob.

```
[edit]
lab@Arcturus# 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]
lab@Arcturus# 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 interconnect 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@Arcturus# set policy-options policy-statement C1-vpn-export
term 1 from protocol direct
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 1 from route-filter 192.168.0.72/30 exact
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 1 from route-filter 172.30.5.21/32 exact
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 1 then community set CE1
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 1 then accept
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 2 from protocol ospf
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 2 then community set CE1
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 2 then accept
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 3 then reject
```

NOTE: term 3 with reject is needed, because the routes copied from C2 customer are BGP routes and will automatically advertised.

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-import
term 1 from protocol bgp
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-import
term 1 from community CE1
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-import
term 1 then accept
```

```
[edit]
lab@Arcturus# set policy-options community CE1 members
target:54591:100
```

```
[edit]
lab@Arcturus# delete routing-instances C1 vrf-target
```

```
[edit]
lab@Arcturus# set routing-instances C1 vrf-export C1-vpn-export
```

```
[edit]
lab@Arcturus# set routing-instances C1 vrf-import C1-vpn-import
```

Modifying C2-spoke export policy.

```
edit]
lab@Arcturus# delete policy-options policy-statement C2-spoke-
export term 1 from protocol bgp
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-export
term 1 from route-filter 192.168.0.76/30 exact
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-export
term 1 from route-filter 172.30.5.22/32 exact
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-export
term 2 from protocol bgp
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-export
term 2 then community set CE2-spoke
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-export
term 2 then community add C2-S2-origin
```

```
[edit]
lab@Arcturus# set policy-options policy-statement C2-spoke-export
term 2 then accept
```

```
[edit]
lab@Arcturus# 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@Arcturus# 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@Canopus# set routing-options rib-groups C1-vpn-inet import-rib
[ C1.inet.0 inet.0 ]

[edit]
lab@Canopus# set routing-instances C1 routing-options interface-
routes rib-group inet C1-vpn-inet

[edit]
lab@Canopus# set routing-instances C1 protocols ospf rib-group C1-
vpn-inet

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

The rib-group copies routing instance C1 routes to inet.0 routing table.

```
[edit]
lab@Canopus# 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@Canopus# 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@Canopus# set policy-options policy-statement C1-default-to-
ospf term 1 from protocol static
```

```
[edit]
lab@Canopus# set policy-options policy-statement C1-default-to-
ospf term 1 route-filter 0/0 exact
```

```
[edit]
lab@Canopus# set policy-options policy-statement C1-default-to-
ospf term 1 then accept
```

```
[edit]
lab@Canopus# set routing-instances C1 protocols ospf export C1-
default-to-ospf
```

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

On R3 everything in the routing instance routing table is advertised to Multiprotocol BGP because the **vrf-target** knob is used.

Verify the default route advertisement to other PE routers.

```
[edit]
lab@Procyon# 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 Procyon-to-Canopus-1
                    to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
                    to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-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
--- (more) ---
```

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 Configure **vrf-table-label** will advertise a common VRF label for all routes per routing instance routing table. This way R3 will perform a route lookup matching the static default route pointing to **inet.0**.

```
[edit]
lab@Procyon# 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 Procyon-to-Canopus-1
                    to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-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 Procyon-to-Canopus-1
                    to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
                    to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
--- (more) ---
```

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@Canopus# set routing-options aggregate route 172.31.48.0/20

[edit]
lab@Canopus# set policy-options policy-statement C1-inet-routes
term 1 from protocol aggregate

[edit]
lab@Canopus# set routing-options C1-inet-routes term 1 from route-
filter 172.31.48.0/20 exact

[edit]
lab@Canopus# set routing-options C1-inet-routes term 1 then accept

[edit]
lab@Canopus# set protocols bgp group ibgp export C1-inet-routes

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

Verify the advertisement of the aggregate route.

```
[edit]
lab@Canopus# 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@Arcturus# delete routing-options rib-groups C1-C2-vpn

[edit]
lab@Arcturus# set routing-options rib-groups C1-C2-vpn import-rib
[ C1.inet.0 C2-spoke.inet.0 inet.0 ]

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

Modifying the policies and advertising the default and aggregate routes.

```
[edit]
lab@Arcturus# set routing-instances C1 routing-options static
route 0/0 next-table inet.0

[edit]
lab@Arcturus# set policy-options policy-statement C1-default-to-
ospf term 1 from protocol static

[edit]
lab@Arcturus# set policy-options policy-statement C1-default-to-
ospf term 1 from route-filter 0/0 exact

[edit]
lab@Arcturus# set policy-options policy-statement C1-default-to-
ospf term 1 then accept

[edit]
lab@Arcturus# set routing-instances C1 protocols ospf export C1-
default-to-ospf
```

Modifying the route target export policy.

```
[edit]
lab@Arcturus# delete policy-options policy-statement C1-vpn-export
term 3

[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 3 from protocol static

[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 3 from route-filter 0/0 exact

[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 3 then community set CE1

[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 3 then accept

[edit]
lab@Arcturus# set policy-options policy-statement C1-vpn-export
term 4 then reject

[edit]
lab@Arcturus# set routing-instances C1 vrf-table-label

[edit]
lab@Arcturus# set routing-options aggregate route 172.31.48.0/20

[edit]
lab@Arcturus# set policy-options policy-statement C1-inet-routes
term 1 from protocol aggregate

[edit]
lab@Arcturus# set policy-options policy-statement C1-inet-routes
term 1 from route-filter 172.31.48.0/20 exact

[edit]
lab@Arcturus# set policy-options policy-statement C1-inet-routes
term 1 then accept

[edit]
lab@Arcturus# set protocols bgp group ibgp export C1-inet-routes

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

13) Configure Customer C2 internet access.

Customer C2 should have Internet access on the dedicate physical interface on R1 and R2. You have to configure an EBGP session with the CE router and advertise the default route.

```
[edit]
lab@Sun# 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@Sun# set routing-options aggregate route 0/0

[edit]
lab@Sun# set policy-options policy-statement default-only term 1
from protocol aggregate

[edit]
lab@Sun# set policy-options policy-statement default-only term 1
from route-filter 0/0 exact

[edit]
lab@Sun# set policy-options policy-statement default-only term 1
then accept

[edit]
lab@Sun# set policy-options policy-statement default-only term 2
then reject

[edit]
lab@Sun# set protocols bgp group CE2 type external export default-
only
```

Verify the advertisement of the default route.

```
[edit]
lab@Sirius# 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
{73 235 448 647 783 859 985 1017 1039 1078 1158 1413 1435 1449
1620 1712 217
```

The same configuration steps have to be applied also to R2.

Task 2: Multicast in L3VPNs

The second task requires Configure multicast for L3VPN. Customer C1 uses the draft-Rosen implementation, whereas customer C2 has NG-MVPN.

- 1) Configure PIM on all routers.

The Draft-Rosen implementation does require MPLS and BGP running in the core, although BGP can be used for PE router discovery. PIM is used as the control-plane protocol used in the core network. L3VPN customers attached to the routing instances on the PE routers are running an other PIM protocol instance (the customer PIM domain). Multicast GRE tunnels are used to tunnel the multicast traffic in the core network.

The PE establishes PIM adjacencies with all other PEs in the multicast VPN (mVPN). A single Multicast Distribution Tree (MDT) is created for a permanent tunnel carrying PIM control messages and data stream by default. This Inclusive MDT is called the default MDT. Additionally a data MDT can be created to for the VPN when multicast traffic grows. All PE's in the mVPN are members of the default MDT. PE's which have interested multicast received attached can become member of the Data MDT. Thus when multicast traffic is switched from the default MDT towards the Data MDT, only the PE's with interested receivers will receive the multicast traffic.

Protocol Independent Multicast (PIM) protocol should be configured on all core interfaces. It ensures that the default MDT for each mVPN can be created.

```
[edit]
lab@Sun# set protocols pim interface ge-0/0/4.114

[edit]
lab@Sun# set protocols pim interface ge-0/0/4.118

[edit]
lab@Sun# set protocols pim interface ae0.0

[edit]
lab@Sun# set protocols pim interface lo0.0

[edit]
lab@Sun# 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@Sun# run show pim interfaces
Instance: PIM.master
```

Name	Stat	Mode	IP V	State	NbrCnt	JoinCnt(sg)
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]
lab@Sun# 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@Sun# 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 RPF 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.

```
[edit]
lab@Sun# set routing-options rib-groups rr-inet0-inet3 import-rib
[ inet.0 inet.3 inet.2 ]
```

```
[edit]
lab@Sun# 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@Sun# set routing-options rib-groups mcast-rib import-rib
inet.2
```

```
[edit]
lab@Sun# set protocols pim rib-group inet mcast-rib
```

```
[edit]
lab@Sun# 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@Sun# 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@Sun# run show route table inet.2

inet.2: 57 destinations, 57 routes (57 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

172.30.0.12/30      *[IS-IS/15] 04:11:18, metric 15
                   > to 172.30.0.2 via ae0.0
172.30.0.16/30      *[IS-IS/15] 04:11:18, metric 15
                   > to 172.30.0.2 via ae0.0
172.30.0.20/30      *[IS-IS/15] 00:01:47, metric 20
                   > to 172.30.0.6 via ge-0/0/4.114
172.30.0.24/30      *[IS-IS/15] 00:07:18, metric 25
                   > to 172.30.0.2 via ae0.0
172.30.0.28/30      *[IS-IS/15] 00:01:47, metric 20
                   > to 172.30.0.6 via ge-0/0/4.114
172.30.0.32/30      *[IS-IS/15] 00:01:47, metric 25
                   > to 172.30.0.6 via ge-0/0/4.114
                   to 172.30.0.10 via ge-0/0/4.118
172.30.0.36/30      *[IS-IS/15] 00:16:42, metric 20
                   > to 172.30.0.10 via ge-0/0/4.118
172.30.0.40/30      *[IS-IS/15] 04:11:18, metric 25
                   > to 172.30.0.2 via ae0.0
172.30.0.44/30      *[IS-IS/15] 00:16:42, metric 20
                   > to 172.30.0.10 via ge-0/0/4.118
172.30.0.48/30      *[IS-IS/15] 00:01:47, metric 20
```

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@Sun# set interfaces lo0.0 family inet address 172.30.5.254/32

[edit]
lab@Sun# set protocols pim rp bootstrap family inet priority 200

[edit]
lab@Sun# set protocols pim rp local family inet address
172.30.5.254

[edit]
lab@Sun# set protocols pim rp local family inet anycast-pim rp-set
address 172.30.5.2
```

```
[edit]
lab@Sun# set protocols pim rp local family inet anycast-pim local-
address 172.30.5.1
```

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

The ip address 172.30.5.254 is used as the anycast-RP. The Local IP address is R1s Loopback IP address. R2 Loopback IP address is configured as RP set address announced by the Bootstrap router.

Verify the RP table and bootstrap table.

```
[edit]
lab@Sun# 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
```

```
[edit]
lab@Sun# 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
0
```

Local router is elected as the Bootstrap router.

b. Configure RP on R2.

```
[edit]
lab@Sirius# set interfaces lo0.0 family inet address
172.30.5.254/32
```

```
[edit]
lab@Sirius# set protocols pim rp bootstrap family inet priority
100
```

```
[edit]
lab@Sirius# set protocols pim rp local family inet address
172.30.5.254
```

```
[edit]
```

```
lab@Sirius# set protocols pim rp local family inet anycast-pim rp-
set address 172.30.5.1
```

```
[edit]
```

```
lab@Sirius# set protocols pim rp local family inet anycast-pim
local-address 172.30.5.2
```

```
[edit]
```

```
lab@Sirius# 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@Sirius# 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	0	None	0	
224.0.0.0/4					

```
[edit]
```

```
lab@Sirius# run show pim bootstrap
```

```
Instance: PIM.master
```

BSR	Pri	Local address	Pri	State
Timeout				
172.30.5.1	200	172.30.5.2	100	Candidate
126				
None	0	fd17:f0f4:f691:5::2	0	InEligible
0				

The elected Bootstrap router is R1, whereas R2 is a candidate to become the Bootstrap router.

4) Configure multicast for customer C1.

Router R3 and R4 are dedicated RPs for the customer PIM instance. Auto-RP is used for discovery of the RPs. The Mapping agents learn the RP from the announcement messages received from the candidate RP routers. These messages are sent to multicast address 224.0.1.39. Then the mapping agent chooses which router will be the RP and sends discovery messages to all other routers which listen on the 224.0.1.40 address.

a. Configure multicast for customer C1 on R3.

```
[edit]
```

```
lab@Canopus# set routing-instances C1 protocols pim interface all
mode sparse-dense
```

Sparse-dense mode is enabled on the PIM enabled interfaces in the routing instance. Dense mode is needed for Auto-RP operation. Sparse mode is usually the preferred mode for multicast routing as it has a better bandwidth utilization. Sparse mode needs the RP to ensure the multicast sources and receivers are able to get connected.

```
[edit]
lab@Canopus# set routing-instances C1 protocols pim rp auto-rp
discovery
```

Auto-RP is enabled.

```
[edit]
lab@Canopus# set routing-instances C1 protocols pim dense-groups
224.0.1.39/32
```

```
[edit]
lab@Canopus# set routing-instances C1 protocols pim dense-groups
224.0.1.40/32
```

Both groups used by Auto-RP are configured as dense groups. This is required for proper Auto-RP operation.

```
[edit]
lab@Canopus# set routing-instances C1 protocols pim vpn-group-
address 239.1.1.1
```

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

The **vpn-group-address** is the group that is used by this VRF for default MDT. It should not overlap with other VPNs.

- b. Configure customer C1 multicast on R4.

The configuration on R4 is the same as on R3

```
[edit]
lab@Arcturus# set routing-instances C1 protocols pim rp auto-rp
discovery
```

```
[edit]
lab@Arcturus# set routing-instances C1 protocols pim interface all
mode sparse-dense
```

```
[edit]
lab@Arcturus# set routing-instances C1 protocols pim dense-groups
224.0.1.39/32
```

```
[edit]
lab@Arcturus# set routing-instances C1 protocols pim dense-groups
224.0.1.40/32
```

```
[edit]
```

```
lab@Arcturus# set routing-instances C1 protocols pim vpn-group-  
address 239.1.1.1
```

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

c. Configure customer C1 multicast on R6.

```
[edit]  
lab@Vega# set routing-instances C1 protocols pim interface all  
mode sparse-dense  
  
[edit]  
lab@Vega# set routing-instances C1 protocols pim protocols pim rp  
auto-rp discovery
```

```
[edit]  
lab@Vega# set routing-instances C1 protocols pim dense-groups  
224.0.1.39/32
```

```
[edit]  
lab@Vega# set routing-instances C1 protocols pim dense-groups  
224.0.1.40/32
```

```
[edit]  
lab@Vega# set routing-instances C1 protocols pim vpn-group-address  
239.1.1.1
```

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

d. Configure customer C1 multicast on R8.

```
[edit]  
lab@Procyon# set routing-instances C1 protocols pim interface all  
mode sparse-dense
```

```
[edit]  
lab@Procyon# set routing-instances C1 protocols pim rp auto-rp  
discovery
```

```
[edit]  
lab@Procyon# set routing-instances C1 protocols pim dense-groups  
239.0.1.39/32
```

```
[edit]  
lab@Procyon# set routing-instances C1 protocols pim dense-groups  
239.0.1.40/32
```

```
[edit]  
lab@Procyon# set routing-instances C1 protocols pim vpn-group-  
address 239.1.1.1
```



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

Verify the operation of the draft-Rosen multicast VPN on R3 (where the source is attached).

```
[edit]
lab@Canopus# run show pim interfaces instance C1
Instance: PIM.C1
```

Name	Stat	Mode	IP V	State	NbrCnt	JoinCnt (sg)
ge-0/0/5.318	Up	SparseDense	4 2	NotDR	1	3
0 192.168.0.70						
lo0.1	Up	SparseDense	4 2	DR	0	0
0 172.30.5.17						
lsi.0	Up	SparseDense	4 2	P2P	0	0
0						
mt-0/0/0.1081344	Up	SparseDense	4 2	P2P	0	0
0						
mt-0/0/0.32768	Up	SparseDense	4 2	P2P	2	0
0						
ppe0.32770	Up	Sparse	4 2	P2P	0	0
0						
lsi.0	Up	SparseDense	6 2	P2P	0	0
0						

mt- interfaces are created when PIM is configured under the routing instance. One interface is used for encapsulation of the multicast traffic over the multicast GRE tunnel, the other is used for de-encapsulation of traffic received from the multicast GRE tunnel. These interfaces are used to communicate between the PIM instance within the VRF and the master PIM instance

ppe0 interface is PIM encapsulation interface. Used for the encapsulation of PIM register messages towards the RP.

```
lab@Canopus# run show pim rps instance C1
Instance: PIM.C1
Address family INET
RP address      Type      Holdtime Timeout Groups Group prefixes
172.31.63.3     auto-rp   150      146      0 239.0.0.0/24
Address family INET6
```

The RP is located on the CEs attached to R3 and R4.

```
[edit]
lab@Canopus# run show pim join extensive instance C1
Instance: PIM.C1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.0.1.39
  Source: 172.31.63.2
  Flags: dense
  Upstream interface: ge-0/0/5.318
  Upstream neighbor: 192.168.0.70
  Uptime: 00:04:40
  Downstream interfaces:
    Local
    lo0.1
    lsi.0
    mt-0/0/0.32768 (assert winner)
    mt-0/0/0.1081344

Group: 224.0.1.39
  Source: 172.31.63.3
  Flags: dense
  Upstream interface: ge-0/0/5.318
  Upstream neighbor: 192.168.0.70
  Uptime: 00:03:59
  Downstream interfaces:
    Local
    lo0.1
    lsi.0
    mt-0/0/0.32768 (assert lost)
    mt-0/0/0.1081344

Group: 224.0.1.40
  Source: 172.31.63.3
  Flags: dense
  Upstream interface: ge-0/0/5.318
  Upstream neighbor: 192.168.0.70 (assert winner)
  Uptime: 00:03:52
  Downstream interfaces:
    Local
    lo0.1
    lsi.0
    mt-0/0/0.32768 (assert lost)
    mt-0/0/0.1081344

Instance: PIM.C1 Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

Only PIM joins are received for the dense groups. Dense mode floods traffic on all interfaces. The interfaces are all listed as downstream interfaces.

```
[edit]
lab@Canopus# run show pim mdt instance C1
Instance: PIM.C1
```

```

Tunnel direction: Outgoing
Tunnel mode: PIM-SM
Default group address: 239.1.1.1
Default source address: 0.0.0.0
Default tunnel interface: mt-0/0/0.32768
Default tunnel source: 0.0.0.0

Instance: PIM.C1
Tunnel direction: Incoming
Tunnel mode: PIM-SM
Default group address: 239.1.1.1
Default source address: 0.0.0.0
Default tunnel interface: mt-0/0/0.1081344
Default tunnel source: 0.0.0.0

```

The default MDT uses the configured 239.1.1.1 multicast group.

```

[edit]
lab@Canopus# run show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 172.30.5.254
Flags: sparse,rptree,wildcard
Upstream interface: ge-0/0/4.123
Upstream neighbor: 172.30.0.13
Upstream state: Join to RP
Uptime: 00:04:00
Downstream neighbors:
  Interface: ge-0/0/4.136
    172.30.0.26 State: Join Flags: SRW Timeout: 159
    Uptime: 00:03:51 Time since last Join: 00:00:51
Downstream neighbors:
  Interface: mt-0/0/0.1081344
    172.30.5.17 State: Join Flags: SRW Timeout: Infinity
    Uptime: 00:04:00 Time since last Join: 00:04:00

Group: 239.1.1.1
Source: 172.30.5.3
Flags: sparse,spt
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local Source, Prune to RP
Keepalive timeout: 355
Uptime: 00:35:32
Downstream neighbors:
  Interface: ge-0/0/4.123
    172.30.0.13 State: Join Flags: S Timeout: 150
    Uptime: 00:03:49 Time since last Join: 00:00:59
Downstream neighbors:

```

```

    Interface: ge-0/0/4.134
      172.30.0.22 State: Join Flags: S Timeout: 154
      Uptime: 00:03:55 Time since last Join: 00:00:55
    Downstream neighbors:
      Interface: ge-0/0/4.136
        172.30.0.26 State: Join Flags: S Timeout: 159
        Uptime: 00:05:05 Time since last Join: 00:00:51
    Downstream neighbors:
      Interface: mt-0/0/0.1081344
        172.30.5.17 State: Join Flags: S Timeout: Infinity
        Uptime: 00:04:00 Time since last Join: 00:04:00

Group: 239.1.1.1
  Source: 172.30.5.4
  Flags: sparse,spt
  Upstream interface: ge-0/0/4.134
  Upstream neighbor: 172.30.0.22
  Upstream state: Join to Source, Prune to RP
  Keepalive timeout: 355
  Uptime: 00:03:56
  Downstream neighbors:
    Interface: ge-0/0/4.136 (pruned)
      172.30.0.26 State: Prune Flags: SR Timeout: 159
      Uptime: 00:03:49 Time since last Prune: 00:00:51
  Downstream neighbors:
    Interface: mt-0/0/0.1081344
      172.30.5.17 State: Join Flags: S Timeout: Infinity
      Uptime: 00:03:56 Time since last Join: 00:03:56

Group: 239.1.1.1
  Source: 172.30.5.6
  Flags: sparse,spt
  Upstream interface: ge-0/0/4.136
  Upstream neighbor: 172.30.0.26
  Upstream state: Join to Source, Prune to RP
  Keepalive timeout: 355
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: ge-0/0/4.136 (pruned)
      172.30.0.26 State: Prune Flags: SR Timeout: 159
      Uptime: 00:03:49 Time since last Prune: 00:00:51
  Downstream neighbors:
    Interface: mt-0/0/0.1081344
      172.30.5.17 State: Join Flags: S Timeout: Infinity
      Uptime: 00:03:49 Time since last Join: 00:03:49

Group: 239.1.1.1
  Source: 172.30.5.8
  Flags: sparse,spt
  Upstream interface: ge-0/0/4.123
  Upstream neighbor: 172.30.0.13
  Upstream state: None, Join to Source
  Keepalive timeout: 355
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: ge-0/0/4.136 (pruned)

```

```

172.30.0.26 State: Prune Flags: SR Timeout: 159
Uptime: 00:03:49 Time since last Prune: 00:00:51
Downstream neighbors:
Interface: mt-0/0/0.1081344
172.30.5.17 State: Join Flags: S Timeout: Infinity
Uptime: 00:03:49 Time since last Join: 00:03:49

```

In the Provider PIM instance all PEs are a source for the VPN group address. The first (star comma g) state (*,239.1.1.1) is send to the RP.

Verify the operation of draft-Rosen on R6, which is one of the PEs with attached receivers.

```

[edit]
lab@Vega# run show pim neighbors instance C1
Instance: PIM.C1
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
mt-0/0/0.32768      4 2           HPLGT       00:01:44
172.30.5.17
mt-0/0/0.32768      4 2           HPLGT       00:00:47
172.30.5.21
mt-0/0/0.32768      4 2           HPLGT       00:01:44
172.30.5.37

```

The customer PIM neighbors are visible over the multicast GRE interfaces.

```

lab@Vega# run show pim rps instance C1
Instance: PIM.C1
Address family INET
RP address      Type      Holdtime Timeout Groups Group prefixes
172.31.63.3     auto-rp   150      146      0 239.0.0.0/24
Address family INET6

```

The source of the dense groups used for Autop-RP are also reachable via the default MDT/multicast GRE interface.

```
[edit]
lab@Vega# run show pim join extensive instance C1
Instance: PIM.C1 Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.0.1.39
  Source: 172.31.63.2
  Flags: dense
  Upstream interface: mt-0/0/0.32769
  Upstream neighbor: 172.30.5.17
  Downstream interfaces:
    Local
    lo0.1
    ge-0/0/5.322
    mt-0/0/0.1081345

Group: 224.0.1.39
  Source: 172.31.63.3
  Flags: dense
  Upstream interface: mt-0/0/0.32769
  Upstream neighbor: 172.30.5.21
  Downstream interfaces:
    Local
    lo0.1
    ge-0/0/5.322
    mt-0/0/0.1081345

Group: 224.0.1.40
  Source: 172.31.63.3
  Flags: dense
  Upstream interface: mt-0/0/0.32769
  Upstream neighbor: 172.30.5.21
  Downstream interfaces:
    Local
    lo0.1
    ge-0/0/5.322
    mt-0/0/0.1081345

Instance: PIM.C1 Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
[edit]
lab@Vega# run show pim mdt instance C1
Instance: PIM.C1
Tunnel direction: Outgoing
Tunnel mode: PIM-SM
Default group address: 239.1.1.1
Default source address: 0.0.0.0
Default tunnel interface: mt-0/0/0.32769
Default tunnel source: 0.0.0.0

Instance: PIM.C1
Tunnel direction: Incoming
Tunnel mode: PIM-SM
Default group address: 239.1.1.1
Default source address: 0.0.0.0
Default tunnel interface: mt-0/0/0.1081345
Default tunnel source: 0.0.0.0
```

List the PIM join states in the Provider PIM domain.

```
[edit]
lab@Vega# run show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 172.30.5.254
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-0/0/4.167
  Upstream neighbor: 172.30.0.42
  Upstream state: Join to RP
  Downstream neighbors:
    Interface: mt-0/0/0.32769
    172.30.5.29 State: Join Flags: SRW Timeout: Infinity

Group: 239.1.1.1
  Source: 172.30.5.3
  Flags: sparse,spt
  Upstream interface: ge-0/0/4.136
  Upstream neighbor: 172.30.0.25
  Upstream state: Join to Source, Prune to RP
  Keepalive timeout: 308
  Downstream neighbors:
    Interface: mt-0/0/0.32769
    172.30.5.29 State: Join Flags: S Timeout: Infinity

Group: 239.1.1.1
  Source: 172.30.5.4
  Flags: sparse,spt
  Upstream interface: ae0.0
  Upstream neighbor: 172.30.0.33
  Upstream state: Join to Source, Prune to RP
  Keepalive timeout: 308
  Downstream neighbors:
```

```

        Interface: mt-0/0/0.32769
          172.30.5.29 State: Join Flags: S    Timeout: Infinity
Group: 239.1.1.1
  Source: 172.30.5.6
  Flags: sparse,spt
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local Source, Prune to RP
  Keepalive timeout: 308
  Downstream neighbors:
    Interface: ae0.0
      172.30.0.33 State: Join Flags: S Timeout: 201
    Interface: ge-0/0/4.136
      172.30.0.25 State: Join Flags: S Timeout: 204
    Interface: mt-0/0/0.32769
      172.30.5.29 State: Join Flags: S    Timeout: Infinity

Group: 239.1.1.1
  Source: 172.30.5.8
  Flags: sparse,spt
  Upstream interface: ae0.0
  Upstream neighbor: 172.30.0.33
  Upstream state: Join to Source, Prune to RP
  Keepalive timeout: 308
  Downstream neighbors:
    Interface: mt-0/0/0.32769
      172.30.5.29 State: Join Flags: S    Timeout: Infinity

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

Like R3, each PE which participates in the mVPN is a source for the VPN group.

5) Configure a data MDT on the PEs with attached sources.

An additional data MDTs should be configured for groups 239.0.0.1 and 239.0.0.2 when the traffic reaches more than 30Mbps.

```

[edit]
lab@Canopus# set routing-instances C1 protocols pim mdt threshold
group 239.0.0.1/32 source 0/0 rate 30000

[edit]
lab@Canopus# set routing-instances C1 protocols pim mdt threshold
group 239.0.0.2/32 source 0/0 rate 30000

[edit]
lab@Canopus# set routing-instances C1 protocols pim mdt tunnel-
limit 5

[edit]
lab@Canopus# set routing-instances C1 protocols pim mdt group-
range 239.0.0.0/24

```



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

Not more than 5 data MDTs are allowed. The **group-range** statement defines a range in the Provider PIM domain used for signaling the data MDTs. The address range should not overlap with any other default or data MDTs from other mVPNs.

The same configuration should be applied to R4.

6) 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 P-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.

First configure PIM in the customer spoke instance.

```
[edit]
lab@Sun# set routing-instances C2-spoke protocols pim interface
all
```

```
[edit]
```

```
lab@Sun# set routing-instances C2-spoke protocols pim rp local
address 172.30.5.253
```

```
[edit]
```

```
lab@Sun# set routing-instances C2-spoke protocols pim rp local
group-ranges 239.0.0.0/24
```

```
[edit]
```

```
lab@Sun# 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@Sun# set interfaces lo0.2 family inet address 172.30.5.253/32
```

```
[edit]
```

```
lab@Sun# set interfaces lo0.2 family inet address 172.30.5.10/32
primary
```

The shared IP address is configure on the Loopback interface as secondary.

```
[edit]
```

```
lab@Sun# set policy-options policy-statement C2-direct-routes term
1 from protocol direct
```

```
[edit]
```

```
lab@Sun# set policy-options policy-statement C2-direct-routes term
1 from route-filter 172.30.5.253/32 exact
```

```
[edit]
```

```
lab@Sun# set policy-options policy-statement C2-direct-routes term
1 then metric 10
```

```
[edit]
```

```
lab@Sun# set policy-options policy-statement C2-direct-routes term
1 then accept
```

```
[edit]
```

```
lab@Sun# set policy-options policy-statement C2-direct-routes term
2 from protocol direct
```

```
[edit]
```

```
lab@Sun# set policy-options policy-statement C2-direct-routes term
2 then accept
```

```
[edit]
```

```
lab@Sun# set routing-instances C2-spoke protocols bgp group ce
export C2-direct-routes
```

```
[edit]
```

```
lab@Sun# commit
commit complete
```

```
[edit]
lab@Sun# run show route advertising-protocol bgp 192.168.0.46

C2-spoke.inet.0: 19 destinations, 23 routes (19 active, 0
holdddown, 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
```

Configure **inet-mvpn signaling** on the IBGP session with the Route Reflector.

```
[edit]
lab@Sun# set protocols bgp group ibgp family inet-mvpn signaling

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

Configure **inet-mvpn signaling** on the Route Reflectors IBGP groups.

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 family
inet-mvpn signaling

[edit]
root@route-reflector# set protocols bgp group cluster-2 family
inet-mvpn signaling

[edit]
root@route-reflector# commit
[edit protocols bgp group cluster-1 cluster]

commit complete
```

25)

26) Verifying the `rt-import` community that should be advertised by the C2-hub routing-instance.

27)

```

28) [edit]
29)      lab@Sirius# run show policy __vrf-mvpn-import-cmcast-C2-
      spoke-internal__
30)      Policy __vrf-mvpn-import-cmcast-C2-spoke-internal__:
31)          Term unnamed:
32)              from community __vrf-mvpn-community-rt_import-target-
      C2-spoke-internal__ [target:172.30.5.2:7]
33)              then accept
34)          Term unnamed:
35)              then reject

```

36)

```

37)      [edit]
38)      lab@Sirius# set policy-options community src-as members src-
      as:54591:0
39)
40)      [edit]
41)      lab@Sirius# set policy-options community rt-import-spoke
      members rt-import:172.30.5.2:7
42)
43)      [edit]
44)      lab@Sirius# set policy-options policy-statement C2-hub-export
      term 1 then community add rt-import-spoke
45)
46)      [edit]
47)      lab@Sirius# 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 send without replication. Please note that unicast traffic still uses a regular point-to-point LSP.

Configure P2MP RSVP LSP.

[edit]

```
lab@Sun# set protocols mpls label-switched-path mcast-p2mp-
template template
```

```
[edit]
```

```
lab@Sun# set protocols mpls label-switched-path mcast-p2mp-
template bandwidth 30
```

```
[edit]
```

```
lab@Sun# set protocols mpls label-switched-path mcast-p2mp-
template hop-limit 5
```

```
[edit]
```

```
lab@Sun# set protocols mpls label-switched-path mcast-p2mp-
template priority 5 5
```

```
[edit]
```

```
lab@Sun# set protocols mpls label-switched-path mcast-p2mp-
template link-protection
```

```
[edit]
```

```
lab@Sun# 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@Sun# set routing-instances C2-spoke protocols mvpn mvpn-mode
spt-only
```

```
[edit]
```

```
lab@Sun# set routing-instances C2-spoke protocols mvpn sender-site
```

```
[edit]
```

```
lab@Sun# set routing-instances C2-spoke protocols mvpn route-
target import-target target target:54591:202
```

```
[edit]
```

```
lab@Sun# 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. 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@Sun# set routing-instances C2-spoke provider-tunnel rsvp-te
label-switched-path-template mcast-p2mp-template
```

```
[edit]
```

```
lab@Sun# commit
```

```
commit complete
```

The dynamic template (P2MP RSVP LSP) is configured as Inclusive P-tunnel.

- b. Configure NG MVPN on R2.

R2 configuration is the same as on R1 (redundancy).

```
[edit]
lab@Sirius# set interfaces lo0.2 family inet address
172.30.5.253/32

[edit]
lab@Sirius# set interfaces lo0.2 family inet address
172.30.5.14/32 primary

[edit]
lab@Sirius# set policy-options policy-statement C2-direct-routes
term 1 from protocol direct

[edit]
lab@Sirius# set policy-options policy-statement C2-direct-routes
term 1 from route-filter 172.30.5.253/32 exact

[edit]
lab@Sirius# set policy-options policy-statement C2-direct-routes
term 1 then metric 100

[edit]
lab@Sirius# set policy-options policy-statement C2-direct-routes
term 1 then accept

[edit]
lab@Sirius# set policy-options policy-statement C2-direct-routes
term 2 from protocol direct

[edit]
lab@Sirius# set policy-options policy-statement C2-direct-routes
term 2 then accept

[edit]
lab@Sirius# set routing-instances C2-spoke protocols bgp group ce
export C2-direct-routes

[edit]
lab@Sirius# 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@Sirius# run show route advertising-protocol bgp 192.168.0.58
table C2-spoke.inet.0
```

```

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@Sirius# set routing-instances C2-spoke protocols pim interface
all

```

```

[edit]
lab@Sirius# set routing-instances C2-spoke protocols pim rp local
address 172.30.5.253

```

```

[edit]
lab@Sirius# set routing-instances C2-spoke protocols pim rp local
group-ranges 239.0.0.0/24

```

```

[edit]
lab@Sirius# 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@Sirius# set protocols mpls label-switched-path mcast-p2mp-
template template

```

```

[edit]

```

```
lab@Sirius# set protocols mpls label-switched-path mcast-p2mp-  
template bandwidth 30
```

```
[edit]
```

```
lab@Sirius# set protocols mpls label-switched-path mcast-p2mp-  
template hop-limit 5
```

```
[edit]
```

```
lab@Sirius# set protocols mpls label-switched-path mcast-p2mp-  
template priority 5 5
```

```
[edit]
```

```
lab@Sirius# set protocols mpls label-switched-path mcast-p2mp-  
template link-protection
```

```
[edit]
```

```
lab@Sirius# set protocols mpls label-switched-path mcast-p2mp-  
template p2mp
```

Configure **protocols mvpn** for routing instance C2-spoke.

```
[edit]
```

```
lab@Sirius# set routing-instances C2-spoke protocols mvpn sender-  
site
```

```
[edit]
```

```
lab@Sirius# set routing-instances C2-spoke protocols mvpn mvpn-  
mode spt-only
```

```
[edit]
```

```
lab@Sirius# set routing-instances C2-spoke protocols mvpn route-  
target import-target target target:54591:202
```

```
[edit]
```

```
lab@Sirius# set routing-instances C2-spoke protocols mvpn route-  
target export-target target target:54591:202
```

```
[edit]
```

```
lab@Sirius# set routing-instances C2-spoke provider-tunnel rsvp-te  
label-switched-path-template mcast-p2mp-template
```

```
[edit]
```

```
lab@Sirius# commit  
commit complete
```

c. Configure NG MVPN support on R4.

```
[edit]
```

```
lab@Arcturus# set protocols bgp group ibgp family inet-mvpn  
signaling
```

```
[edit]
```

```
lab@Arcturus# set routing-instances C2-spoke protocols pim  
interface all
```



```

[edit]
lab@Arcturus# set routing-instances C2-spoke protocols mvpn
receiver-site

[edit]
lab@Arcturus# set routing-instances C2-spoke protocols mvpn mvpn-
mode spt-only

[edit]
lab@Arcturus# set routing-instances C2-spoke protocols mvpn route-
target import-target target target:54591:202

[edit]
lab@Arcturus# set routing-instances C2-spoke protocols mvpn route-
target export-target target target:54591:202

[edit]
lab@Arcturus# 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@A-Centauri# set protocols bgp group ibgp family inet-mvpn
signaling

[edit]
lab@A-Centauri# set routing-instances C2-spoke protocols pim
interface all

[edit]
lab@A-Centauri# set routing-instances C2-spoke protocols mvpn
mvpn-mode spt-only

[edit]
lab@A-Centauri# set routing-instances C2-spoke protocols mvpn
receiver-site

[edit]
lab@A-Centauri# set routing-instances C2-spoke protocols mvpn
route-target import-target target target:54591:202

[edit]
lab@A-Centauri# set routing-instances C2-spoke protocols mvpn
route-target export-target target target:54591:202

[edit]
lab@A-Centauri# 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.

```
[edit]
lab@Rigel# set protocols bgp group ibgp family inet-mvpn signaling

[edit]
lab@Rigel# set routing-instances C2-spoke protocols pim interface
all

[edit]
lab@Rigel# set routing-instances C2-spoke protocols mvpn mvpn-mode
spt-only

[edit]
lab@Rigel# set routing-instances C2-spoke protocols mvpn receiver-
site

[edit]
lab@Rigel# set routing-instances C2-spoke protocols mvpn route-
target import-target target target:54591:202

[edit]
lab@Rigel# set routing-instances C2-spoke protocols mvpn route-
target export-target target target:54591:202

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

Verify the operation of NG MVPN for customer C2 on R1.

```
[edit]
lab@Sun# 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@Sun# run show pim interfaces instance C2-spoke
Instance: PIM.C2-spoke
```

Name	Stat	Mode	IP V	State	NbrCnt	JoinCnt (sg)
JoinCnt(*g) DR address						
ge-0/0/5.312	Up	Sparse	4 2	NotDR	1	0
0 192.168.0.46						
lo0.2	Up	Sparse	4 2	DR	0	0
0 172.30.5.10						
lsi.1	Up	SparseDense	4 2	P2P	0	0
0						
ppd0.32770	Up	Sparse	4 2	P2P	0	0
0						
lsi.1	Up	SparseDense	6 2	P2P	0	0
0						

A Virtual LSI interface is created because of the **vrf-table-label** knob under the routing instance. **ppd0** is a PIM de-encapsulation interface. The local router is the RP and needs to de-encapsulate the register messages from the sources.

```
[edit]
lab@Sun# 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    0        None    0
239.0.0.0/24
Address family INET6
```

Listing the BGP neighbors with all the NLRI/ families enabled.

```
[edit]
lab@Sun# run show bgp summary
Groups: 5 Peers: 6 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp
State          Pending
inet.0
0              0          1344          599          0          0
bgp.l3vpn.0
0              0          16           16          0          0
inet6.0
0              0          48           48          0          0
bgp.mvpn.0
0              0          3            3          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
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@Sun# 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 Sun-to-Procyon
    to 172.30.0.10 via ge-0/0/4.118, label-
switched-path Sun-to-Procyon
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

```

Extensive output of the local originated Type 1 route shows the signaled PMSI Tunnel attribute.

```

[edit]
lab@Sun# 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[172.30.5.1:0:24
98:172.30.5.1]
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
--- (more) ---

```

Verify the P2MP RSVP sessions.

```
[edit]
lab@Sun# 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
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@Sun# 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.

```
[edit]
lab@Arcturus# 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-0/0/5.320       4 2                HPLGT           01:01:22
192.168.0.78
```

One PIM neighbor is established on the CE interface.

```
[edit]
lab@Arcturus# 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 stp-only mode.


```
[edit]
lab@Arcturus# 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 Arcturus-to-A-Centauri
    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 Arcturus-to-Rigel-1
    > to 172.30.0.30 via ge-0/0/4.145, label-
switched-path Arcturus-to-Rigel-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@Arcturus# 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@Arcturus# 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:invalid:
Neighbor                                     Inclusive Provider Tunnel
172.30.5.1                                   RSVP-TE P2MP:172.30.5.1,
2498,172.30.5.1
172.30.5.2                                   RSVP-TE P2MP:172.30.5.2,
53319,172.30.5.2
172.30.5.5
172.30.5.7

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 : C2-spoke
MVPN Mode : SPT-ONLY
Provider tunnel: I-P-tnl:invalid:
```

The output above displays the MVPN instances, Inclusive P-tunnel and the neighbors.

7) 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@Sun# set protocols mpls label-switched-path mcast-selective-
template template
```

```
[edit]
lab@Sun# set protocols mpls label-switched-path mcast-selective-
template bandwidth 60m
```

```
[edit]
lab@Sun# set protocols mpls label-switched-path mcast-selective-
template hop-limit 5
```

```
[edit]
```

```

lab@Sun# set protocols mpls label-switched-path mcast-selective-
template priority 5 5

[edit]
lab@Sun# set protocols mpls label-switched-path mcast-selective-
template link-protection

[edit]
lab@Sun# set protocols mpls label-switched-path mcast-selective-
template p2mp

[edit]
lab@Sun# set routing-instances C2-spoke provider-tunnel selective
tunnel-limit 5

[edit]
lab@Sun# 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@Sun# set routing-instances C2-spoke provider-tunnel selective
group 239.0.0.1/32 rsvp-te label-switched-path-template mcast-
selective-template

[edit]
lab@Sun# 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@Sun# set routing-instances C2-spoke provider-tunnel selective
group 239.0.0.2/32 rsvp-te label-switched-path-template mcast-
selective-template

[edit]
lab@Sun# commit
commit complete

```

R2 should be configured with same configuration for redundancy.

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** knob under **protocols mpls** stanza 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@Canopus# 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 Canopus-to-Vega
    to 172.30.0.13 via ge-0/0/4.123, label-
switched-path Canopus-to-Vega
::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 Canopus-to-Vega
    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 Canopus-to-Vega
    to 172.30.0.13 via ge-0/0/4.123, label-
switched-path Canopus-to-Vega
::ffff:172.30.5.7/128
    *[LDP/9] 00:08:55, metric 20
    > to 172.30.0.22 via ge-0/0/4.134, label-
switched-path Canopus-to-Vega
    to 172.30.0.13 via ge-0/0/4.123, label-
switched-path Canopus-to-Vega
::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 Canopus-to-Procyon-1
    to 172.30.0.26 via ge-0/0/4.136, label-
switched-path Canopus-to-Procyon-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@Canopus# set protocols bgp group ibgp family inet6-vpn unicast

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

On R8 the same configuration should be applied.

- 2) Configure the IPv6 BGP family on the Route Reflector.

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 family
inet6-vpn unicast

[edit]
root@route-reflector# set protocols bgp group cluster-2 family
inet6-vpn unicast

[edit]
root@route-reflector# commit
[edit protocols bgp group cluster-1 cluster]

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.

- a. Configure C3 VPN on R3.

```
[edit]
lab@Canopus# set interfaces lo0.2 family inet address
172.30.5.18/32

[edit]
lab@Canopus# set interfaces lo0.2 family inet6 address
fd17:f0f4:f691:5::12/128

[edit]
lab@Canopus# set routing-instances C3 instance-type vrf

[edit]
lab@Canopus# set routing-instances C3 interface ge-0/0/5.317

[edit]
lab@Canopus# set routing-instances C3 interface lo0.2
```

```

[edit]
lab@Canopus# set routing-instances C3 vrf-target target:54591:300

[edit]
lab@Canopus# set routing-instances C3 protocols bgp group ce peer-
as 64601

[edit]
lab@Canopus# set routing-instances C3 protocols bgp group ce
neighbor fc09:c0:ffee::a

[edit]
lab@Canopus# set routing-instances C3 protocols bgp group ce as-
override

[edit]
lab@Canopus# 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@Procyon# set interfaces lo0.2 family inet address
172.30.5.38/32

[edit]
lab@Procyon# set interfaces lo0.2 family inet6 address
fd17:f0f4:f691:5::26/128

[edit]
lab@Procyon# set routing-instances C3 instance-type vrf

[edit]
lab@Procyon# set routing-instances C3 interface lo0.2

[edit]
lab@Procyon# set routing-instances C3 interface ge-0/0/5.325

[edit]
lab@Procyon# set routing-instances C3 vrf-target target:54591:300

[edit]
lab@Procyon# set routing-instances C3 protocols bgp group ce peer-
as 64601

[edit]
lab@Procyon# set routing-instances C3 protocols bgp group ce
neighbor fc09:c0:ffee::e

[edit]

```

```
lab@Procyon# set routing-instances C3 protocols bgp group ce as-override
```

```
[edit]
```

```
lab@Procyon# commit
commit complete
```

Verify the operation of 6VPE sites between R3 and R8.

```
[edit]
lab@Canopus# run show bgp summary
Groups: 5 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp
State          Pending
inet.0
0              0          660          598          0          0
bgp.l3vpn.0
0              0          19          19          0          0
inet6.0
0              0          48          48          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
  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@Canopus# 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 Canopus-to-Procyon-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 Canopus-to-Procyon-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 Canopus-to-Procyon-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:5:2::/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 Canopus-to-Procyon-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 Canopus-to-Procyon-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 Canopus-to-Procyon-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 Canopus-to-Procyon-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 Canopus-to-Procyon-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 Canopus-to-Procyon-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
fd18:cccc:dddd:5:9::/80
        *[BGP/170] 00:21:40, localpref 100
        AS path: 64601 I
        > to fc09:c0:ffee::a via ge-0/0/5.317
fd18:cccc:dddd:5:a::/80
        *[BGP/170] 00:21:40, localpref 100
        AS path: 64601 I
        > to fc09:c0:ffee::a via ge-0/0/5.317
fd18:cccc:dddd:5:b::/80
        *[BGP/170] 00:21:40, localpref 100
        AS path: 64601 I
        > to fc09:c0:ffee::a via ge-0/0/5.317
fd18:cccc:dddd:5:c::/80
        *[BGP/170] 00:21:40, localpref 100
        AS path: 64601 I
        > to fc09:c0:ffee::a via ge-0/0/5.317
fd18:cccc:dddd:5:d::/80
        *[BGP/170] 00:21:40, localpref 100
        AS path: 64601 I
        > to fc09:c0:ffee::a via ge-0/0/5.317
fd18:cccc:dddd:5:e::/80
        *[BGP/170] 00:21:40, localpref 100
        AS path: 64601 I
        > to fc09:c0:ffee::a via ge-0/0/5.317
fd18:cccc:dddd:5:f::/80
        *[BGP/170] 00:21:40, localpref 100

```

```

AS path: 64601 I
> to fc09:c0:ffee::a via ge-0/0/5.317
fd17:f0f4:f691:5::26/128
[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 Canopus-to-Procyon-1
to 172.30.0.26 via ge-0/0/4.136, label-
switched-path Bypass->172.30.0.13->172.30.0.1
fe80::/64
*[Direct/0] 00:22:02
> via ge-0/0/5.317
fe80::2e21:7201:3dcd:2685/128
*[Local/0] 00:22:02
Local via ge-0/0/5.317
fe80::2e21:720f:fccd:2680/128
*[Direct/0] 00:22:02

```

The routes are advertised to the CE routers via the native IPv6 EBGp session.

```

[edit]
lab@Canopus# run show route advertising-protocol bgp
fc09:c0:ffee::a

C3.inet6.0: 23 destinations, 24 routes (23 active, 0 holddown, 0
hidden)
  Prefix                               Nexthop          MED      Lclpref    AS
path
* fc09:c0:ffee::c/126                 Self              I
  fd18:cccc:dddd:5::/80
*
54591 I
  fd18:cccc:dddd:5:1::/80
*
54591 I
  fd18:cccc:dddd:5:2::/80
*
54591 I
  fd18:cccc:dddd:5:3::/80
*
54591 I
  fd18:cccc:dddd:5:4::/80
*
54591 I
  fd18:cccc:dddd:5:5::/80
*
54591 I
  fd18:cccc:dddd:5:6::/80
*
54591 I
  fd18:cccc:dddd:5:7::/80
*
54591 I
  fd18:cccc:dddd:5:8::/80

```

```
*                                fc09:c0:ffee::a
54591 I
  fd18:cccc:dddd:5:9::/80
*                                fc09:c0:ffee::a
54591 I
  fd18:cccc:dddd:5:a::/80
*                                fc09:c0:ffee::a
54591 I
  fd18:cccc:dddd:5:b::/80
*                                fc09:c0:ffee::a
54591 I
  fd18:cccc:dddd:5:c::/80
*                                fc09:c0:ffee::a
54591 I
  fd18:cccc:dddd:5:d::/80
*                                fc09:c0:ffee::a
54591 I
  fd18:cccc:dddd:5:e::/80
*                                fc09:c0:ffee::a
54591 I
  fd18:cccc:dddd:5:f::/80
*                                fc09:c0:ffee::a
54591 I
```

Verify if the routes are present in the **C3.inet6.0** table at R8.

```
[edit]
lab@Procyon# 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 Procyon-to-Canopus-1
    > to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-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
    > 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 Procyon-to-Canopus-1
> to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
fd18:cccc:dddd:5:9::/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 Procyon-to-Canopus-1
> to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
fd18:cccc:dddd:5:a::/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 Procyon-to-Canopus-1
> to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
fd18:cccc:dddd:5:b::/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 Procyon-to-Canopus-1
> to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
fd18:cccc:dddd:5:c::/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 Procyon-to-Canopus-1
> to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
fd18:cccc:dddd:5:d::/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 Procyon-to-Canopus-1
> to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
fd18:cccc:dddd:5:e::/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 Procyon-to-Canopus-1
> to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
fd18:cccc:dddd:5:f::/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 Procyon-to-Canopus-1
> to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
fd17:f0f4:f691:5::12/128
[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 Procyon-to-Canopus-1
> to 172.30.0.37 via ge-0/0/4.158, label-
switched-path Procyon-to-Canopus-2
to 172.30.0.45 via ge-0/0/4.178, label-
switched-path Bypass->172.30.0.9->172.30.0.2
fe80::/64
*[Direct/0] 00:18:39
> via ge-0/0/5.325
fe80::fac0:101:45dc:2e85/128
*[Local/0] 00:18:39
Local via ge-0/0/5.325
fe80::fac0:10f:fcde:2e80/128
*[Direct/0] 00:18:39

```

Verify route advertisement to the CE router on R8.

```

[edit]
lab@Procyon# run show route advertising-protocol bgp
fc09:c0:ffee::e

C3.inet6.0: 23 destinations, 24 routes (23 active, 0 holddown, 0
hidden)
  Prefix                               Nexthop          MED      Lclpref    AS
path

```

```

* fc09:c0:ffee::8/126      Self      I
  fd18:cccc:dddd:5::/80
*                          fc09:c0:ffee::e
54591 I
  fd18:cccc:dddd:5:1::/80
*                          fc09:c0:ffee::e
54591 I
  fd18:cccc:dddd:5:2::/80
*                          fc09:c0:ffee::e
54591 I
  fd18:cccc:dddd:5:3::/80
*                          fc09:c0:ffee::e
54591 I
  fd18:cccc:dddd:5:4::/80
*                          fc09:c0:ffee::e
54591 I
  fd18:cccc:dddd:5:5::/80
*                          fc09:c0:ffee::e
54591 I
  fd18:cccc:dddd:5:6::/80
*                          fc09:c0:ffee::e
54591 I
  fd18:cccc:dddd:5:7::/80
*                          fc09:c0:ffee::e
54591 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. You can save the configuration on all devices.

Chapter Six solutions: L2VPN and VPLS configuration

The tasks included in this chapter introduce L2VPN and VPLS services to the MPLS network. LDP and BGP signaled L2VPN and VPLS 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 ends are the same, but may also be different. For example it is possible to have 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@Sun# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@Sun# set interfaces ge-0/0/3 encapsulation flexible-ethernet-
services

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@Sun# set interfaces ge-0/0/3.512 encapsulation vlan-ccc

[edit]
lab@Sun# set interfaces ge-0/0/3.512 vlan-id 512

[edit]
lab@Sun# set interfaces ge-0/0/3.513 encapsulation vlan-ccc

[edit]
lab@Sun# set interfaces ge-0/0/3.513 vlan-id 513

[edit]
lab@Sun# set interfaces ge-0/0/3.514 encapsulation vlan-ccc

[edit]
lab@Sun# set interfaces ge-0/0/3.514 vlan-id 514
```

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

vlan-ccc encapsulation configures the port for VLAN circuit cross connect.

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@Canopus# set protocols bgp group ibgp family l2vpn signaling

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

Configure both IBGP groups on the route reflector also.

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 family l2vpn signaling

[edit]
root@route-reflector# set protocols bgp group cluster-2 family l2vpn signaling

[edit]
root@route-reflector# commit
[edit protocols bgp group cluster-1 cluster]

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@Rigel# set routing-instances C2-l2vpn instance-type l2vpn

[edit]
lab@Rigel# set routing-instances C2-l2vpn interface ge-0/0/3.512

[edit]
lab@Rigel# set routing-instances C2-l2vpn interface ge-0/0/3.513

[edit]
lab@Rigel# set routing-instances C2-l2vpn vrf-target
target:54591:500
```

The route distinguisher is automatically generated because of the **route-distinguisher-id** knob configured under the *routing-options* stanza.

The actual L2VPN configuration is configured under the **protocols l2vpn** stanza.

```
[edit]
lab@Rigel# set routing-instances C2-l2vpn protocols l2vpn
encapsulation-type ethernet-vlan

[edit]
lab@Rigel# set routing-instances C2-l2vpn protocols l2vpn site
site-1 site-identifier 1

[edit]
lab@Rigel# set routing-instances C2-l2vpn protocols l2vpn site
site-1 interface ge-0/0/3.512

[edit]
lab@Rigel# set routing-instances C2-l2vpn protocols l2vpn site
site-1 interface ge-0/0/3.513

[edit]
lab@Rigel# 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@Canopus# set routing-instances C5-l2vpn instance-type l2vpn

[edit]
lab@Canopus# set routing-instances C5-l2vpn interface ge-0/0/3.512

[edit]
lab@Canopus# set routing-instances C5-l2vpn interface ge-0/0/3.514

[edit]
lab@Canopus# set routing-instances C5-l2vpn vrf-target
target:54591:500

[edit]
lab@Canopus# set routing-instances C5-l2vpn protocols l2vpn
encapsulation-type ethernet-vlan

[edit]
lab@Canopus# set routing-instances C5-l2vpn protocols l2vpn site
site-2 site-identifier 2

[edit]
lab@Canopus# set routing-instances C5-l2vpn protocols l2vpn site
site-2 interface ge-0/0/3.512

[edit]
lab@Canopus# set routing-instances C5-l2vpn protocols l2vpn site
site-2 interface ge-0/0/3.514

[edit]
lab@Canopus# 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@A-Centauri# set routing-instances C5-l2vpn instance-type l2vpn

[edit]
lab@A-Centauri# set routing-instances C5-l2vpn interface ge-
0/0/3.513

[edit]
lab@A-Centauri# set routing-instances C5-l2vpn interface ge-
0/0/3.514
```

```
[edit]
lab@A-Centauri# set routing-instances C5-l2vpn vrf-target
target:54591:500

[edit]
lab@A-Centauri# set routing-instances C5-l2vpn protocols l2vpn
encapsulation-type ethernet-vlan

[edit]
lab@A-Centauri# set routing-instances C5-l2vpn protocols l2vpn
site site-3 site-identifier 3

[edit]
lab@A-Centauri# set routing-instances C5-l2vpn protocols l2vpn
site site-3 interface ge-0/0/3.513

[edit]
lab@A-Centauri# set routing-instances C5-l2vpn protocols l2vpn
site site-3 interface ge-0/0/3.514

[edit]
lab@A-Centauri# 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@Rigel# run show l2vpn connections extensive
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS                    WE -- interface and instance
EM -- encapsulation mismatch     encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not
present                         -> -- only outbound connection is
CM -- control-word mismatch     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: C2-l2vpn
  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
Up trans
  2                                Type  St      Time last up      #
1                                rmt  Up      Oct 23 05:42:32 2013
    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
  Connection History:
```

```

Oct 23 05:42:32 2013 status update timer
Oct 23 05:42:32 2013 PE route changed
Oct 23 05:42:32 2013 Out lbl Update
800000
Oct 23 05:42:32 2013 In lbl Update
800001
Oct 23 05:42:32 2013 loc intf up ge-
0/0/3.512
3 rmt Up Oct 23 05:35:51 2013
1
Remote PE: 172.30.5.5, Negotiated control-word: Yes (Null)
Incoming label: 800002, Outgoing label: 800000
Local interface: ge-0/0/3.513, Status: Up, Encapsulation:
VLAN
Connection History:
Oct 23 05:35:51 2013 status update timer
Oct 23 05:35:51 2013 PE route changed
Oct 23 05:35:51 2013 Out lbl Update
800000
Oct 23 05:35:51 2013 In lbl Update
800002
Oct 23 05:35:51 2013 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)

Let's looking at R3's L2VPN connection table for the calculated labels.

```
[edit]
lab@Canopus# run show l2vpn connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS                    WE -- interface and instance
EM -- encapsulation mismatch     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-l2vpn
  Local site: site-2 (2)
    connection-site      Type  St      Time last up      #
Up trans
  1                      rmt  Up      Oct 23 05:39:31 2013
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
  3                      rmt  Up      Oct 23 05:35:51 2013
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
```


To find the Label-Base and the Label-offset you have to look at the L2VPN prefixes in **C5-12vpn.12vpn.0** routing table.

```
[edit]
lab@Canopus# run show route table C5-12vpn.12vpn.0 extensive

C5-12vpn.12vpn.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: C5-12vpn-12vpn
    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-C5-12vpn-12vpn
      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
```

```

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

```

```

                                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.

$800000 \text{ (Label-Base)} + 2 \text{ (Local-site-Id)} - 1 \text{ (Remote-label-offset)} = 800001$

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** stanza.

Because LDP tunneling over RSVP is also using extended LDP sessions, Loopback interface is configured as LDP interface in Chapter Four.

```

[edit]
lab@Sun# 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
}
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 L2 circuit 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 as listed in Table 22.

```
[edit]
lab@Sun# set protocols l2circuit neighbor 172.30.5.8 interface ge-
0/0/3.512 virtual-circuit-id 512
```

```
[edit]
lab@Sun# set protocols l2circuit neighbor 172.30.5.6 interface ge-
0/0/3.513 virtual-circuit-id 513
```

```
[edit]
lab@Sun# 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@Vega# set protocols l2circuit neighbor 172.30.5.1 interface
ge-0/0/3.513 virtual-circuit-id 513
```

```
[edit]
lab@Vega# set protocols l2circuit neighbor 172.30.5.8 interface
ge-0/0/3.514 virtual-circuit-id 514
```

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

c. R8.

R8 should have a l2circuit to R1 and R6.

```
[edit]
lab@Procyon# set protocols l2circuit neighbor 172.30.5.1 interface
ge-0/0/3.512 virtual-circuit-id 512
```

```
[edit]
lab@Procyon# set protocols l2circuit neighbor 172.30.5.6 interface
ge-0/0/3.514 virtual-circuit-id 514
```

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

Verify the operation of the LDP L2VPN on R1.

```
[edit]
lab@Sun# 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 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      Oct 23 04:46:34 2013
1
  Remote PE: 172.30.5.6, Negotiated control-word: Yes (Null)
  Incoming label: 303984, Outgoing label: 306448
  Negotiated PW status TLV: No
  Local interface: ge-0/0/3.513, Status: Up, Encapsulation:
VLAN
Neighbor: 172.30.5.8
  Interface          Type  St      Time last up      #
Up trans
  ge-0/0/3.512(vc 512)  rmt  Up      Oct 23 04:50:01 2013
1
  Remote PE: 172.30.5.8, Negotiated control-word: Yes (Null)
  Incoming label: 303968, Outgoing label: 305600
  Negotiated PW status TLV: No
```

```
Local interface: ge-0/0/3.512, Status: Up, Encapsulation:
VLAN
```

Both l2circuit connections are UP. Also the Incoming and Outgoing labels are shown.

```
[edit]
lab@Sun# run show ldp database session 172.30.5.6 l2circuit
extensive
Input label database, 172.30.5.1:0--172.30.5.6:0
  Label      Prefix
  306448     L2CKT CtrlWord VLAN VC 513
              MTU: 1500 Requested VLAN ID: 513
              State: Active
              Age: 22:30
              VCCV Control Channel types:
                  MPLS router alert label
                  MPLS PW label with TTL=1
              VCCV Control Verification types:
                  LSP ping

Output label database, 172.30.5.1:0--172.30.5.6:0
  Label      Prefix
  303984     L2CKT CtrlWord VLAN VC 513
              MTU: 1500 Requested VLAN ID: 513
              State: Active
              Age: 23:58
              VCCV Control Channel types:
                  MPLS router alert label
                  MPLS PW label with TTL=1
              VCCV Control Verification types:
                  LSP ping
```

```
[edit]
lab@Sun# run show ldp database session 172.30.5.8 l2circuit
extensive
Input label database, 172.30.5.1:0--172.30.5.8:0
  Label      Prefix
  305600     L2CKT CtrlWord VLAN VC 512
              MTU: 1500 Requested VLAN ID: 512
              State: Active
              Age: 19:09
              VCCV Control Channel types:
                  MPLS router alert label
                  MPLS PW label with TTL=1
              VCCV Control Verification types:
                  LSP ping

Output label database, 172.30.5.1:0--172.30.5.8:0
  Label      Prefix
  303968     L2CKT CtrlWord VLAN VC 512
              MTU: 1500 Requested VLAN ID: 512
              State: Active
              Age: 19:14
              VCCV Control Channel types:
```

```

MPLS router alert label
MPLS PW label with TTL=1
VCCV Control Verification types:
LSP ping

```

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.

```

[edit]
lab@Vega# 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 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      Oct 23 04:46:33 2013
1
    Remote PE: 172.30.5.1, Negotiated control-word: Yes (Null)
    Incoming label: 306448, Outgoing label: 303984
    Negotiated PW status TLV: No
    Local interface: ge-0/0/3.513, Status: Up, Encapsulation:
VLAN
Neighbor: 172.30.5.8
  Interface                Type  St      Time last up      #
Up trans
  ge-0/0/3.514(vc 514)     rmt   Up      Oct 23 04:49:44 2013
1
    Remote PE: 172.30.5.8, Negotiated control-word: Yes (Null)
    Incoming label: 306464, Outgoing label: 305616
    Negotiated PW status TLV: No
    Local interface: ge-0/0/3.514, Status: Up, Encapsulation:
VLAN
[edit]
lab@Vega# 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

```

```

Label      Prefix
303984     L2CKT CtrlWord VLAN VC 513
           MTU: 1500 Requested VLAN ID: 513
           State: Active
           Age: 19:58
           VCCV Control Channel types:
             MPLS router alert label
             MPLS PW label with TTL=1
           VCCV Control Verification types:
             LSP ping

Output label database, 172.30.5.6:0--172.30.5.1:0
Label      Prefix
306448     L2CKT CtrlWord VLAN VC 513
           MTU: 1500 Requested VLAN ID: 513
           State: Active
           Age: 20:05
           VCCV Control Channel types:
             MPLS router alert label
             MPLS PW label with TTL=1
           VCCV Control Verification types:
             LSP ping

```

```

[edit]
lab@Vega# 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
Label      Prefix
305616     L2CKT CtrlWord VLAN VC 514
           MTU: 1500 Requested VLAN ID: 514
           State: Active
           Age: 17:04
           VCCV Control Channel types:
             MPLS router alert label
             MPLS PW label with TTL=1
           VCCV Control Verification types:
             LSP ping

Output label database, 172.30.5.6:0--172.30.5.8:0
Label      Prefix
306464     L2CKT CtrlWord VLAN VC 514
           MTU: 1500 Requested VLAN ID: 514
           State: Active
           Age: 19:56
           VCCV Control Channel types:
             MPLS router alert label
             MPLS PW label with TTL=1
           VCCV Control Verification types:
             LSP ping

```

Verify the operation of the LDP L2VPN on R8.

```

[edit]
lab@Procyon# 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	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	Oct 23 04:50:01 2013	1

Remote PE: 172.30.5.1, Negotiated control-word: Yes (Null)
Incoming label: 305600, Outgoing label: 303968
Negotiated PW status TLV: No
Local interface: ge-0/0/3.512, Status: Up, Encapsulation:

VLAN

Neighbor: 172.30.5.6

Interface	Type	St	Time last up	#
Up trans				
ge-0/0/3.514(vc 514)	rmt	Up	Oct 23 04:49:43 2013	2

Remote PE: 172.30.5.6, Negotiated control-word: Yes (Null)
Incoming label: 305616, Outgoing label: 306464
Negotiated PW status TLV: No
Local interface: ge-0/0/3.514, Status: Up, Encapsulation:

VLAN

[edit]

lab@Procyon# **run show ldp database session 172.30.5.1 l2circuit**

Input label database, 172.30.5.8:0--172.30.5.1:0

Label	Prefix
303968	L2CKT CtrlWord VLAN VC 512

Output label database, 172.30.5.8:0--172.30.5.1:0

Label	Prefix
305600	L2CKT CtrlWord VLAN VC 512

[edit]

lab@Procyon# **run show ldp database session 172.30.5.6 l2circuit**

Input label database, 172.30.5.8:0--172.30.5.6:0

Label	Prefix
306464	L2CKT CtrlWord VLAN VC 514

```
Output label database, 172.30.5.8:0--172.30.5.6:0
Label      Prefix
305616     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.

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.

The PE-CE interfaces have to be configured for vpls encapsulation.

```
[edit]
lab@Sirius# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@Sirius# set interfaces ge-0/0/3 encapsulation flexible-
ethernet-services

[edit]
lab@Sirius# set interface ge-0/0/3 unit 600 vlan-id 600

[edit]
lab@Sirius# set interface ge-0/0/3 unit 600 encapsulation vlan-
vpls

[edit]
lab@Sirius# set interface ge-0/0/3 unit 601 vlan-id 601

[edit]
lab@Sirius# set interface ge-0/0/3 unit 601 encapsulation vlan-
vpls

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

- 2) Configure customer C5 VPLS.

- a. Configure VPLS routing instance on R2.

```
[edit]
lab@Sirius# set routing-instances C5-vpls instance-type vpls

[edit]
lab@Sirius# set routing-instances C5-vpls vlan-id all
```

vlan-id all creates a separate bridge domain for all VLANs. This will automatically enable L2 learning for VLAN600 and VLAN601.

```
[edit]
lab@Sirius# set routing-instances C5-vpls interface ge-0/0/3.600

[edit]
lab@Sirius# set routing-instances C5-vpls interface ge-0/0/3.601

[edit]
lab@Sirius# set routing-instances C5-vpls vrf-target
target:54591:501
```

Configure **protocol vpls** under the routing instance.

```
[edit]
lab@Sirius# 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.

```
[edit]
lab@Sirius# set routing-instances C5-vpls protocols vpls no-
tunnel-services
```

no-tunnel-services 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@Sirius# set routing-instances C5-vpls protocols vpls site
site-4 site-identifier 4
```

```
[edit]
lab@Sirius# 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@Canopus# set routing-instances C5-vpls instance-type vpls

[edit]
lab@Canopus# set routing-instances C5-vpls vlan-id all

[edit]
lab@Canopus# set routing-instances C5-vpls interface ge-0/0/3.600
```

```
[edit]
lab@Canopus# set routing-instances C5-vpls interface ge-0/0/3.601
```

```
[edit]
lab@Canopus# set routing-instances C5-vpls vrf-target
target:54591:501
```

```
[edit]
lab@Canopus# set routing-instances C5-vpls protocols vpls site-
range 8
```

```
[edit]
lab@Canopus# set routing-instances C5-vpls protocols vpls no-
tunnel-services
```

```
[edit]
lab@Canopus# set routing-instances C5-vpls protocols vpls site
site-5 site-identifier 5
```

```
[edit]
lab@Canopus# set routing-instances C5-vpls protocols vpls site
site-5 multi-homing
```

multi-homing 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@Canopus# 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 at actually 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@Canopus# commit
commit complete
```

- c. Configure VPLS routing instance on R4.

R4 is the backup for the multi-homed site.

```
[edit]
lab@Arcturus# set routing-instances C5-vpls instance-type vpls
```

```
[edit]
lab@Arcturus# set routing-instances C5-vpls vlan-id all
```

```
[edit]
lab@Arcturus# set routing-instances C5-vpls interface ge-0/0/3.601
```

```
[edit]
lab@Arcturus# set routing-instances C5-vpls interface ge-0/0/3.600

[edit]
lab@Arcturus# set routing-instances C5-vpls vrf-target
target:54591:501

[edit]
lab@Arcturus# set routing-instances C5-vpls protocols vpls site-
range 8

[edit]
lab@Arcturus# set routing-instances C5-vpls protocols vpls no-
tunnel-services

[edit]
lab@Arcturus# set routing-instances C5-vpls protocols vpls site
site-5 site-identifier 5
```

The same site-identifier as on R3 is configured on R4.

```
[edit]
lab@Arcturus# set routing-instances C5-vpls protocols vpls site
site-5 multi-homing

[edit]
lab@Arcturus# set routing-instances C5-vpls protocols vpls site
site-5 site-preference backup
```

The **backup** knob sets the preference and respectively the local preference of the announcements from R4 to 1, which is lower then the value used by R3.

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

d. Configure VPLS routing instance on R5.

```
[edit]
lab@A-Centauri# set routing-instances C5-vpls instance-type vpls

[edit]
lab@A-Centauri# set routing-instances C5-vpls vlan-id all

[edit]
lab@A-Centauri# set routing-instances C5-vpls interface ge-
0/0/3.600

[edit]
lab@A-Centauri# set routing-instances C5-vpls interface ge-
0/0/3.601

[edit]
lab@A-Centauri# set routing-instances C5-vpls vrf-target
target:54591:501
```

```
[edit]
lab@A-Centauri# set routing-instances C5-vpls protocols vpls site-
range 8

[edit]
lab@A-Centauri# set routing-instances C5-vpls protocols vpls no-
tunnel-services

[edit]
lab@A-Centauri# set routing-instances C5-vpls protocols vpls site
site-6 site-identifier 6

[edit]
lab@A-Centauri# commit
commit complete
```

Verify the operation of customer C5 BGP VPLS on R2

```
[edit]
lab@Sirius# run show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS                    WE -- interface and instance
EM -- encapsulation mismatch     encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not
present                         -> -- only outbound connection is
CM -- control-word mismatch     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
Up trans
  5
1
  Type  St      Time last up      #
  rmt   Up      Oct 23 07:33:33 2013
  Remote PE: 172.30.5.3, Negotiated control-word: No
  Incoming label: 262149, Outgoing label: 262148
  Local interface: lsi.1048590, Status: Up, Encapsulation:
VPLS
  Description: Intf - vpls C5-vpls local site 4 remote site
  5
  6
1
  Type  St      Time last up      #
  rmt   Up      Oct 23 07:25:51 2013
  Remote PE: 172.30.5.5, Negotiated control-word: No
  Incoming label: 262150, Outgoing label: 262148
```



```

Local interface: lsi.1048588, Status: Up, Encapsulation:
VPLS
Description: Intf - vpls C5-vpls local site 4 remote site
6

```

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 have a lower local preference value of 1.

```

[edit]
lab@Arcturus# 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)
* 172.30.5.4:6:5:1/96 (1 entry, 1 announced)
  BGP group ibgp type Internal
    Route Distinguisher: 172.30.5.4:6
    Label-base: 262145, range: 8
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 1
    AS path: [54591] I
    Communities: target:54591:501 Layer2-info: encaps:VPLS,
control flags:, mtu: 0, site preference: 1

[edit]
lab@Canopus# 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)
* 172.30.5.3:6:5:1/96 (1 entry, 1 announced)
  BGP group ibgp type Internal
    Route Distinguisher: 172.30.5.3:6
    Label-base: 262145, range: 8
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 65535
    AS path: [54591] I
    Communities: target:54591:501 Layer2-info: encaps:VPLS,
control flags:, mtu: 0, site preference: 65535

```

The announce from R3 is with the maximum local preference value of 65535. The same value is set for the site preference.

Verify the operation of C5 VPLS on R3.

```
[edit]
lab@Canopus# run show vpls connections extensive
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS                    WE -- interface and instance
EM -- encapsulation mismatch     NP -- interface hardware not
encaps not same                 present
VC-Dn -- Virtual circuit down    -> -- only outbound connection is
CM -- control-word mismatch      up
CN -- circuit not provisioned    <- -- only inbound connection is
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
LN -- local site not designated  LM -- local site ID not minimum
designated                      RN -- remote site not designated
RM -- remote site ID not minimum
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-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          100
  connection-site      Type      St          Time last up      #
Up trans
  4                rmt      Up          Oct 23 06:59:46 2013
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
    Connection History:
        Oct 23 06:59:46 2013  status update timer
        Oct 23 06:59:46 2013  loc intf up
lsi.1048576
        Oct 23 06:59:46 2013  PE route changed
        Oct 23 06:59:46 2013  Out lbl Update
262149
        Oct 23 06:59:46 2013  In lbl Update
262148
        Oct 23 06:59:46 2013  loc intf down
5
        rmt      RN
6
        rmt      Up      Oct 23 07:25:52 2013
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
    Connection History:
        Oct 23 07:25:52 2013  status update timer
        Oct 23 07:25:52 2013  loc intf up
lsi.1048577
        Oct 23 07:25:52 2013  PE route changed
        Oct 23 07:25:52 2013  Out lbl Update
262149
        Oct 23 07:25:52 2013  In lbl Update
262150
        Oct 23 07:25:52 2013  loc intf down

```

The connections the the remote sites to R2 and R5 are UP. The connection to site 5, with R4 is shown as **RN -- remote site not designated**

```

[edit]
lab@Canopus# run show route forwarding-table family vpls
Routing table: C5-vpls.vpls
VPLS:
Destination          Type RtRef Next hop          Type Index NhRef
Netif
default              perm    0
ge-0/0/3.600         user    0
ge-0/0/3.601         user    0
lsi.1048631          user    0
lsi.1048635          user    0
00:23:9c:8b:6c:96/48 dynm    0
                                172.30.0.13      Push 262149 751
2 ge-0/0/4.123
00:23:9c:8b:6c:97/48 dynm    0
                                ucst   696      3
ge-0/0/3.601

```

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 Unknown Unicast, Broadcasts 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@Arcturus# run show vpls connections extensive
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid          NC -- interface encapsulation not
CCC/TCC/VPLS                        WE -- interface and instance
EM -- encapsulation mismatch         encaps not same
VC-Dn -- Virtual circuit down       NP -- interface hardware not
present                             -> -- only outbound connection is
CM -- control-word mismatch          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
RD -- remote site signaled down     SC -- local and remote site ID
collision                           LN -- local site not designated
LN -- local site not designated     LM -- local site ID not minimum
designated                           RN -- remote site not designated
RN -- remote site not designated    RM -- remote site ID not minimum
designated                           XX -- unknown connection status
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-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.

```
[edit]
lab@Canopus# deactivate routing-instances C5-vpls

[edit]
lab@Canopus# commit
[edit interfaces ge-0/0/3]
'unit 600'
    interface needs to be in a VPLS routing instance to support
family VPLS
[edit interfaces ge-0/0/3]
'unit 601'
    interface needs to be in a VPLS routing instance to support
family VPLS
commit complete
```

Verify again the table with VPLS connections on R4.

```
[edit]
lab@Arcturus# run show vpls connections extensive
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS                    WE -- interface and instance
EM -- encapsulation mismatch     encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not
present                          -> -- only outbound connection is
CM -- control-word mismatch     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-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.1049858          4          Intf - vpls C5-vpls local site 5
remote site 4
    lsi.1049859          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      Oct 23 07:36:59 2013
1
  Remote PE: 172.30.5.2, Negotiated control-word: No
  Incoming label: 262148, Outgoing label: 262149
  Local interface: lsi.1049858, Status: Up, Encapsulation:
VPLS
  Description: Intf - vpls C5-vpls local site 5 remote site
4
  Connection History:
    Oct 23 07:36:59 2013  status update timer
    Oct 23 07:36:59 2013  loc intf up
lsi.1049858
    Oct 23 07:36:59 2013  PE route changed
    Oct 23 07:36:59 2013  Out lbl Update
262149
    Oct 23 07:36:59 2013  In lbl Update
262148
    Oct 23 07:36:59 2013  loc intf down
  6                  rmt  Up      Oct 23 07:36:59 2013
1
  Remote PE: 172.30.5.5, Negotiated control-word: No
  Incoming label: 262150, Outgoing label: 262149
  Local interface: lsi.1049859, Status: Up, Encapsulation:
VPLS
  Description: Intf - vpls C5-vpls local site 5 remote site
6
  Connection History:
    Oct 23 07:36:59 2013  status update timer
    Oct 23 07:36:59 2013  loc intf up
lsi.1049859

```

	Oct 23 07:36:59 2013	PE route changed
	Oct 23 07:36:59 2013	Out lbl Update
262149	Oct 23 07:36:59 2013	In lbl Update
262150	Oct 23 07:36:59 2013	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@A-Centauri# run show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS                    WE -- interface and instance
EM -- encapsulation mismatch     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-6 (6)
    connection-site
Up trans
  4
1
  Type  St      Time last up  #
  rmt   Up      Oct 23 07:25:51 2013
  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
  5
1
  Type  St      Time last up  #
  rmt   Up      Oct 23 07:28:34 2013
  Remote PE: 172.30.5.3, Negotiated control-word: No
  Incoming label: 262149, Outgoing label: 262150
```



```

    Local interface: lsi.1048578, Status: Up, Encapsulation:
VPLS
    Description: Intf - vpls C5-vpls local site 6 remote site
5

```

3) Configure the additional PE-CE interfaces on R1, R6, R7 and R8 for LDP VPLS.

```

[edit]
lab@Sun# set interfaces ge-0/0/3.700 vlan-id 700

[edit]
lab@Sun# set interfaces ge-0/0/3.700 encapsulation vlan-vpls

[edit]
lab@Sun# set interfaces ge-0/0/3.701 vlan-id 701

[edit]
lab@Sun# set interfaces ge-0/0/3.701 encapsulation vlan-vpls

[edit]
lab@Sun# commit
commit complete

```

4) Configure the customer C6 LDP VPLS instances.

Unlike the l2circuit configuration, LDP VPLS is configured as a routing-instance **type vpls**.

a. Configure R1.

```

[edit]
lab@Sun# set routing-instances C6-vpls instance-type vpls

[edit]
lab@Sun# set routing-instances C6-vpls vlan-id all

[edit]
lab@Sun# set routing-instances C6-vpls interface ge-0/0/3.700

[edit]
lab@Sun# set routing-instances C6-vpls interface ge-0/0/3.701

[edit]
lab@Sun# set routing-instances C6-vpls protocols vpls no-tunnel-
services

```

NOTE: **no-tunnel-services** is needed also.

```

[edit]
lab@Sun# set routing-instances C6-vpls protocols vpls
encapsulation-type ethernet-vlan

[edit]
lab@Sun# set routing-instances C6-vpls protocols vpls vpls-id 600

```

vpls-id is needed to identify the virtual circuit for the VPLS site. It should match on the remote sites.

```
[edit]
lab@Sun# set routing-instances C6-vpls protocols vpls neighbor
172.30.5.6
```

The neighbors should be configured manually.

```
[edit]
lab@Sun# set routing-instances C6-vpls protocols vpls neighbor
172.30.5.7 backup-neighbor 172.30.5.8
```

The site attached to R7 is multi-homed with R8. With LDP VPLS the backup is configured on the remote sites. With LDP VPLS, multi-homed 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@Sun# set routing-instances C6-vpls protocols vpls neighbor
172.30.5.7 revert-time 60
```

The **revert-time** will delay the switching from the backup to the primary site, once the primary becomes available again (preempt delay).

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

b. Configure R6.

```
[edit]
lab@Vega# set routing-instances C6-vpls instance-type vpls

[edit]
lab@Vega# set routing-instances C6-vpls vlan-id all

[edit]
lab@Vega# set routing-instances C6-vpls interface ge-0/0/3.700

[edit]
lab@Vega# set routing-instances C6-vpls interface ge-0/0/3.701

[edit]
lab@Vega# set routing-instances C6-vpls protocols vpls no-tunnel-
services

[edit]
lab@Vega# set routing-instances C6-vpls protocols vpls vpls-id 600

[edit]
lab@Vega# set routing-instances C6-vpls protocols vpls
encapsulation-type ethernet-vlan
```

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@Vega# set routing-instances C6-vpls protocols vpls neighbor
172.30.5.1
```

```
[edit]
lab@Vega# set routing-instances C6-vpls protocols vpls neighbor
172.30.5.7 backup-neighbor 172.30.5.8
```

```
[edit]
lab@Vega# set routing-instances C6-vpls protocols neighbor
172.30.5.7 revert-time 60
```

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

c. Configure R7.

CE attached to R7 is multi-homed on R8 for redundancy.

```
[edit]
lab@Rigel# set routing-instances C6-vpls instance-type vpls
```

```
[edit]
lab@Rigel# set routing-instances C6-vpls vlan-id all
```

```
[edit]
lab@Rigel# set routing-instances C6-vpls interface ge-0/0/3.700
```

```
[edit]
lab@Rigel# set routing-instances C6-vpls interface ge-0/0/3.701
```

```
[edit]
lab@Rigel# set routing-instances C6-vpls protocols vpls no-tunnel-
services
```

```
[edit]
lab@Rigel# set routing-instances C6-vpls protocols vpls
encapsulation-type ethernet-vlan
```

```
[edit]
lab@Rigel# set routing-instances C6-vpls protocols vpls vpls-id
600
```

```
[edit]
lab@Rigel# set routing-instances C6-vpls protocols vpls neighbor
172.30.5.1
```

```
[edit]
lab@Rigel# set routing-instances C6-vpls protocols vpls neighbor
172.30.5.6
```

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

Unlike BGP VPLS with LDP VPLS no multi-homed configuration is needed on multi-homed sites.

d. Configure R8.

R8 is the other multi-homed site PE.

```
[edit]
lab@Procyon# set routing-instances C6-vpls instance-type vpls

[edit]
lab@Procyon# set routing-instances C6-vpls vlan-id all

[edit]
lab@Procyon# set routing-instances C6-vpls interface ge-0/0/3.700

[edit]
lab@Procyon# set routing-instances C6-vpls interface ge-0/0/3.701

[edit]
lab@Procyon# set routing-instances C6-vpls protocols vpls no-
tunnel-services

[edit]
lab@Procyon# set routing-instances C6-vpls protocols vpls vpls-id
600

[edit]
lab@Procyon# set routing-instances C6-vpls protocols vpls neighbor
172.30.5.1

[edit]
lab@Procyon# set routing-instances C6-vpls protocols vpls neighbor
172.30.5.6

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

Again as with R7, on R8 no configuration for multi-homing is applied.

Verify the operation of customer C6 LDP VPLS on R1.

```
[edit]
lab@Sun# run show ldp neighbor
```

Address	Interface	Label space ID	Hold time
172.30.5.6	100.0	172.30.5.6:0	43
172.30.5.7	100.0	172.30.5.7:0	37
172.30.5.8	100.0	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@Sun# run show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS                    WE -- interface and instance
EM -- encapsulation mismatch     encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not
present                         -> -- only outbound connection is
CM -- control-word mismatch      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: C6-vpls
VPLS-id: 600
Neighbor                Type  St      Time last up      #
Up trans
172.30.5.6(vpls-id 600)  rmt  Up      Nov  1 05:35:19 2013
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
172.30.5.7(vpls-id 600)  rmt  Up      Nov  1 05:35:24 2013
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
    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@Sun# run show ldp database l2circuit session 172.30.5.6
Input label database, 172.30.5.1:0--172.30.5.6:0
  Label      Prefix
262145      L2CKT NoCtrlWord VLAN VC 600
299776      L2CKT CtrlWord VLAN VC 513

Output label database, 172.30.5.1:0--172.30.5.6:0
  Label      Prefix
262145      L2CKT NoCtrlWord VLAN VC 600
299776      L2CKT CtrlWord VLAN VC 513

```

```

[edit]
lab@Sun# run show ldp database l2circuit session 172.30.5.7
Input label database, 172.30.5.1:0--172.30.5.7:0
  Label      Prefix
262145      L2CKT NoCtrlWord VLAN VC 600

Output label database, 172.30.5.1:0--172.30.5.7:0
  Label      Prefix
262148      L2CKT NoCtrlWord VLAN VC 600

```

```

[edit]
lab@Sun# run show ldp database l2circuit session 172.30.5.8
Input label database, 172.30.5.1:0--172.30.5.8:0
  Label      Prefix
262145      L2CKT NoCtrlWord VLAN VC 600
299776      L2CKT CtrlWord VLAN VC 512

Output label database, 172.30.5.1:0--172.30.5.8:0
  Label      Prefix
299792      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@Sun# run show route forwarding-table family vpls
Routing table: C6-vpls.vpls
VPLS:

```

Destination	Type	RtRef	Next hop	Type	Index	NhRef
Netif						
default	perm	0		rjct	760	1
ge-0/0/3.700	user	0		comp	709	3
ge-0/0/3.701	user	0		comp	709	3
lsi.1048576	user	0		comp	768	3
lsi.1048579	user	0		comp	768	3
00:23:9c:8b:6c:95/48 dynm		0		ucst	710	3
ge-0/0/3.701						

The forwarding table of the customer VPLS table is shown above.

Verify the operation of customer C6 LDP VPLS on R7.

```
[edit]
lab@Rigel# 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: C6-vpls
VPLS-id: 600
Neighbor                                Type  St      Time last up      #
Up trans
```

```

172.30.5.1(vpls-id 600)    rmt    Up      Nov  1 05:35:28 2013
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.1048578, Status: Up, Encapsulation:
VLAN
  Description: Intf - vpls C6-vpls neighbor 172.30.5.1 vpls-
id 600
172.30.5.6(vpls-id 600)    rmt    Up      Nov  1 05:26:07 2013
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.1048576, Status: Up, Encapsulation:
VLAN
  Description: Intf - vpls C6-vpls neighbor 172.30.5.6 vpls-
id 600

```

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@Procyon# run show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid          NC -- interface encapsulation not
CCC/TCC/VPLS                        WE -- interface and instance
EM -- encapsulation mismatch         NP -- interface hardware not
encaps not same                     present
VC-Dn -- Virtual circuit down       -> -- only outbound connection is
CM -- control-word mismatch          up
CN -- circuit not provisioned        <- -- only inbound connection is
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
LN -- local site not designated      LM -- local site ID not minimum
RN -- remote site not designated     RM -- remote site ID not minimum
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: C6-vpls
VPLS-id: 600
Neighbor      Type  St      Time last up  #
Up trans
172.30.5.1(vpls-id 600)  rmt  OL      17:05:00      1
172.30.5.6(vpls-id 600)  rmt  OL      17:05:00      2
```

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@Rigel# deactivate routing-instances C6-vpls

[edit]
```

```

lab@Rigel# commit
[edit interfaces ge-0/0/3]
  'unit 700'
    interface needs to be in a VPLS routing instance to support
family VPLS
[edit interfaces ge-0/0/3]
  'unit 701'
    interface needs to be in a VPLS routing instance to support
family VPLS
commit complete

```

Verify the vpls connection table on R8 again.

```

[edit]
lab@Procyon# 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
VM -- VLAN ID mismatch

Legend for interface status
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.1048576                               Intf - vpls C6-vpls neighbor
172.30.5.1 vpls-id 600
lsi.1048577                               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      Nov  1 05:37:58 2013
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.1048576, Status: Up, Encapsulation:
VLAN
  Description: Intf - vpls C6-vpls neighbor 172.30.5.1 vpls-
id 600
  Connection History:
    Nov  1 05:37:58 2013  status update timer
    Nov  1 05:37:58 2013  PE route changed
    Nov  1 05:37:58 2013  Out lbl Update
262147
    Nov  1 05:37:58 2013  In lbl Update
262145
    Nov  1 05:37:58 2013  loc intf up
lsi.1048576                               Up
  172.30.5.6(vpls-id 600)    rmt   Nov  1 05:37:59 2013
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.1048577, Status: Up, Encapsulation:
VLAN
  Description: Intf - vpls C6-vpls neighbor 172.30.5.6 vpls-
id 600
  Connection History:
    Nov  1 05:37:59 2013  status update timer
    Nov  1 05:37:59 2013  PE route changed
    Nov  1 05:37:59 2013  Out lbl Update
262147
    Nov  1 05:37:59 2013  In lbl Update
262146
    Nov  1 05:37:59 2013  loc intf up
lsi.1048577

```

Now both neighbors to R1 and R6 are in UP state.

5) 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.

[edit]

```
lab@Sirius# set routing-instances C5-vpls protocols vpls mac-  
table-size 200
```

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

- 6) 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.

```
[edit]  
lab@Sun# set routing-instances C6-vpls protocols vpls mac-table-  
size 100
```

```
[edit]  
lab@Sun# set routing-instances C6-vpls protocols vpls mac-table-  
size packet-action drop
```

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

- 7) 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@Rigel# set interfaces ge-0/0/3.600 vlan-id 600  
  
[edit]  
lab@Rigel# set interfaces ge-0/0/3.600 encapsulation vlan-ccc  
  
[edit]  
lab@Rigel# commit  
commit complete
```

Modify the L2VPN routing instance.

```
[edit]  
lab@Rigel# set routing-instances C5-l2vpn interface ge-0/0/3.600  
  
[edit]  
lab@Rigel# 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@Rigel# commit
commit complete
```

8) 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

a. Configure lt- interface.

You need to 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@Sirius# set interfaces lt-0/0/0.0 vlan-id 600
```

Both lt- units are configured to accept 802.1q tagged packets.

```
[edit]
lab@Sirius# set interfaces lt-0/0/0.0 encapsulation vlan-ccc
```

unit 0 is configured for **vlan-ccc** encapsulation. This unit will be attached to the L2VPN routing instance.

```
[edit]
lab@Sirius# set interfaces lt-0/0/0.0 peer-unit 1
```

unit 0 is peered with **unit 1**.

```
[edit]
lab@Sirius# set interfaces lt-0/0/0.1 vlan-id 600
```

```
[edit]
lab@Sirius# set interfaces lt-0/0/0.1 encapsulation vlan-vpls
```

unit 1 is configured for **vlan-vpls** encapsulation. This unit will be attached to the VPLS routing instance.

```
[edit]
lab@Sirius# 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@Sirius# set routing-instances C5-l2vpn instance-type l2vpn

[edit]
lab@Sirius# set routing-instances C5-l2vpn interface lt-0/0/0.0
```

unit 0 of **lt-** is the only interface assigned to that routing instance

```
[edit]
lab@Sirius# set routing-instances C5-l2vpn vrf-target
target:54591:500
```

The route target matches the route target configured on R7.

```
[edit]
lab@Sirius# set routing-instances C5-l2vpn protocols l2vpn
encapsulation-type ethernet-vlan

[edit]
lab@Sirius# set routing-instances C5-l2vpn protocols l2vpn site
site-4 site-identifier 4
```

The local site identifier is configured as 4.

```
[edit]
lab@Sirius# set routing-instances C5-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 forth interface on R7 site 1.

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

Verify the l2vpn connection table.

```
[edit]
lab@Sirius# 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
RS -- remote site standby       SN -- Static Neighbor
VM -- VLAN ID mismatch

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

Instance: C5-l2vpn
  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      Nov  1 06:30:56 2013
  2
    Remote PE: 172.30.5.7, Negotiated control-word: Yes (Null)
    Incoming label: 800000, Outgoing label: 800003
    Local interface: lt-0/0/0.0, Status: Up, Encapsulation: VLAN
```

Connection History:						
	Nov	1	06:30:56	2013	status update timer	
	Nov	1	06:30:47	2013	rmt side marked VC Dn	
	Nov	1	06:30:38	2013	status update timer	
	Nov	1	06:30:38	2013	PE route changed	
	Nov	1	06:30:38	2013	Out lbl Update	
800003						
	Nov	1	06:30:38	2013	In lbl Update	
800000						
	Nov	1	06:30:38	2013	loc intf up	lt-
0/0/0.0						
2					rmt	OR
3					rmt	OR

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@Sirius# set routing-instances C5-vpls interface lt-0/0/0.1

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

This concludes Chapter Six. You can save the configuration on all devices.

Chapter Seven solutions: 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@Canopus# set interfaces ge-0/0/5.302 family mpls
```

- 2) Configure the P3-1 facing interface under the protocols mpls stanza.

```
[edit]
lab@Canopus# set protocols mpls interface ge-0/0/5.302
```

- 3) Configure **family inet-vpn unicast** on P3-1 EBGP session.

```
[edit]
```

```
lab@Canopus# set protocols bgp group P3-1 family inet-vpn unicast

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

Verify the BGP session table.

```
[edit]
lab@Canopus# run show bgp summary
Groups: 5 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp
State          Pending
inet.0
0              0
546            518            0            0
bgp.13vpn.0
22            21            0            0
0              0
inet6.0
48            48            0            0
0              0
bgp.12vpn.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.13vpn.0: 19/19/19/0
  inet6.0: 32/32/32/0
  bgp.13vpn-inet6.0: 10/10/10/0
  bgp.12vpn.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
  C5-12vpn.12vpn.0: 3/3/3/0
  C5-vpls.12vpn.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.13vpn.0: 2/3/2/0
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.13vpn.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@Canopus# set protocols bgp group P3-1 family inet unicast
```

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

```
[edit]
lab@Canopus# run show bgp summary
Groups: 5 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp
State          Pending
inet.0
0              0          660          598          0          0
bgp.l3vpn.0
0              0          22          21          0          0
inet6.0
0              0          48          48          0          0
bgp.l2vpn.0
0              0          6           6          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
  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 have to modify the policy to also allow **family inet-vpn** unicast prefixes.

```
lab@Canopus# show policy-options policy-statement P3-filter
term 1 {
    from {
        route-filter 0.0.0.0/0 prefix-length-range /8-/24;
    }
    then {
        local-preference 200;
        community set P3;
        accept;
    }
}
term 2 {
    then reject;
}
```

See above how the current import policy looks. You have to restrict both existing terms only for **family inet**. And for security reasons only accept /32 prefixes from the strictly defined AS path.

[edit]

```
lab@Canopus# set policy-options policy-statement P3-filter term 1
from family inet
```

[edit]

```
lab@Canopus# rename policy-options policy-statement P3-filter term
2 to term 3
```

[edit]

```
lab@Canopus# rename policy-options policy-statement P3-filter term
1 to term 2
```

Renaming the existing rules, so that you can create new term 1 and inserts it before existing terms.

[edit]

```
lab@Canopus# set policy-options policy-statement P3-filter term 1
from protocol bgp
```

[edit]

```
lab@Canopus# set policy-options policy-statement P3-filter term 1
from route-filter 0/0 prefix-length-range /32-/32
```

[edit]

```
lab@Canopus# set policy-options policy-statement P3-filter term 1
from as-path P3-local-routes
```

[edit]

```
lab@Canopus# set policy-options policy-statement P3-filter term 1
then accept
```

```
[edit]
```

```
lab@Canopus# insert policy-options policy-statement P3-filter term
1 before term 2
```

Insert the new term 1 before term 2 (old term 1).

```
[edit]
```

```
lab@Canopus# set policy-options as-path P3-local-routes
"2831679853 64600"
```

Let's verify the routes received from neighbor 192.168.0.6

```
[edit]
lab@Canopus# run show route received-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
```

Two prefixes are received and accepted from neighbor P3-1 in the **bgp.l3vpn.0** table.

The routes in **bgp.13vpn.0** match the default BGP export policy and are advertised to P3-1.

```
[edit]
lab@Canopus# run show route advertising-protocol bgp 192.168.0.6
table bgp.13vpn.0

bgp.13vpn.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
*                               Self                      2
I
  172.30.5.3:4:172.31.48.12/30
*                               Self                      3
I
  172.30.5.3:4:172.31.48.16/30
*                               Self                      7
I
  172.30.5.3:4:172.31.63.1/32
*                               Self                      5
I
  172.30.5.3:4:172.31.63.2/32
*                               Self                      1
I
  172.30.5.3:4:172.31.63.3/32
*                               Self                      2
I
--- (more) ---
```

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@Canopus# 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
--- (more) ---
```

Both 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 rReflector because no other PE requested these route targets.

Activating the advertise-default 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@Canopus# set protocols ibgp family route-target advertise-
default

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

```
[edit]
lab@Canopus# run show route advertising-protocol bgp 172.30.5.41
table bgp.l3vpn.0

bgp.l3vpn.0: 68 destinations, 68 routes (68 active, 0 holddown, 0
hidden)
  Prefix                               Nexthop                MED      Lclpref
AS path
172.17.47.2:200:172.31.78.0/24
*                                     Self                    100
2831679853 64600 I
  172.17.47.2:200:172.31.79.0/24
*                                     Self                    100
2831679853 64600 I
  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
--- (more) ---
```

After announcing the default Route-target route to the Route Reflector, the VPNv4 routes are sent.

6) Check the received route target.

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@Canopus# run show route received-protocol bgp 192.168.0.6
table bgp.l3vpn.0 extensive

bgp.l3vpn.0: 56 destinations, 56 routes (56 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.

7) 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 have to change the extended community on the routes in each directions. 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@Canopus# set policy-options policy-statement C2-vpn-target-
import term 1 from protocol bgp
```

```
[edit]
lab@Canopus# set policy-options policy-statement C2-vpn-target-
import term 1 from community CE2-remote
```

```
[edit]
lab@Canopus# set policy-options policy-statement C2-vpn-target-
import term 1 then community delete CE2-remote
```

Delete the existing community.

```
[edit]
lab@Canopus# set policy-options policy-statement C2-vpn-target-
import term 1 then community add CE2-spoke
```

Apply the new, normalized community.

```
[edit]
lab@Canopus# set policy-options policy-statement C2-vpn-target-
import term 1 then accept
```

```
[edit]
lab@Canopus# set policy-options policy-statement C2-vpn-target-
export term 1 from protocol bgp
```

```
[edit]
lab@Canopus# set policy-options policy-statement C2-vpn-target-
export term 1 from community CE2-hub
```

```
[edit]
lab@Canopus# set policy-options policy-statement C2-vpn-target-
export term 1 then community delete CE2-hub
```

```
[edit]
lab@Canopus# set policy-options policy-statement C2-vpn-target-
export term 1 then community add CE2-remote
```

```
[edit]
lab@Canopus# set policy-options policy-statement C2-vpn-target-
export term 1 then accept
```

Configure needed extended communities.

```
[edit]
lab@Canopus# set policy-options community CE2-remote members
target:43208:200
```

```
[edit]
lab@Canopus# set policy-options community CE2-spoke members
target:54591:201
```

```
[edit]
lab@Canopus# set policy-options community CE2-hub members
target:54591:200
```

Apply the policy on EBGp session to P3-1.

```
[edit]
lab@Canopus# set protocols bgp group P3-1 import C2-vpn-target-
import
```

```
[edit]
lab@Canopus# set protocols bgp group P3-1 export C2-vpn-target-
export
```

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

Verify the normalization of the route targets.

```
[edit]
lab@Canopus# 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
```

The routes received from P3-1 and sent to the Route Reflector are normalized with the spoke route target.

```
[edit]
lab@Canopus# run show route advertising-protocol bgp 192.168.0.6
table bgp.l3vpn.0

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
```

The remote route target is replaced on the advertised routes to P3-1.

8) 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@Canopus# run show route received-protocol bgp 192.168.0.6
table bgp.13vpn.0 extensive

bgp.13vpn.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 next-hop is changed the label will also be changed.

Verify the label advertised to the route reflector.

```
[edit]
lab@Canopus# 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@Canopus# run show route table mpls.0 label 314208

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 have to 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@Sirius# set protocols bgp group ibgp family inet labeled-
unicast rib inet.3
```

```
[edit]
lab@Sirius# 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]
root@rr# set protocols bgp group cluster-1 family inet labeled-
unicast rib inet.3
```

```
[edit]
root@rr# set protocols bgp group cluster-2 family inet labeled-
unicast rib inet.3
```

```
[edit]
root@rr# commit
[edit protocols bgp group cluster-1 cluster]
commit complete
```

The Route Reflector is also configured for the BGP family inet labeled unicast.

On R3 you have to 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@Canopus# set protocol bgp group ibgp family inet labeled-
unicast rib inet.3
```

```
[edit]
```

```
lab@Canopus# set protocol bgp group P3-1 family inet labeled-unicast rib inet.3
```

```
[edit]
```

```
lab@Canopus# commit
commit complete
```

You need to change the P3-filter policy to accept the labeled BGP routes from P3-1.

```
[edit]
```

```
lab@Canopus# set policy-options policy-statement P3-filter term 3 from family inet
```

```
[edit]
```

```
lab@Canopus# commit
commit complete
```

Verify the received labeled BGP route, which is the the remote PE Loopback IP address from P3-1.

```
[edit]
lab@Canopus# 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
                   > to 192.168.0.6 via ge-0/0/5.302
```

2) Modify the P3-1 export policy on R3.

You need to 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@Canopus# 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 Canopus-to-Vega
                   to 172.30.0.26 via ge-0/0/4.136, label-
switched-path Canopus-to-Vega
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 Canopus-to-Vega
                   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 Canopus-to-Vega
                   to 172.30.0.26 via ge-0/0/4.136, label-
switched-path Canopus-to-Vega
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 Canopus-to-Vega
                   to 172.30.0.26 via ge-0/0/4.136, label-
switched-path Canopus-to-Vega
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 Canopus-to-Procyon-1
                   > to 172.30.0.26 via ge-0/0/4.136, label-
switched-path Canopus-to-Procyon-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 have to modify the policy to include the Loopback IP address.

```
[edit]
lab@Canopus# set policy-options policy-statement rr-loopback-to-
inet3 term 1 from route-filter 172.30.5.3/32 exact
```

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

```
[edit]
lab@Canopus# 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@Canopus# set policy-options policy-statement P3-export term 2
from rib inet.3
```

```
[edit]
lab@Canopus# set policy-options policy-statement P3-export term 2
route-filter 172.30.5.0/24 prefix-length-range /32-/32
```

```
[edit]
lab@Canopus# set policy-options policy-statement P3-export term 2
then accept
```

```
[edit]
lab@Canopus# 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 I2vpn signalled routes will be reflected to the local PE routers.

```
[edit]
root@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]
root@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]
root@route-reflector# set protocols bgp group P3-remote-pe no-
nexthop-change
```

By default BGP changes the next-hop attribute on all EBGp sessions. Changing the next-hop in this particular 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]
root@route-reflector# set protocols bgp group P3-remote-pe local-
address 172.30.5.41
```

Specify the source address the BGP TCP session.

```
[edit]
root@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]
root@route-reflector# set protocols bgp group P3-remote-pe peer-as
23456
```

Because the peer AS number is 4-byte, and the Router Reflectors JUNOS version 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]
root@route-reflector# set protocols bgp group P3-remote-pe
neighbor 172.17.47.3
```

```
[edit]
root@route-reflector# commit
[edit protocols bgp group cluster-1 cluster]
```

commit complete

Verify the BGP session table for the remote EBGp peer.

```
[edit]
root@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.13vpn.0     105         105          0          0
0              0
inet6.0         64          48          0          0
0              0
bgp.12vpn.0     9           9           0          0
0              0
bgp.mvpn.0      5           5           0          0
0              0
inet.3          1           1           0          0
0              0
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.12vpn.0: 1/1/0
--- (more) ---
```

Verify the single prefix received from the remote neighbor in the **bgp.12vpn.0** table.

```
[edit]
root@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)
  Route Distinguisher: 172.17.47.3:500
  Label-base: 262145, range: 8
  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:43208:500 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
```

The route target community target:43208:500 attached to the route is different then the route target used for customer C5 VPLS service. You have to normalize the route targets attached to the prefixes in both directions.

```
[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-import term 1 from protocol bgp
```

```
[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-import term 1 from community CE5-remote
```

```
[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-import term 1 then community delete CE5-remote
```

```
[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-import term 1 then community add CE5
```

```
[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-import term 1 then accept
```

```
[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-export term 1 from protocol bgp
```

```
[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-export term 1 from community CE5
```

```
[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-export term 1 then community delete CE5
```

```
[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-export term 1 then community add CE5-remote
```

```

[edit]
root@route-reflector# set policy-options policy-statement C5-vpn-
target-export term 1 then accept

[edit]
root@route-reflector# set policy-options community CE5-remote
members target:43208:500

[edit]
root@route-reflector# set policy-options community CE5 members
target:54591:501

[edit]
root@route-reflector# set protocols bgp group P3-remote-pe import
C5-vpn-target-import

[edit]
root@route-reflector# set protocols bgp group P3-remote-pe export
C5-vpn-target-export

[edit]
root@route-reflector# commit
[edit protocols bgp group cluster-1 cluster]

commit complete

```

The policies delete the extended route target community and set the correct one between the remote and local AS.

Verify the operation of the Inter-provider option C VPLS service.

```
[edit]
lab@Sirius# run show route table bgp.12vpn.0 extensive

bgp.12vpn.0: 7 destinations, 7 routes (7 active, 0 holddown, 0
hidden)
  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
              Address: 0x155cfcc
              Next-hop reference count: 5
              Source: 172.30.5.41
              Protocol next hop: 172.17.47.3
              Indirect next hop: 2 no-forward
              State: <Active Int Ext>
              Local AS: 54591 Peer AS: 54591
              Age: 27          Metric2: 1
              Task: BGP_54591.172.30.5.41+1181
              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:, mtu: 0, site preference: 25600
              Import Accepted
              Label-base: 262145, range: 8
              Localpref: 100
              Router ID: 172.30.5.41
              Secondary Tables: C5-vpls.12vpn.0
              Indirect next hops: 1
                Protocol next hop: 172.17.47.3 Metric: 10
                Indirect next hop: 2 no-forward
                Indirect path forwarding next hops: 1
                  Next hop type: Router
                  Next hop: 172.30.0.14 via ge-
0/0/4.123
                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 318064
                  Indirect nexthop: 1856188 -
                  Indirect path forwarding nexthops:
1
                    Nexthop: 172.30.0.14 via
ge-0/0/4.123
                    172.30.5.3/32 Originating RIB:
inet.3
                      Metric: 10
Node path count: 1
                      Forwarding nexthops: 1
```

ge-0/0/4.123 Nexthop: 172.30.0.14 via

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@Sirius# 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.

Verify the VPLS connection table.

```
[edit]
lab@Sirius# run show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS                    WE -- interface and instance
EM -- encapsulation mismatch     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
```



```

6      Description: Intf - vpls C5-vpls local site 4 remote site
7      rmt      Up      Nov  4 20:28:18 2013
1      [redacted]
      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.

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@Canopus# 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@Canopus# 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
    AS path: [54591] 2831679853 I
```

This label is used by other the PE routers to reaching the remote PE router.

```
[edit]
lab@Canopus# 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@Canopus# run show vpls connections
Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not
CCC/TCC/VPLS                    WE -- interface and instance
EM -- encapsulation mismatch     encaps not same
VC-Dn -- Virtual circuit down    NP -- interface hardware not
present                          -> -- only outbound connection is
CM -- control-word mismatch      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                         SC -- local and remote site ID
RD -- remote site signaled down 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-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! You can save the configuration on all devices.

Chapter Eight solutionst: 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.

Four predefined classes are available on each JUNOS devices.

```
[edit]
lab@Sun# run show class-of-service forwarding-class
Forwarding class      ID   Queue Policing priority  SPU
priority
  best-effort          0    0          normal        low
  expedited-forwarding 1    1          normal        low
  assured-forwarding   2    2          normal        low
  network-control      3    3          normal        low
```

You have to change the name of some the classes based on Table 25.

```
[edit]
lab@Sun# set class-of-service forwarding-classes queue 0 best-
effort

[edit]
lab@Sun# set class-of-service forwarding-classes queue 1 vpn

[edit]
lab@Sun# set class-of-service forwarding-classes queue 2 vpn-
priority

[edit]
lab@Sun# set class-of-service forwarding-classes queue 3 nc

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

Again, verify the forwarding classes configuration.

```
[edit]
lab@Sun# run show class-of-service forwarding-class
Forwarding class      ID   Queue Policing priority  SPU
priority
  best-effort          0    0          normal        low
  vpn                  1    1          normal        low
  vpn-priority         2    2          normal        low
  nc                   3    3          normal        low
```

- 2) Configure drop profiles.

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.

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.

```
[edit]
lab@Sun# set class-of-service drop-profiles low-drop interpolate
fill-level [ 25 50 75 ]
```

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@Sun# set class-of-service drop-profiles low-drop interpolate
drop-probability [ 5 15 40 ]
```

```
[edit]
lab@Sun# set class-of-service drop-profiles high-drop interpolate
fill-level [ 25 50 75 ]
```

```
[edit]
lab@Sun# set class-of-service drop-profiles high-drop interpolate
drop-probability [ 10 30 65 ]
```

3) Configure the schedulers

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 be 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.

```
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@Sun# set class-of-services 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@Sun# set class-of-services schedulers be-sc-q0 buffer-size
remainder
```

It can use the remaining buffer of the queue.

```
[edit]
lab@Sun# 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@Sun# set class-of-services 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@Sun# set class-of-services 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 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@Sun# set class-of-service schedulers vpn-sc-q1 buffer-size
percent 20
```

20 percent of the buffer is dedicated for that queue.

```
[edit]
lab@Sun# 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@Sun# set class-of-services 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@Sun# set class-of-services schedulers vpn-sc-q1 drop-profile-
map loss-priority high protocol any drop-profile high-drop
```

```
[edit]
lab@Sun# set class-of-services schedulers vpn-prio-sc-q2 transmit-
rate percent 10
```

```
[edit]
lab@Sun# set class-of-services 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** knob, the buffer has a strict upper limit and cannot expand.

```
[edit]
```



```
lab@Sun# set class-of-services schedulers vpn-prio-sc-q2 priority
medium-high
```

```
[edit]
```

```
lab@Sun# set class-of-service schedulers nc-sc-q3 transmit-rate
percent 5
```

```
[edit]
```

```
lab@Sun# set class-of-service schedulers nc-sc-q3 buffer-size
percent 5
```

```
[edit]
```

```
lab@Sun# set class-of-service schedulers nc-sc-q3 priority high
```

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@Sun# set class-of-services scheduler-maps core-interfaces
forwarding-class best-effort scheduler be-sc-q0
```

```
[edit]
```

```
lab@Sun# set class-of-services scheduler-maps core-interfaces
forwarding-class nc scheduler nc-sc-q3
```

```
[edit]
```

```
lab@Sun# set class-of-services scheduler-maps core-interfaces
forwarding-class vpn scheduler vpn-sc-q1
```

```
[edit]
```

```
lab@Sun# set class-of-services scheduler-maps core-interfaces
forwarding-class vpn-priority scheduler vpn-prio-sc-q2
```

```
[edit]
```

```
lab@Sun# commit
commit complete
```

Verify the configured scheduler-map

```
[edit]
lab@Sun# 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
      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@Sun# set class-of-services interfaces ge-0/0/4 unit *
scheduler-map core-interfaces

[edit]
lab@Sun# set class-of-services 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.

```
[edit]
lab@Sun# 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@Sun# set interfaces ge-0/0/4 per-unit-scheduler

[edit]
lab@Sun# set interfaces ae0 per-unit-scheduler

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

After enabling **per-unit-scheduler** commit is successful.

Verify the interface queues configured in the scheduler-map.

```
[edit]
lab@Sun# 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
    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.

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.

```
[edit]
lab@Canopus# set firewall family inet filter C2-classifier term 1
from dscp be
```

```
[edit]
lab@Canopus# set firewall family inet filter r C2-classifier term
1 then forwarding-class vpn
```

```
[edit]
lab@Canopus# set firewall family inet filter C2-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@Canopus# set firewall family inet filter C2-classifier term 2
then forwarding-class vpn-priority
```

```
[edit]
lab@Canopus# set firewall family inet filter C2-classifier term 2
then accept
```

Rest of the traffic is assigned to the vpn-priority class.

```
[edit]
lab@Canopus# set interface ge-0/0/5.317 family inet filter input
C2-classifier
```

- 2) Configure the core point aliases on all routers.

Code-point-aliases assign friendly names to class-of-service code-points.

```
[edit]
lab@Canopus# 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@Canopus# set class-of-service code-point-aliases dscp vpn-high
001100
```

```
[edit]
```

```

lab@Canopus# set class-of-service code-point-aliases dscp vpn-
priority 101110

[edit]
lab@Canopus# set class-of-service code-point-aliases dscp be
000000

[edit]
lab@Canopus# set class-of-service code-point-aliases dscp nc
110000

[edit]
lab@Canopus# set class-of-service code-point-aliases exp vpn-low
010

[edit]
lab@Canopus# set class-of-service code-point-aliases exp vpn-high
011

[edit]
lab@Canopus# set class-of-service code-point-aliases exp vpn-
priority 101

[edit]
lab@Canopus# set class-of-service code-point-aliases exp be 000

[edit]
lab@Canopus# 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.

```

[edit]
lab@Canopus# set class-of-services forwarding-policy next-hop-map
cbf-map forwarding-class vpn lsp-next-hop Canopus-to-Procyon-2

[edit]
lab@Canopus# set class-of-services forwarding-policy next-hop-map
cbf-map forwarding-class vpn-priority lsp-next-hop Canopus-to-
Procyon-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**.

```

[edit]
lab@Canopus# set policy-options policy-statement cbf-map term 1
from route-filter fd18:cccc:dddd:5:0::/77 longer

```

```
[edit]
lab@Canopus# 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@Canopus# set routing-options forwarding-table export cbf-map
```

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

Verify the next-hop installed in the forwarding table.

```
[edit]
lab@Canopus# 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:5:0::/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
                        idx:2      172.30.0.22  Push 349008, Push
309712, Push 310976(top)  780      2 ge-0/0/4.134
```


324720 (top)	idx:xx	172.30.0.26	Push 349008, Push
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@Canopus# 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@Canopus# set firewall policer vpn-policer if-exceeding burst-
size-limit 62k
```

burst-size is 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:

$$\text{burst-size} = \text{interface bandwidth} \times \text{allowable time for bust traffic} / 8$$

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@Canopus# set firewall policer vpn-policer then loss-priority
high
```

The soft policers change the loss priority of the exceeding traffic.

```
[edit]
lab@Canopus# set firewall policer vpn-priority-policer if-
exceeding bandwidth-limit 60m
```

```
[edit]
lab@Canopus# set firewall policer vpn-priority-policer if-
exceeding burst-size-limit 62k
```

```
[edit]
lab@Canopus# set firewall policer vpn-priority-policer then
discard
```

The hard policers drop the exceeding traffic.

```
[edit]
```

```
lab@Canopus# set firewall family any filter vpn-filter term 1 then
policer vpn-policer
```

```
[edit]
```

```
lab@Canopus# set firewall family any filter vpn-filter term 1 then
accept
```

```
[edit]
```

```
lab@Canopus# set firewall family any filter vpn-priority-filter
term 1 then policer vpn-priority-policer
```

```
[edit]
```

```
lab@Canopus# 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.

```
[edit]
```

```
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-1 policing filter vpn-priority-filter
```

```
[edit]
```

```
lab@Canopus# set protocols mpls label-switched-path Canopus-to-
Procyon-2 policing filter vpn-filter
```

```
[edit]
```

```
lab@Canopus# commit
commit complete
```

8) Configure the behavior aggregate classifiers on all routers.

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 have to configure the DSCP and EXP BA classifiers on all routers.

```
[edit]
```

```
lab@Sun# set class-of-service classifiers dscp dscp-classifier
forwarding-class best-effort loss-priority low code-points be
```

Classifier **dscp-classifier** from type **dscp** classifies code-points **be** to forwarding-class **best-effort**.

```
[edit]
```

```
lab@Sun# set class-of-service classifiers dscp dscp-classifier
forwarding-class vpn loss-priority low code-points vpn-low
```

```
[edit]
```

```

lab@Sun# set class-of-service classifiers dscp dscp-classifier
forwarding-class vpn loss-priority high code-points vpn-high

[edit]
lab@Sun# set class-of-service classifiers dscp dscp-classifier
forwarding-class vpn-priority loss-priority low code-points vpn-
priority

[edit]
lab@Sun# set class-of-service classifiers dscp dscp-classifier
forwarding-class nc loss-priority low code-points nc

[edit]
lab@Sun# set class-of-service classifiers exp mpls-classifier
forwarding-class best-effort loss-priority low code-points be

[edit]
lab@Sun# set class-of-service classifiers exp mpls-classifier
forwarding-class vpn loss-priority low code-points vpn-low

[edit]
lab@Sun# set class-of-service classifiers exp mpls-classifier
forwarding-class vpn loss-priority high code-points vpn-high

[edit]
lab@Sun# set class-of-service classifiers exp mpls-classifier
forwarding-class vpn-priority loss-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@Sun# set class-of-service interfaces ge-0/0/4 unit *
classifiers dscp dscp-classifier

[edit]
lab@Sun# set class-of-service interfaces ge-0/0/4 unit *
classifiers exp mpls-classifier

[edit]
lab@Sun# set class-of-service interfaces ae0 unit * classifiers
dscp dscp-classifier

[edit]
lab@Sun# set class-of-service interfaces ae0 unit * classifiers
exp mpls-classifier

[edit]
lab@Sun# 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.

```
[edit]
lab@Sun# set class-of-service rewrite-rules dscp dscp-rewriter
forwarding-class best-effort loss-priority low code-point be
```

```
[edit]
lab@Sun# set class-of-service rewrite-rules dscp dscp-rewriter
forwarding-class vpn loss-priority low code-point vpn-low
```

```
[edit]
lab@Sun# set class-of-service rewrite-rules dscp dscp-rewriter
forwarding-class vpn loss-priority high code-point vpn-high
```

```
[edit]
lab@Sun# set class-of-service rewrite-rules dscp dscp-rewriter
forwarding-class vpn-priority loss-priority low code-point vpn-
priority
```

```
[edit]
lab@Sun# set class-of-service rewrite-rules dscp dscp-rewriter
forwarding-class nc loss-priority low code-point nc
```

```
[edit]
lab@Sun# set class-of-service rewrite-rules exp mpls-rewriter
forwarding-class best-effort loss-priority low code-point be
```

```
[edit]
lab@Sun# set class-of-service rewrite-rules exp mpls-rewriter
forwarding-class vpn loss-priority low code-point vpn-low
```

```
[edit]
lab@Sun# set class-of-service rewrite-rules exp mpls-rewriter
forwarding-class vpn loss-priority high code-point vpn-high
```

```
[edit]
```

```
lab@Sun# 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@Sun# set class-of-service interfaces ge-0/0/4 unit * rewrite-
rules dscp dscp-rewriter
```

```
[edit]
lab@Sun# set class-of-service interfaces ge-0/0/4 unit * rewrite-
rules exp mpls-rewriter protocol mpls-inet-both
```

```
[edit]
lab@Sun# set class-of-service interfaces ae0 unit * rewrite-rules
dscp dscp-rewriter
```

```
[edit]
lab@Sun# set class-of-service interfaces ae0 unit * rewrite-rules
exp mpls-rewriter protocol mpls-inet-both
```

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

Verify the interface class-of-services configuration.

```
[edit]
lab@Sun# 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
  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! You can save the configuration on all devices.

Chapter Nine solutions: Full Day Lab Challenge I

This last chapter will simulate a full day 8 hour lab challenge. You have to configure everything from scratch and build the network according to the requirements. You will be tested on the technologies practiced in the previous eight chapters. All tasks should be completed in 8 hours.

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 providedn for each device in Table 31 and 32.

- 1) Configure the hostnames on the routers according to Table 31.

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

- 2) Configure the OoB management interface according to Table 32.

```
[edit]
lab@Sun# set interfaces ge-0/0/0.0 description "OoB management"
```

```
[edit]
lab@Sun# set interfaces ge-0/0/0.0 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@Sun# set system services ftp
```

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

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

Configure a back-router ensures connectivity to the management subnet when RDP is not running.

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

- 5) Configure server S1 as the DNS server.

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

- 6) Set the time zone to Europe/Amsterdam on all your devices.

```
[edit]
lab@Sun# set system time-zone Europe/Amsterdam
```

- 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@Sun# set system ntp server 10.10.1.100
```

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

- 8) Configure the authentication method that first tries 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@Sun# set system authentication-order [ radius password ]
```

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

```
[edit]
lab@Sun# 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@Sun# set system login user lab class super-user authentication
plain-text-password
```

- 10) Configure additional users on all devices according to Table 33.

a 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 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@Sun# set system login class op-plus permissions [ clear
maintenance network reset snmp-control trace view ]
```

```
[edit]
```

```
lab@Sun# set system login class op-plus deny-commands "start
shell"
```

You have to configure a new class based on *thesuper-user* class. However the users of this class are not allowed to execute "*clear*", "*configure*" and "*edit*" command.

```
[edit]
```

```
lab@Sun# set system login class su-minus permissions all deny-
commands "(clear) | (configure) | (edit)"
```

```
[edit]
```

```
lab@Sun# set system login user noc class op-plus authentication
plain-text-password
```

```
[edit]
```

```
lab@Sun# set system login user tac class su-minus authentication
plain-text-password
```

11) Configure syslog settings on all your devices as described in Table 34.

```
[edit]
```

```
lab@Sun# set system syslog user noc any warning
```

```
[edit]
```

```
lab@Sun# set system syslog user lab any emergency
```

```
[edit]
```

```
lab@Sun# set system syslog host 10.10.1.100 change-log any
```

```
[edit]
```

```
lab@Sun# set system syslog file jncie-sp-messages any info
```

```
[edit]
```

```
lab@Sun# 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
```

```
[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 user name *lab* and password *_lab123*.

```
[edit]
lab@Sun# set system archival configuration transfer-interval 1440
```

```
[edit]
lab@Sun# 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.

```
[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 triggered 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

```

Task 2: Building the network

The second task requires the configuration of the interfaces on all routers according to Table 35 and 36. As a result the topology on Figure 16 will be built.

The aggregated interfaces should use LACP. All interfaces should have descriptions.

```

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

```

You have to configure how many **ae** interface you will use under **chassis** stanza.

```
[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
```

You have to add physical interfaces to the Aggregated Ethernet bundle interfaces.

```
[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
```

```
[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
```

Enabling LACP on Aggregated Ethernet interface.

```
[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

[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

```

After Configure all routers according to the given Tables you should have the topology shown on Figure 16.

Task 3: IGP Configuration

You have to configure both OSPFv2 and OSPFv3 according to Table 37.

- 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.

```

[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface ae0 interface-
type p2p

```

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 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
```

Enable OSPFv2 to the Route Reflector router.

```
[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
```

```
[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 a manual Router ID based on the Loopback IP address.

```
[edit]
lab@R1# set protocols ospf reference-bandwidth 10g
```

```
[edit]
lab@R1# set protocols ospf3 reference-bandwidth 10g
```

Configure the reference bandwidth to 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
```

```
[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
```

```
[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
```

```
[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

[edit]
lab@R4# set protocols ospf area 0.0.0.1 interface lo0.0

[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

[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]
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]
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

[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]
root@route-reflector# set protocols ospf area 0.0.0.0 interface
ge-0/0/1.206 interface-type p2p

[edit]
root@route-reflector# set protocols ospf area 0.0.0.0 interface
ge-0/0/1.207 interface-type p2p

[edit]
root@route-reflector# set protocols ospf area 0.0.0.0 interface
lo0.0

[edit]
root@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
ae0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0
ge-0/0/4.117	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0
ge-0/0/4.118	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0
ge-0/0/4.206	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0
lo0.0	DR	0.0.0.0	172.30.5.1	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
172.30.5.7	ge-0/0/4.117	Full	128	32
172.30.5.8	ge-0/0/4.118	Full	128	32

All neighbors for OSPFv3 are UP.

```
[edit]
lab@R1# run show ospf3 interface
```

Interface	State	Area	DR ID	BDR ID
ae0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0
ge-0/0/4.117	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0
ge-0/0/4.118	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0

```
100.0          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
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
Type      ID      Adv Rtr      Seq      Age
Checksum 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 1	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0
ge-0/0/4.123 1	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0
ge-0/0/4.126 1	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0
lo0.0 0	PtToPt	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
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
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
Neighbor-address	fe80::fac0:100:87dd:384			

```
[edit]
lab@R3# run show ospf3 interface
Interface          State   Area      DR ID      BDR ID
Nbtrs
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				
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
0xdc8b	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
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			
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
Nbrrs
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:fece: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
Nbrrs
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.

```
[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
0xdc8    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.

```
[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]
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
0xdcbb  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
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
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
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				
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
0xdcbb	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
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			
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
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
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]
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

```

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	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
```

```
[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]
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
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 presens.

```
[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
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::fac0:1fff:fedc:2c80 via ae0.0
fd17:f0f4:f691:5::2/128
                   *[OSPF3/10] 00:00:10, metric 15
                   > to fe80::fac0:1fff: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::fac0:1fff: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::fac0:1fff:fedc:2c80 via ae0.0
fd17:f0f4:f691:5::7/128
                   *[OSPF3/10] 00:00:10, metric 15
                   > to fe80::fac0:1fff: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 stanza, not under **nsa**.

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]
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. 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]
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
                  > 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

In Neighbor Met      Local  Source      Destination      Send  Receive
State Address      Address          Mode          Mode
-----
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
                  > 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.

```

[edit]
lab@R3# set protocols ospf area 0.0.0.1 nssa default-lsa default-
metric 10

```

```

[edit]
lab@R3# commit

```



```
commit complete
```

```
[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 a LSA Type 7 default on both ABRs.

Verify if the default route is present on R5.

```
[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
--- (more) ---
```

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
--- (more) ---
```

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.

```
[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 policy-statement ospf-
to-rip term 1 then tag 1234

[edit]
lab@R5# set policy-options policy-statement 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

[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
```

R6 is advertising the statically configured aggregate route.

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

[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
```

```
[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 0.0.0.0						
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.

RIP route in R6 are known by RIP and by OSPF.

```
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 have to 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 have to implement an other solution for advertising the summary route because the active routes in the routing table are known by RIP. The RIP routes will not be summarized by adding the area-range under area nssa. However you still have to 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.

```
[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 configred 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 protocol ospf stanza 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				
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.

8) Verify loop free connectivity in the network.

Perform 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

Task 4: BGP Configuration

In this task you have to configure External BGP sessions with different peers and Customers. First you have to build the Internal BGP network. According to the requirements routing policy should be applied to control the exchange of the routes.

First you have to configure IBGP and ensure your configuration follows the requirements. You have to 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.

```
[edit]
lab@R1# set routing-options autonomous-system 54591
```

The Autonomous system of the network is 54591

```
[edit]
lab@R1# set protocols bgp group ibgp type internal
```

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 log-updown
```

Configure neighbor changes to be logged by syslog configured under the system syslog stanza.

```
[edit]
lab@R1# set protocols bgp group ibgp local-address 172.30.5.1
```

IBGP sessions are sourced from the Loopback IP address.

```
[edit]
lab@R1# set protocols bgp group ibgp neighbor 172.30.5.41
```

The RR loopback address is used in the neighbor statement.

```
[edit]
lab@R1# set protocols bgp group ibgp authentication-key workbook
```

Configure BGP authentication for the session with the Route Reflector.

```
[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

[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

[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
```

```
[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.

```
[edit]
root@route-reflector# set routing-options autonomous-system 54591
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 type
internal
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 log-updown
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 local-
address 172.30.5.41
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 cluster
0.0.0.1
```

Configure the BGP groups with cluster-ids to reflect the routes between the internal BGP neighbors.

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 neighbor
172.30.5.1
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 neighbor
172.30.5.6
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 neighbor
172.30.5.7
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-1 neighbor
172.30.5.8
```

Configure routers R1, R6, R7 and R8 as neighbors to group cluster-1.

```
[edit]
root@route-reflector# set protocols bgp group cluster-2 type
internal
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-2 log-updown
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-2 local-
address 172.30.5.41
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-2 cluster
0.0.0.2
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-2 neighbor
172.30.5.2
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-2 neighbor
172.30.5.3
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-2 neighbor
172.30.5.4
```

```
[edit]
root@route-reflector# set protocols bgp group cluster-2 neighbor
172.30.5.5
```

Configure routers R2, R3, R4, R5 as neighbors in group cluster-2.

Verify the iBGP operation between the routers and the Route Reflector.

```
[edit]
root@route-reflector# run show bgp summary
Groups: 3 Peers: 9 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/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          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 39.

According to Table 39 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 39.

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.

```
[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
0              0          564          564          0          0
inet6.0
0              0          16          16          0          0
Peer           AS           InPkt      OutPkt      OutQ      Flaps
Last Up/Dwn State|#Active/Received/Accepted/Damped...
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

```

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 do esnot process IPv4-mapped addresses because of security measures. To enable the processing of these packets you have to explicitly enable it under the system stanza 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]
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
0              0          16          16          0          0
inet6.0
0              0          16          16          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				
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
*                               ::ffff:192.168.0.10
64514 I
    fd18:cccc:dddd:2:9::/80
*                               ::ffff:192.168.0.10
64514 I
--- (more) ---

```

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
  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.

```

[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 multi-hop knob 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.

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

[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

[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

[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
0              680        669          0          0
inet6.0
0              16         16           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 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

[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
```

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.

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.

```
[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
```

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
--- (more) ---
```

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
```

```
[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 have to 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 r1
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
```

```
[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 router 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 have to apply them on the neighbors in the right direction.

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
```

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 ]
```

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 neighborhood with the customer C3-1 router.

The default community list is configured on R2 together with policies. Also the remove-private knob is used.

```
[edit]
lab@R2# set protocols bgp group C3-1 remove-private

[edit]
lab@R2# set policy-options community customer members 54591:200

[edit]
lab@R2# set policy-options community peer members 54591:100

[edit]
lab@R2# set policy-options community wildcard members *:*

[edit]
lab@R2# set policy-options community rtbh members 6451.:666

[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-option policy-statement delete-communities term
1 from protocol bgp

[edit]
lab@R2# set policy-option policy-statement delete-communities term
1 then community delete wildcard

[edit]
lab@R2# set policy-option policy-statement ebgp-import-filter term
1 from route-filter 0/0 upto /7

[edit]
lab@R2# set policy-option policy-statement ebgp-import-filter term
1 then reject

[edit]
lab@R2# set policy-option policy-statement ebgp-import-filter term
2 from route-filter 0/0 prefix-length-range /25-/32

[edit]
lab@R2# set policy-option policy-statement ebgp-import-filter term
2 then reject

[edit]
lab@R2# set policy-option policy-statement ebgp-import-filter term
3 from route-filter 0/0 through 0/32

[edit]
lab@R2# set policy-option 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

[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.

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 have to 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 have to 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 and communities on R3.

```
[edit]
```

```
lab@R3# set policy-options community customer members 54591:200
```

```
[edit]
```

```
lab@R3# set policy-options community peer members 54591:100
```

```
[edit]
```

```
lab@R3# set policy-options community wildcard members *:*
```

```
[edit]
```

```
lab@R3# set policy-options community rtbh members 6451.:666
```

```
[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
```

```
[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.

```
[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 C3-1 import [ ebgp-import-filter
customer-routes customer-preferred ]

[edit]
lab@R3# set protocols bgp group C3-1 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 community customer members 54591:200

[edit]
lab@R4# set policy-options community peer members 54591:100

[edit]
lab@R4# set policy-options community wildcard members *:*

[edit]
lab@R4# set policy-options community rtbh members 6451.:666

[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.

```
[edit]
lab@R5# set protocols bgp group C1-1 remove-private

[edit]
lab@R5# set policy-options community customer members 54591:200

[edit]
lab@R5# set policy-options community peer members 54591:100

[edit]
lab@R5# set policy-options community wildcard members *:*

[edit]
lab@R5# set policy-options community rtbh members 6451.:666

[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
```

```
[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

[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 community list.

```
[edit]
lab@R6# set policy-options community customer members 54591:200

[edit]
lab@R6# set policy-options community peer members 54591:100

[edit]
lab@R6# set policy-options community wildcard members *:*

[edit]
lab@R6# set policy-options community rtbh members 6451.:666

[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 community customer members 54591:200

[edit]
lab@R7# set policy-options community peer members 54591:100

[edit]
lab@R7# set policy-options community wildcard members *:*

[edit]
lab@R7# set policy-options community rtbh members 6451.:666

[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

[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 Task.

```
[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 customer members 54591:200

[edit]
lab@R8# set policy-options community peer members 54591:100

[edit]
lab@R8# set policy-options community wildcard members *:*

[edit]
lab@R8# set policy-options community rtbh members 6451.:666

[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

[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 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 have to 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.

```
[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 ]

```

```
[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 ]

```

- 4) Verify the operation of the routing policies according to the requirements.

You have to 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      Lc1pref
AS path
* 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
--- (more) ---
```

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
* 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
```

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
```

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
```

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

--- (more) ---
```

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
```

```
[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

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)
  Prefix                               Nexthop                          MED      Lc1pref
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

```

```

[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
* 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

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] )"
* 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 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 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 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)
  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 ?
--- (more) ---
```

```
[edit]
lab@R8# run show 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
* 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
```

[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 40. Configure LDP on the Route Reflector. Configure MD5 authentication for all LDP sessions. Configure OSPF to track the status of LDP on the interfaces and Configure the LDP routes to use the IGP metric.

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.

family mpls should be configured on all interfaces forwarding labeled traffic. Additionally the interfaces should be added to **protocols mpls** stanza.

```
[edit]
lab@R1# set protocols mpls interface ae0.0

[edit]
lab@R1# set protocols mpls interface ge-0/0/4.117

[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 Configure 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 a 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
```

Verify LDP operation.

```
[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
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.

b. Configure R2.

The configuration on R2 is identical as on R1.

```
[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

[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 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
```

```
[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

c. Configure R3.

On R3 both interfaces to R4 and R5 must be enabled for LDP.

```
[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]
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

```

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

d. Configure router R4.

```

[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]
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

```

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

```

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

```

NOTE: The output is displayed after LDP is enabled on R6 also.

e. Configure router R5.

```

[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

[edit]
lab@R5# commit
commit complete

```

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
  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
```

NOTE: The output is displayed after LDP is enabled on R6 also.

f. Configure router R6.

```
[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

[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 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
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 R3's 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
```

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. A 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
```

g. Configure the Route Reflector.

```
[edit]
root@route-reflector# set interfaces ge-0/0/1.206 family mpls

[edit]
root@route-reflector# set interfaces ge-0/0/1.207 family mpls

[edit]
root@route-reflector# set protocols mpls interface ge-0/0/1.206

[edit]
root@route-reflector# set protocols mpls interface ge-0/0/1.207

[edit]
root@route-reflector# set protocols mpls track-igp-metric

[edit]
root@route-reflector# set protocols mpls session 172.30.5.1
authentication-key workbook

[edit]
root@route-reflector# set protocols mpls session 172.30.5.2
authentication-key workbook

[edit]
root@route-reflector# commit
commit complete
```

Verify the operation of LDP on the Route Reflector.

```
[edit]
root@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]
root@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

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
```

- 2) Configure the RSVP interfaces and coloring as described in Table 41. Enable RSVP message aggregation and path MTU discovery. Configure MD5 authentication for all RSVP sessions.

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.

family mpls should be configured on all interfaces which need to support MPLS. Additionally the interfaces should be added to **protocols mpls** stanza. 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 41.

```
[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]
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>

```
[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 a 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
  Protection: Off
```

b. Configure RSVP on R2.

```
[edit]
lab@R2# set protocols mpls path-mtu rsvp mtu-signaling

[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
```

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	ption	BW	BW
mark					
ae0.0	Up	0	100%	2Gbps	0bps
ge-0/0/4.123	Up	0	120%	1000Mbps	0bps
ge-0/0/4.126	Up	0	120%	1000Mbps	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]
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

[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

```

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` stanza.

```
[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

[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
Interface	State	resv	ption	BW	BW
Highwater					
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 rsvp neighbor
RSVP neighbor: 0 learned

[edit]
lab@R7# run show mpls interface
Interface      State      Administrative groups (x: extended)
ge-0/0/4.117   Up         blue
ge-0/0/4.137   Up         green
ge-0/0/4.167   Up         purple

```

f. Configure RSVP on R8.

```

[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

[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
    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

```

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]
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
```

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 have to configure two LSP from R2 to R1 and R8 using a different physical paths. 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
```

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
```

auto-bandwidth 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.

```
[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] set protocols mpls label-switched-path R2-to-R6 bandwidth
100m
lab@R2#

[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

[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
To          From          State Rt P      ActivePath
LSPname
172.30.5.1  172.30.5.2  Up    0 *    path-1
R2-to-R1-second
172.30.5.1  172.30.5.2  Up    0 *    primary-path
R2-to-R1-first
172.30.5.3  172.30.5.2  Up    0 *    primary-path
R2-to-R3
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 *    path-1
R2-to-R8-second
172.30.5.8  172.30.5.2  Up    0 *    primary-path
R2-to-R8-first
Total 7 displayed, Up 7, Down 0

Egress LSP: 5 sessions
To          From          State Rt Style Labelin Labelout
LSPname
172.30.5.2  172.30.5.8  Up    0  1 FF      3      -
R8-to-R2
172.30.5.2  172.30.5.3  Up    0  1 FF      3      -
R3-to-R2
172.30.5.2  172.30.5.6  Up    0  1 FF      3      -
R6-to-R2
172.30.5.2  172.30.5.1  Up    0  1 FF      3      -
R1-to-R2
172.30.5.2  172.30.5.7  Up    0  1 FF      3      -
R7-to-R2
Total 5 displayed, Up 5, Down 0

Transit LSP: 3 sessions
To          From          State Rt Style Labelin Labelout
LSPname
172.30.5.3  172.30.5.6  Up    0  1 FF  299856      3
R6-to-R3
172.30.5.6  172.30.5.1  Up    0  1 FF  299872      3
R1-to-R6
172.30.5.8  172.30.5.1  Up    0  1 FF  299888  299888
R1-to-R8
Total 3 displayed, Up 3, Down 0
```

Verify the content of the **inet.3** routing table.

```
[edit]
lab@R2# run show route table inet.3

inet.3: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.30.5.1/32      *[RSVP/7/1] 00:31:59, metric 11
                  > to 172.30.0.1 via ae0.0, label-switched-path
R2-to-R1-first
                  [LDP/9] 02:16:12, metric 11
                  > to 172.30.0.70 via ge-0/0/4.207, Push 100000
172.30.5.3/32      *[RSVP/7/1] 00:31:58, metric 10
                  > to 172.30.0.1 via ae0.0, label-switched-path
R2-to-R3
172.30.5.6/32      *[RSVP/7/1] 00:31:59, metric 20
                  > 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] 00:31:59, metric 10
                  > to 172.30.0.1 via ae0.0, label-switched-path
R2-to-R7
172.30.5.8/32      *[RSVP/7/1] 00:31:59, metric 10
                  > to 172.30.0.18 via ge-0/0/4.126, label-
switched-path R2-to-R8-first
172.30.5.41/32     *[LDP/9] 02:16:22, metric 10
                  > to 172.30.0.70 via ge-0/0/4.207
```

c. Configure RSVP LSPs on R3.

You have to 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-R1 admin-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-R6 admin-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 ]

[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

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

[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

[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)
+ = 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]
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.

```
[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
                  > 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]  
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: 5 sessions
To          From          State Rt Style Labelin Labelout
LSPname
172.30.5.8   172.30.5.7   Up    0  1 FF      3      -
R7-to-R8
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
172.30.5.8   172.30.5.2   Up    0  1 FF      3      -
R2-to-R8-first
172.30.5.8   172.30.5.6   Up    0  1 FF      3      -
R6-to-R8
Total 5 displayed, Up 5, 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 have to 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]
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 of 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

[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: 11 sessions, 5 detours
```

To	From	State	Rt	Style	Labelin	Labelout
LSPname						
172.30.5.1	172.30.5.8	Up	0	1 SE	3	-
R8-to-R1						
172.30.5.1	172.30.5.8	Up	0	1 SE	3	-
Bypass->172.30.0.9						
172.30.5.1	172.30.5.3	Up	0	1 SE	3	-
R3-to-R1						
172.30.5.1	172.30.5.3	Up	0	1 SE	3	-
R3-to-R1						
172.30.5.1	172.30.5.6	Up	0	1 SE	3	-
R6-to-R1						
172.30.5.1	172.30.5.6	Up	0	1 SE	3	-
R6-to-R1						
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	3	-
R2-to-R1-second						
172.30.5.1	172.30.5.2	Up	0	1 SE	3	-
Bypass->172.30.0.1						
172.30.5.1	172.30.5.7	Up	0	1 SE	3	-
R7-to-R1						
172.30.5.1	172.30.5.7	Up	0	1 SE	3	-
Bypass->172.30.0.5						

Total 11 displayed, Up 11, Down 0

```
Transit LSP: 30 sessions, 7 detours
```

To	From	State	Rt	Style	Labelin	Labelout
LSPname						
172.30.5.2	172.30.5.3	Up	0	1 SE	300368	3
R3-to-R2						

172.30.5.2	172.30.5.3	Up	0	1	SE	300464	3
R3-to-R2							
172.30.5.2	172.30.5.6	Up	0	1	SE	300608	300288
R6-to-R2							
172.30.5.2	172.30.5.6	Up	0	1	SE	300736	3
R6-to-R2							
172.30.5.2	172.30.5.6	Up	0	1	SE	300832	3
Bypass->172.30.0.17							
172.30.5.3	172.30.5.8	Up	0	1	SE	300880	300752
Bypass->172.30.0.33							
172.30.5.3	172.30.5.6	Up	0	1	SE	300704	300320
R6-to-R3							
172.30.5.3	172.30.5.6	Up	0	1	SE	300592	300256
R6-to-R3							
172.30.5.3	172.30.5.2	Up	0	1	SE	300320	300176
R2-to-R3							
172.30.5.3	172.30.5.2	Up	0	1	SE	300336	300144
Bypass->172.30.0.14							
172.30.5.3	172.30.5.7	Up	0	1	SE	300784	300416
R7-to-R3							
172.30.5.3	172.30.5.7	Up	0	1	SE	300800	300672
Bypass->172.30.0.29							
172.30.5.6	172.30.5.8	Up	0	1	SE	300864	300400
R8-to-R6							
172.30.5.6	172.30.5.3	Up	0	1	SE	300400	300192
R3-to-R6							
172.30.5.6	172.30.5.3	Up	0	1	SE	300512	300256
R3-to-R6							
172.30.5.6	172.30.5.2	Up	0	1	SE	300352	300160
Bypass->172.30.0.18							
172.30.5.6	172.30.5.7	Up	0	1	SE	300816	300720
Bypass->172.30.0.45							
172.30.5.7	172.30.5.8	Up	0	1	SE	300848	3
R8-to-R7							
172.30.5.7	172.30.5.3	Up	0	1	SE	300544	3
R3-to-R7							
172.30.5.7	172.30.5.3	Up	0	1	SE	300384	3
R3-to-R7							
172.30.5.7	172.30.5.6	Up	0	1	SE	300720	3
R6-to-R7							
172.30.5.7	172.30.5.6	Up	0	1	SE	300576	3
R6-to-R7							
172.30.5.7	172.30.5.2	Up	0	1	SE	300304	3
R2-to-R7							
172.30.5.8	172.30.5.3	Up	0	1	SE	300528	3
R3-to-R8							
172.30.5.8	172.30.5.3	Up	0	1	SE	300416	300224
R3-to-R8							
172.30.5.8	172.30.5.3	Up	0	1	SE	300560	3
Bypass->172.30.0.34							
172.30.5.8	172.30.5.6	Up	0	1	SE	300656	3
R6-to-R8							
172.30.5.8	172.30.5.6	Up	0	1	SE	300624	3
R6-to-R8							
172.30.5.8	172.30.5.6	Up	0	1	SE	300752	3
Bypass->172.30.0.22							

```

172.30.5.8      172.30.5.7      Up      0   1 SE   300768      3
R7-to-R8
Total 30 displayed, Up 30, 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: 8 sessions
To          From          State   Rt  Style  Labelin  Labelout
LSPname
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    -      300096
Bypass->172.30.0.2
172.30.5.3   172.30.5.1   Up      0   1 SE    -      300160
R1-to-R3
172.30.5.6   172.30.5.1   Up      0   1 SE    -      300128
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    -      300160
Bypass->172.30.0.6
172.30.5.8   172.30.5.1   Up      0   1 SE    -      300144
R1-to-R8
172.30.5.8   172.30.5.1   Up      0   1 SE    -      300688
Bypass->172.30.0.10
Total 8 displayed, Up 8, Down 0

Egress RSVP: 11 sessions, 5 detours
To          From          State   Rt  Style  Labelin  Labelout
LSPname
172.30.5.1   172.30.5.8   Up      0   1 SE    3        -
R8-to-R1
172.30.5.1   172.30.5.8   Up      0   1 SE    3        -
Bypass->172.30.0.9
172.30.5.1   172.30.5.3   Up      0   1 SE    3        -
R3-to-R1
172.30.5.1   172.30.5.3   Up      0   1 SE    3        -
R3-to-R1
172.30.5.1   172.30.5.6   Up      0   1 SE    3        -
R6-to-R1
172.30.5.1   172.30.5.6   Up      0   1 SE    3        -
R6-to-R1
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    3        -
R2-to-R1-second
172.30.5.1   172.30.5.2   Up      0   1 SE    3        -
Bypass->172.30.0.1
172.30.5.1   172.30.5.7   Up      0   1 SE    3        -
R7-to-R1
172.30.5.1   172.30.5.7   Up      0   1 SE    3        -
Bypass->172.30.0.5
Total 11 displayed, Up 11, Down 0

```

```

Transit RSVP: 31 sessions, 7 detours
To          From          State   Rt  Style  Labelin Labelout
LSPname
172.30.5.2  172.30.5.3      Up      0   1 SE   300368      3
R3-to-R2
172.30.5.2  172.30.5.3      Up      0   1 SE   300464      3
R3-to-R2
172.30.5.2  172.30.5.3      Up      0   1 SE   300896      3
Bypass->172.30.0.13
172.30.5.2  172.30.5.6      Up      0   1 SE   300608    300288
R6-to-R2
172.30.5.2  172.30.5.6      Up      0   1 SE   300736      3
R6-to-R2
172.30.5.2  172.30.5.6      Up      0   1 SE   300832      3
Bypass->172.30.0.17
172.30.5.3  172.30.5.8      Up      0   1 SE   300880    300752
Bypass->172.30.0.33
172.30.5.3  172.30.5.6      Up      0   1 SE   300704    300320
R6-to-R3
172.30.5.3  172.30.5.6      Up      0   1 SE   300592    300256
R6-to-R3
172.30.5.3  172.30.5.2      Up      0   1 SE   300320    300176
R2-to-R3
172.30.5.3  172.30.5.2      Up      0   1 SE   300336    300144
Bypass->172.30.0.14
172.30.5.3  172.30.5.7      Up      0   1 SE   300784    300416
R7-to-R3
172.30.5.3  172.30.5.7      Up      0   1 SE   300800    300672
Bypass->172.30.0.29
172.30.5.6  172.30.5.8      Up      0   1 SE   300864    300400
R8-to-R6
172.30.5.6  172.30.5.3      Up      0   1 SE   300400    300192
R3-to-R6
172.30.5.6  172.30.5.3      Up      0   1 SE   300512    300256
R3-to-R6
172.30.5.6  172.30.5.2      Up      0   1 SE   300352    300160
Bypass->172.30.0.18
172.30.5.6  172.30.5.7      Up      0   1 SE   300816    300720
Bypass->172.30.0.45
172.30.5.7  172.30.5.8      Up      0   1 SE   300848      3
R8-to-R7
172.30.5.7  172.30.5.3      Up      0   1 SE   300544      3
R3-to-R7
172.30.5.7  172.30.5.3      Up      0   1 SE   300384      3
R3-to-R7
172.30.5.7  172.30.5.6      Up      0   1 SE   300720      3
R6-to-R7
172.30.5.7  172.30.5.6      Up      0   1 SE   300576      3
R6-to-R7
172.30.5.7  172.30.5.2      Up      0   1 SE   300304      3
R2-to-R7
172.30.5.8  172.30.5.3      Up      0   1 SE   300528      3
R3-to-R8
172.30.5.8  172.30.5.3      Up      0   1 SE   300416    300224
R3-to-R8

```

172.30.5.8	172.30.5.3	Up	0	1	SE	300560	3
Bypass->172.30.0.34							
172.30.5.8	172.30.5.6	Up	0	1	SE	300656	3
R6-to-R8							
172.30.5.8	172.30.5.6	Up	0	1	SE	300624	3
R6-to-R8							
172.30.5.8	172.30.5.6	Up	0	1	SE	300752	3
Bypass->172.30.0.22							
172.30.5.8	172.30.5.7	Up	0	1	SE	300768	3
R7-to-R8							
Total 31 displayed, Up 31, 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 69, State Ena/Up
  Authentication, Aggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 172.30.0.1
  ActiveResv 18, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 2Gbps
  ct0: StaticBW 2Gbps, AvailableBW 900Mbps
    MaxAvailableBW 2Gbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 800Mbps[4] 0bps[5]
300Mbps[6] 0bps[7] 0bps
  Protection: On, Bypass: 1, LSP: 3, Protected LSP: 3, Unprotected
LSP: 0
    1 Nov 13 19:52:21 New bypass Bypass->172.30.0.2
    Bypass: Bypass->172.30.0.2, State: Up, Type: LP, LSP: 3,
Backup: 0
    3 Nov 13 19:52:23 Record Route: 172.30.0.10 172.30.0.33
172.30.0.13
    2 Nov 13 19:52:23 Up
    1 Nov 13 19:52:23 CSPF: computation result accepted
--- (more) ---
```

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]
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 *    path-1
R2-to-R1-second
172.30.5.1   172.30.5.2   Up    0 *    primary-path
R2-to-R1-first
172.30.5.3   172.30.5.2   Up    0 *    primary-path
R2-to-R3
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 *    path-1
R2-to-R8-second
172.30.5.8   172.30.5.2   Up    0 *    primary-path
R2-to-R8-first
Total 7 displayed, Up 7, Down 0

Egress LSP: 5 sessions
To          From          State Rt Style Labelin Labelout
LSPname
172.30.5.2   172.30.5.8   Up    0  1 FF      3      -
R8-to-R2
172.30.5.2   172.30.5.3   Up    0  1 FF      3      -
R3-to-R2
172.30.5.2   172.30.5.6   Up    0  1 FF      3      -
R6-to-R2
172.30.5.2   172.30.5.1   Up    0  1 FF      3      -
R1-to-R2
172.30.5.2   172.30.5.7   Up    0  1 FF      3      -
R7-to-R2
Total 5 displayed, Up 5, Down 0
```

```

Transit LSP: 3 sessions
To          From          State   Rt Style Labelin Labelout
LSPname
172.30.5.3   172.30.5.6   Up      0  1 FF  299856      3
R6-to-R3
172.30.5.6   172.30.5.1   Up      0  1 FF  299872      3
R1-to-R6
172.30.5.8   172.30.5.1   Up      0  1 FF  299888    299888
R1-to-R8
Total 3 displayed, Up 3, 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)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 50Mbps MaxBW: 100Mbps
  AdjustTimer: 86400 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 64523 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-path State: Up
    Priorities: 5 5
    Bandwidth: 50Mbps
    SmartOptimizeTimer: 180
    Include Any: blue green
    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):
      172.30.0.1
        5 Nov 13 18:44:19.225 Selected as active path
        4 Nov 13 18:44:19.223 Record Route: 172.30.0.1
        3 Nov 13 18:44:19.222 Up
        2 Nov 13 18:44:19.027 Originate Call
        1 Nov 13 18:44:19.027 CSPF: computation result accepted
  172.30.0.1
    Secondary secondary-path State: Dn
    Priorities: 5 5
    Bandwidth: 50Mbps
    SmartOptimizeTimer: 180
    Include Any: blue green
    No computed ERO.

```

Created: Wed Nov 13 18:44:19 2013
Total 1 displayed, Up 1, Down 0

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    -        300032
R2-to-R1-second
172.30.5.1   172.30.5.2   Up      0   1 SE    -        300144
Bypass->172.30.0.1
172.30.5.3   172.30.5.2   Up      0   1 SE    -        300320
R2-to-R3
172.30.5.3   172.30.5.2   Up      0   1 SE    -        300336
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    -        300352
Bypass->172.30.0.18
172.30.5.7   172.30.5.2   Up      0   1 SE    -        300304
R2-to-R7
172.30.5.8   172.30.5.2   Up      0   1 SE    -        300128
R2-to-R8-first
172.30.5.8   172.30.5.2   Up      0   1 SE    -        300016
R2-to-R8-second
Total 10 displayed, Up 10, Down 0

Egress RSVP: 10 sessions, 7 detours
To          From          State   Rt  Style  Labelin  Labelout
LSPname
172.30.5.2   172.30.5.8   Up      0   1 SE    3         -
R8-to-R2
172.30.5.2   172.30.5.3   Up      0   1 SE    3         -
R3-to-R2
172.30.5.2   172.30.5.3   Up      0   1 SE    3         -
R3-to-R2
172.30.5.2   172.30.5.3   Up      0   1 SE    3         -
Bypass->172.30.0.13
172.30.5.2   172.30.5.6   Up      0   1 SE    3         -
R6-to-R2
172.30.5.2   172.30.5.6   Up      0   1 SE    3         -
R6-to-R2
172.30.5.2   172.30.5.6   Up      0   1 SE    3         -
Bypass->172.30.0.17
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   172.30.5.7   Up      0   1 SE    3         -
R7-to-R2
```

Total 10 displayed, Up 10, Down 0						
Transit RSVP: 28 sessions, 8 detours						
To	From	State	Rt	Style	Labelin	Labelout
LSPname						
172.30.5.1	172.30.5.8	Up	0	1 SE	300736	3
Bypass->172.30.0.9						
172.30.5.1	172.30.5.3	Up	0	1 SE	300240	3
R3-to-R1						
172.30.5.1	172.30.5.3	Up	0	1 SE	300176	3
R3-to-R1						
172.30.5.1	172.30.5.6	Up	0	1 SE	300480	3
R6-to-R1						
172.30.5.1	172.30.5.6	Up	0	1 SE	300464	300192
R6-to-R1						
172.30.5.1	172.30.5.7	Up	0	1 SE	300704	3
Bypass->172.30.0.5						
172.30.5.3	172.30.5.8	Up	0	1 SE	300752	3
Bypass->172.30.0.33						
172.30.5.3	172.30.5.6	Up	0	1 SE	300512	3
R6-to-R3						
172.30.5.3	172.30.5.6	Up	0	1 SE	300416	3
R6-to-R3						
172.30.5.3	172.30.5.7	Up	0	1 SE	300672	3
Bypass->172.30.0.29						
172.30.5.6	172.30.5.3	Up	0	1 SE	300256	3
R3-to-R6						
172.30.5.6	172.30.5.3	Up	0	1 SE	300192	3
R3-to-R6						
172.30.5.6	172.30.5.1	Up	0	1 SE	300128	3
R1-to-R6						
172.30.5.6	172.30.5.7	Up	0	1 SE	300720	3
Bypass->172.30.0.45						
172.30.5.7	172.30.5.3	Up	0	1 SE	300336	300544
R3-to-R7						
172.30.5.7	172.30.5.3	Up	0	1 SE	300208	300384
R3-to-R7						
172.30.5.7	172.30.5.6	Up	0	1 SE	300656	300304
R6-to-R7						
172.30.5.7	172.30.5.6	Up	0	1 SE	300624	300720
R6-to-R7						
172.30.5.7	172.30.5.6	Up	0	1 SE	300448	300176
R6-to-R7						
172.30.5.7	172.30.5.1	Up	0	1 SE	300160	300112
Bypass->172.30.0.6						
172.30.5.8	172.30.5.3	Up	0	1 SE	300224	300192
R3-to-R8						
172.30.5.8	172.30.5.3	Up	0	1 SE	300320	300528
R3-to-R8						
172.30.5.8	172.30.5.3	Up	0	1 SE	300352	300560
Bypass->172.30.0.34						
172.30.5.8	172.30.5.6	Up	0	1 SE	300576	300256
R6-to-R8						
172.30.5.8	172.30.5.6	Up	0	1 SE	300432	300144
R6-to-R8						

172.30.5.8	172.30.5.6	Up	0	1	SE	300640	300752
Bypass->172.30.0.22							
172.30.5.8	172.30.5.1	Up	0	1	SE	300144	300096
R1-to-R8							
172.30.5.8	172.30.5.1	Up	0	1	SE	300688	300336
Bypass->172.30.0.10							
Total 28 displayed, Up 28, 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
```

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.

```
[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.

```

[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: 9 sessions, 3 detours

```

To	From	State	Rt	Style	Labelin	Labelout
LSPname						
172.30.5.3	172.30.5.8	Up	0	1 SE	3	-
R8-to-R3						
172.30.5.3	172.30.5.8	Up	0	1 SE	3	-
Bypass->172.30.0.33						

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	-
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	172.30.5.1	Up	0	1	SE	3	-
R1-to-R3							
172.30.5.3	172.30.5.7	Up	0	1	SE	3	-
R7-to-R3							
172.30.5.3	172.30.5.7	Up	0	1	SE	3	-
Bypass->172.30.0.29							
Total 9 displayed, Up 9, Down 0							
Transit LSP: 12 sessions, 6 detours							
To	From	State	Rt	Style	Labelin	Labelout	
LSPname							
172.30.5.1	172.30.5.8	Up	0	1	SE	300336	300736
Bypass->172.30.0.9							
172.30.5.1	172.30.5.6	Up	0	1	SE	300192	300304
R6-to-R1							
172.30.5.1	172.30.5.6	Up	0	1	SE	300224	300480
R6-to-R1							
172.30.5.1	172.30.5.2	Up	0	1	SE	300032	300112
R2-to-R1-second							
172.30.5.1	172.30.5.7	Up	0	1	SE	300320	300704
Bypass->172.30.0.5							
172.30.5.2	172.30.5.8	Up	0	1	SE	300352	3
R8-to-R2							
172.30.5.2	172.30.5.6	Up	0	1	SE	300160	3
R6-to-R2							
172.30.5.2	172.30.5.1	Up	0	1	SE	300000	3
Bypass->172.30.0.2							
172.30.5.7	172.30.5.6	Up	0	1	SE	300176	300272
R6-to-R7							
172.30.5.8	172.30.5.6	Up	0	1	SE	300144	3
R6-to-R8							
172.30.5.8	172.30.5.6	Up	0	1	SE	300272	3
R6-to-R8							
172.30.5.8	172.30.5.2	Up	0	1	SE	300016	3
R2-to-R8-second							
Total 12 displayed, Up 12, Down 0							

Verify the RSVP session table.

```
lab@R3# run show rsvp session | no-more
```

Ingress RSVP: 12 sessions

To	From	State	Rt	Style	Labelin	Labelout
LSPname						
172.30.5.1	172.30.5.3	Up	0	1 SE	-	300192
R3-to-R1						
172.30.5.1	172.30.5.3	Up	0	1 SE	-	300208
R3-to-R1						
172.30.5.2	172.30.5.3	Up	0	1 SE	-	300208
R3-to-R2						
172.30.5.2	172.30.5.3	Up	0	1 SE	-	300192
R3-to-R2						
172.30.5.2	172.30.5.3	Up	0	1 SE	-	300448
Bypass->172.30.0.13						
172.30.5.6	172.30.5.3	Up	0	1 SE	-	300176
R3-to-R6						
172.30.5.6	172.30.5.3	Up	0	1 SE	-	300224
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	-	300224
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	-	300240
R3-to-R8						
172.30.5.8	172.30.5.3	Up	0	1 SE	-	300352
Bypass->172.30.0.34						
Total 12 displayed, Up 12, Down 0						

Egress RSVP: 9 sessions, 3 detours

To	From	State	Rt	Style	Labelin	Labelout
LSPname						
172.30.5.3	172.30.5.8	Up	0	1 SE	3	-
R8-to-R3						
172.30.5.3	172.30.5.8	Up	0	1 SE	3	-
Bypass->172.30.0.33						
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	-
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	172.30.5.1	Up	0	1 SE	3	-
R1-to-R3						
172.30.5.3	172.30.5.7	Up	0	1 SE	3	-
R7-to-R3						

172.30.5.3	172.30.5.7	Up	0	1	SE	3	-
Bypass->172.30.0.29							
Total 9 displayed, Up 9, Down 0							
Transit RSVP: 12 sessions, 6 detours							
To	From	State	Rt	Style	Labelin	Labelout	
LSPname							
172.30.5.1	172.30.5.8	Up	0	1	SE	300336	300736
Bypass->172.30.0.9							
172.30.5.1	172.30.5.6	Up	0	1	SE	300192	300304
R6-to-R1							
172.30.5.1	172.30.5.6	Up	0	1	SE	300224	300480
R6-to-R1							
172.30.5.1	172.30.5.2	Up	0	1	SE	300032	300112
R2-to-R1-second							
172.30.5.1	172.30.5.7	Up	0	1	SE	300320	300704
Bypass->172.30.0.5							
172.30.5.2	172.30.5.8	Up	0	1	SE	300352	3
R8-to-R2							
172.30.5.2	172.30.5.6	Up	0	1	SE	300160	3
R6-to-R2							
172.30.5.2	172.30.5.1	Up	0	1	SE	300000	3
Bypass->172.30.0.2							
172.30.5.7	172.30.5.6	Up	0	1	SE	300176	300272
R6-to-R7							
172.30.5.8	172.30.5.6	Up	0	1	SE	300144	3
R6-to-R8							
172.30.5.8	172.30.5.6	Up	0	1	SE	300272	3
R6-to-R8							
172.30.5.8	172.30.5.2	Up	0	1	SE	300016	3
R2-to-R8-second							
Total 12 displayed, Up 12, 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 74, State Ena/Up
  Authentication, Aggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 172.30.0.14
  ActiveResv 13, PreemptionCnt 0, Update threshold 10%
  Subscription 120%,
  bc0 = ct0, StaticBW 1000Mbps
  ct0: StaticBW 1000Mbps, AvailableBW 300Mbps
    MaxAvailableBW 1.2Gbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 800Mbps[4] 0bps[5]
    100Mbps[6] 0bps[7] 0bps
  Protection: On, Bypass: 1, LSP: 1, Protected LSP: 1, Unprotected
  LSP: 0
    1 Nov 13 20:05:51 New bypass Bypass->172.30.0.13
    Bypass: Bypass->172.30.0.13, State: Up, Type: LP, LSP: 1,
  Backup: 0
    3 Nov 13 20:05:53 Record Route: 172.30.0.34 172.30.0.9
    172.30.0.2
    2 Nov 13 20:05:53 Up
    1 Nov 13 20:05:52 CSPF: computation result accepted
```

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
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary primary-path State: Up
    Priorities: 3 3
    Bandwidth: 100Mbps
    SmartOptimizeTimer: 180
    Include Any: blue green
    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) 172.30.0.5
    26 Nov 13 20:00:15.857 Fast-reroute Detour Up
    25 Nov 13 20:00:10.393 Record Route: 172.30.0.30(flag=5)
172.30.0.5
  24 Nov 13 20:00:08.861 Make-before-break: Switched to new
instance
  23 Nov 13 20:00:07.081 Record Route: 172.30.0.30 172.30.0.5
  22 Nov 13 20:00:07.081 Up
  21 Nov 13 20:00:06.635 Originate make-before-break call
  20 Nov 13 20:00:06.635 CSPF: computation result accepted
172.30.0.30 172.30.0.5
  19 Nov 13 19:07:13.358 Selected as active path
  18 Nov 13 19:07:13.357 Record Route: 172.30.0.30 172.30.0.5
  17 Nov 13 19:07:13.357 Up
  16 Nov 13 19:07:12.938 Originate Call
  15 Nov 13 19:07:12.938 CSPF: computation result accepted
172.30.0.30 172.30.0.5
  14 Nov 13 19:07:12.916 Clear Call
  13 Nov 13 19:07:12.915 Deselected as active
  12 Nov 13 19:07:06.996 Selected as active path
  11 Nov 13 19:07:06.995 Record Route: 172.30.0.30 172.30.0.5
  10 Nov 13 19:07:06.995 Up
  9 Nov 13 19:07:06.567 Originate Call
  8 Nov 13 19:07:06.567 CSPF: computation result accepted
172.30.0.30 172.30.0.5
  7 Nov 13 19:07:06.524 Clear Call
  6 Nov 13 19:07:06.523 Deselected as active
  5 Nov 13 19:03:50.111 Selected as active path
  4 Nov 13 19:03:50.110 Record Route: 172.30.0.30 172.30.0.5
  3 Nov 13 19:03:50.108 Up
  2 Nov 13 19:03:49.682 Originate Call
```

```

1 Nov 13 19:03:49.682 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
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) 172.30.0.21(flag=0xd)
172.30.0.17(flag=5) 172.30.0.1
18 Nov 13 20:01:54.871 Fast-reroute Detour Up
17 Nov 13 20:00:10.612 Record Route: 172.30.0.34(flag=0xd)
172.30.0.21(flag=0xd) 172.30.0.17(flag=5) 172.30.0.1
16 Nov 13 20:00:10.511 Record Route: 172.30.0.34(flag=0xd)
172.30.0.21(flag=0xd) 172.30.0.17 172.30.0.1
15 Nov 13 20:00:10.073 Record Route: 172.30.0.34(flag=0xd)
172.30.0.21 172.30.0.17 172.30.0.1
14 Nov 13 20:00:07.404 Record Route: 172.30.0.34 172.30.0.21
172.30.0.17 172.30.0.1
13 Nov 13 20:00:07.404 Up
12 Nov 13 20:00:06.642 Originate Call
11 Nov 13 20:00:06.641 Clear Call
10 Nov 13 20:00:06.641 CSPF: computation result accepted
172.30.0.34 172.30.0.21 172.30.0.17 172.30.0.1
9 Nov 13 19:07:42.587 Record Route: 172.30.0.34 172.30.0.21
172.30.0.17 172.30.0.1
8 Nov 13 19:07:42.587 Up
7 Nov 13 19:07:41.742 Originate Call
6 Nov 13 19:07:41.742 CSPF: computation result accepted
172.30.0.34 172.30.0.21 172.30.0.17 172.30.0.1
5 Nov 13 19:07:06.525 Clear Call
4 Nov 13 19:04:19.617 Record Route: 172.30.0.34 172.30.0.21
172.30.0.17 172.30.0.1
3 Nov 13 19:04:19.617 Up
2 Nov 13 19:04:18.833 Originate Call
1 Nov 13 19:04:18.833 CSPF: computation result accepted
172.30.0.34 172.30.0.21 172.30.0.17 172.30.0.1
Created: Wed Nov 13 19:03:49 2013
Total 1 displayed, Up 1, Down 0

```

d. Configure LSP traffic protection on R6.

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

[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

[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: 6 sessions
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   172.30.5.2   Up    0  1 FF      3      -
R2-to-R6
172.30.5.6   172.30.5.1   Up    0  1 FF      3      -
R1-to-R6
172.30.5.6   172.30.5.7   Up    0  1 FF      3      -
R7-to-R6
Total 6 displayed, Up 6, Down 0

Transit LSP: 8 sessions
To          From          State Rt Style Labelin Labelout
LSPname
172.30.5.1   172.30.5.3   Up    0  1 SE  299984  299968
R3-to-R1
172.30.5.2   172.30.5.8   Up    0  1 FF  299840      3
R8-to-R2
172.30.5.2   172.30.5.3   Up    0  1 SE  300000      3
R3-to-R2
172.30.5.2   172.30.5.7   Up    0  1 FF  299792      3
R7-to-R2
172.30.5.7   172.30.5.3   Up    0  1 SE  300032  300016
R3-to-R7
172.30.5.8   172.30.5.3   Up    0  1 SE  300016      3
R3-to-R8
172.30.5.8   172.30.5.2   Up    0  1 FF  299904      3
R2-to-R8-first
172.30.5.8   172.30.5.1   Up    0  1 FF  299888      3
R1-to-R8
Total 8 displayed, Up 8, Down 0
```


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: 12 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 8 sessions, 2 detours
Total 0 displayed, Up 0, Down 0

Transit RSVP: 12 sessions, 3 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: 300240
  Resv style: 1 SE, Label in: 300256, Label out: 300240
  Time left: 131, Since: Wed Nov 13 20:00:09 2013
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
  Port number: sender 7 receiver 18023 protocol 0
  Detour branch from 172.30.0.17, to skip 172.30.5.1, Up
    Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 172.30.0.17 (ge-0/0/4.126) 28 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 172.30.0.22 (ge-0/0/4.168) 3 pkts
    RESV rcvfrom: 172.30.0.22 (ge-0/0/4.168) 3 pkts
    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: 300256, Label out: 300240

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: 300176
  Resv style: 1 SE, Label in: 300176, Label out: 300176
  Time left: 119, Since: Wed Nov 13 20:00:06 2013
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
  Port number: sender 7 receiver 18023 protocol 0
  FastReroute desired
  PATH rcvfrom: 172.30.0.22 (ge-0/0/4.168) 26 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 172.30.0.17 (ge-0/0/4.126) 8 pkts
  RESV rcvfrom: 172.30.0.17 (ge-0/0/4.126) 7 pkts
  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
1500
    Detour adspec: received MTU 1500 sent MTU 1500
    Path MTU: received 1500
    Detour PATH sentto: 172.30.0.46 (ge-0/0/4.167) 4 pkts
    Detour RESV rcvfrom: 172.30.0.46 (ge-0/0/4.167) 1 pkts
    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: 300272
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

[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

```

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: 7 sessions							
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	172.30.5.6	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 FF	3	-	
R2-to-R7							
172.30.5.7	172.30.5.1	Up	0	1 FF	3	-	
R1-to-R7							
Total 7 displayed, Up 7, Down 0							
Transit LSP: 12 sessions							
To	From	State	Rt	Style	Labelin	Labelout	
LSPname							
172.30.5.1	172.30.5.3	Up	0	1 SE	300016	3	
R3-to-R1							
172.30.5.1	172.30.5.6	Up	0	1 SE	300080	3	
R6-to-R1							
172.30.5.2	172.30.5.8	Up	0	1 FF	299808	299840	
R8-to-R2							
172.30.5.2	172.30.5.3	Up	0	1 SE	300032	300096	
R3-to-R2							
172.30.5.2	172.30.5.6	Up	0	1 SE	300128	300208	
R6-to-R2							
172.30.5.3	172.30.5.6	Up	0	1 SE	300112	300176	
R6-to-R3							
172.30.5.3	172.30.5.2	Up	0	1 FF	299904	3	
R2-to-R3							
172.30.5.3	172.30.5.1	Up	0	1 FF	299888	3	
R1-to-R3							
172.30.5.6	172.30.5.8	Up	0	1 FF	299824	3	
R8-to-R6							
172.30.5.6	172.30.5.3	Up	0	1 SE	300048	300128	
R3-to-R6							
172.30.5.8	172.30.5.3	Up	0	1 SE	300064	300144	
R3-to-R8							
172.30.5.8	172.30.5.6	Up	0	1 SE	300096	300160	
R6-to-R8							
Total 12 displayed, Up 12, 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 label-switched-path R8-to-R1 primary
primary-path

[edit]
lab@R8# set protocols label-switched-path R8-to-R2 primary
primary-path

[edit]
lab@R8# set protocols label-switched-path R8-to-R3 primary
primary-path

[edit]
lab@R8# set protocols label-switched-path R8-to-R6 primary
primary-path

[edit]
lab@R8# set protocols label-switched-path R8-to-R7 primary
primary-path

[edit]
lab@R8# set protocols label-switched-path R8-to-R1 secondary
secondary-path

[edit]
lab@R8# set protocols label-switched-path R8-to-R2 secondary
secondary-path

[edit]
lab@R8# set protocols label-switched-path R8-to-R3 secondary
secondary-path

[edit]
lab@R8# set protocols label-switched-path R8-to-R6 secondary
secondary-path

[edit]
lab@R8# set protocols label-switched-path R8-to-R7 secondary
secondary-path

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.118 link-protection

[edit]
lab@R8# set protocols rsvp interface ge-0/0/4.138 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: 8 sessions

```

To	From	State	Rt	Style	Labelin	Labelout
LSPname						
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	172.30.5.6	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	172.30.5.2	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
172.30.5.8      172.30.5.1      Up      0   1  FF      3      -
R1-to-R8
172.30.5.8      172.30.5.7      Up      0   1  FF      3      -
R7-to-R8
Total 8 displayed, Up 8, Down 0

Transit LSP: 10 sessions
To          From          State   Rt  Style  Labelin Labelout
LSPname
172.30.5.1   172.30.5.3      Up      0   1  SE   299952  299984
R3-to-R1
172.30.5.1   172.30.5.6      Up      0   1  SE   300048      3
R6-to-R1
172.30.5.1   172.30.5.2      Up      0   1  FF   299840      3
R2-to-R1-second
172.30.5.1   172.30.5.7      Up      0   1  FF   300064      3
R7-to-R1
172.30.5.2   172.30.5.3      Up      0   1  SE   299968  300000
R3-to-R2
172.30.5.2   172.30.5.6      Up      0   1  SE   300032  299888
R6-to-R2
172.30.5.3   172.30.5.6      Up      0   1  SE   300000      3
R6-to-R3
172.30.5.6   172.30.5.3      Up      0   1  SE   299936      3
R3-to-R6
172.30.5.7   172.30.5.3      Up      0   1  SE   299984  300032
R3-to-R7
172.30.5.7   172.30.5.6      Up      0   1  SE   300016  300192
R6-to-R7
Total 10 displayed, Up 10, Down 0

```

5) Configure LDP tunneling from R4 and R5 to the Route Reflector.

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
```

```
[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
```

```
[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	100.0	172.30.5.2:0	44
172.30.5.7	100.0	172.30.5.7:0	32
172.30.5.8	100.0	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
```

```
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

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
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** knob.

```

[edit]
lab@R6# run show route table inet.3 172.30.5.1 hidden

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.

```

[edit]

```

```
lab@R3# set protocols mpls label-switched-path R3-to-R1 ldp-
tunneling
```

```
[edit]
lab@R3# commit
commit complete
```

```
[edit]
lab@R6# set protocols mpls label-switched-path R6-to-R1 ldp-
tunneling
```

```
[edit]
lab@R6# commit
commit complete
```

```
[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.

```
[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
```

```
[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 via 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 other 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
```

```
[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          172.30.0.1      indr 262143 403
ae0.0                ucst      620         2
```


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.

```
[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
--- (more) ---
```

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.

```
[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          172.30.0.14      indr 262145      17
                                     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
--- (more) ---
```

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.

The load balancing policy should be configured on all routes in the same was a done on R1.

```
[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.

```
[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
      > 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
--- (more) ---
```

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**

```
[edit]
lab@R1# set protocols bgp group ibgp family inet6 labeled-unicast
explicit-null
```

Don't forget 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
Table          Tot Paths  Act Paths Suppressed    History Damp
State          Pending
inet.0
0              0          650          488          0          0
inet6.0
0              0          32          32          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, 53 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(262144)}
Page 0 idx 0 Type 1 val 17f855c
  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
             Address: 0x155cf34
             Next-hop reference count: 48
             Source: 172.30.5.41
             Next hop type: Router, Next hop index: 262143
             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
             Next hop: 172.30.0.10 via ge-0/0/4.118 weight
0x8001
             Label-switched-path Bypass->172.30.0.2
             Label operation: Push 2, Push 300096(top)
             Label TTL action: prop-ttl, prop-ttl(top)
             Protocol next hop: ::ffff:172.30.5.2
             Push 2
             Indirect next hop: 183c000 262144
             State: <Active Int Ext>
             Local AS: 54591 Peer AS: 54591
             Age: 9:48      Metric2: 11
             Task: BGP_54591.172.30.5.41+179
             Announcement bits (3): 0-KRT 2-BGP_RT_Background
3-Resolve tree 3
             AS path: 64514 I (Originator)
             Cluster list: 0.0.0.2
             Originator ID: 172.30.5.2
             Communities: 54591:200
             Accepted
             Route Label: 2
             Localpref: 200
             Router ID: 172.30.5.41
             Indirect next hops: 1
             Protocol next hop: ::ffff:172.30.5.2
Metric: 11
             Push 2
             Indirect next hop: 183c000 262144
             Indirect path forwarding next hops: 2
```

```

                                Next hop type: Router
                                Next hop: 172.30.0.2 via ae0.0
weight 0x1
                                Next hop: 172.30.0.10 via ge-
0/0/4.118 weight 0x8001
                                ::ffff:172.30.5.2/128 Originating RIB:
inet6.3
                                Metric: 11
                                Node
path count: 1
                                Forwarding nexthops: 2
                                Nexthop: 172.30.0.2 via ae0.0
                                Nexthop: 172.30.0.10 via ge-
0/0/4.118

```

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
*
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.

- 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.

```
[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.

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
```

```
[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
```

```
[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]
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
--- (more) ---
[edit]
lab@R7# 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] 00:01:02, 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
--- (more) ---

```

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 a 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.

```

[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
```

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				
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
--- (more) ---
```

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 have to ensure that the routes from both VPN sites are part of a different OSPF domain.

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
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				
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 have to 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

[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
```

You have to 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 spokes 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"

[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.

```
[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 exports 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

[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-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

172.30.5.33/32      *[Direct/0] 00:25:05
                    > via lo0.1
172.31.64.0/24     *[BGP/170] 00:25:01, localpref 100
                    AS path: 64600 I
                    > to 192.168.0.42 via ge-0/0/5.311
172.31.65.0/24     *[BGP/170] 00:25:01, localpref 100
                    AS path: 64600 I
                    > to 192.168.0.42 via ge-0/0/5.311
172.31.66.0/24     *[BGP/170] 00:25:01, localpref 100
                    AS path: 64600 I
                    > to 192.168.0.42 via ge-0/0/5.311
172.31.67.0/24     *[BGP/170] 00:25:01, localpref 100
                    AS path: 64600 I
                    > to 192.168.0.42 via ge-0/0/5.311
--- (more) ---

```

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
--- (more) ---

```

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
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)

```

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	Localpref
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			
* 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			

- 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. As you cannot use **rib-groups** to accomplish this requirement so you have to use another solution. The **auto-export** knob 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

[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
                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

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.

```
[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 have to 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 summarises all the routes used by CE2 customer. This summary route should be advertised to the Internet as well.

You need to 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
```

```
[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.

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 target 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** knob under **protocols ldp** stanza. 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 a 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.

```
[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 have to 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
```

```
[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
```

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)
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.

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 received-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.

```
[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 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
holdddown, 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
              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.

```
[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
7:172.30.5.7:5:54591:32:172.31.65.5:32:232.1.1.1/240

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
    *[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

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.

```
[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
```

```
[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]
root@route-reflector# set protocols bgp group P2-remote-pe local-
address 172.30.5.41

```

```

[edit]
root@route-reflector# set protocols bgp group P2-remote-pe
neighbor 172.17.47.3

```

Because the local Route Reflector router is using an old JUNOS version and the remote AS is using a 4-byte AS, you have to configure the **peer-as** with transition AS 23456.

```

[edit]
root@route-reflector# set protocols bgp group P2-remote-pe peer-as
23456

```

You have to configure the session as multihop without changing the next-hop.

```

[edit]
root@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]
root@route-reflector# set protocols bgp group P2-remote-pe family
inet-vpn unicast

```

The Labeled BGP family should be enabled on the IBGP session.

```

[edit]

```

```
root@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# 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 ibgp family inet labeled-unicast rib
inet.3

[edit]
lab@R7# commit
commit complete
```

Verify BGP operation

```
[edit]
root@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          0          505          505          0          0
0              0
inet.3          0           1           1           0          0
0              0
bgp.l3vpn.0     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]
root@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

```

Routes received from the remote PE are attached with a route target that is not recognized by R7. So you have to normalize the route targets.

```

[edit]
root@route-reflector# set policy-options policy-statement CE2-vpn-
target-export term 1 from protocol bgp

```

```

[edit]

```

```

root@route-reflector# set policy-options policy-statement CE2-vpn-
target-export term 1 from community CE2-hub

[edit]
root@route-reflector# set policy-options policy-statement CE2-vpn-
target-export term 1 then community delete CE2-hub

[edit]
root@route-reflector# set policy-options policy-statement CE2-vpn-
target-export term 1 then community add CE2-remote

[edit]
root@route-reflector# set policy-options policy-statement CE2-vpn-
target-export term 1 then accept

[edit]
root@route-reflector# set policy-options policy-statement CE2-vpn-
target-import term 1 from protocol bgp

[edit]
root@route-reflector# set policy-options policy-statement CE2-vpn-
target-import term 1 from community CE2-remote

[edit]
root@route-reflector# set policy-options policy-statement CE2-vpn-
target-import term 1 then community delete CE2-remote

[edit]
root@route-reflector# set policy-options policy-statement CE2-vpn-
target-import term 1 then community add CE2-spoke

[edit]
root@route-reflector# set policy-options policy-statement CE2-vpn-
target-import term 1 then accept

[edit]
root@route-reflector# set policy-options community CE2-hub members
target:54591:200

[edit]
root@route-reflector# set policy-options community CE2-remote
members target:43208:200

[edit]
root@route-reflector# set policy-options community CE2-spoke
members target:54591:201

[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.0 172.31.7[7-8]

172.31.78.0/24      *[BGP/170] 00:00:03, localpref 100, from
172.30.5.41
                  AS path: 2831679853 64600 I
                  > to 172.30.0.5 via ge-0/0/4.117, label-
switched-path R7-to-R8
                  [BGP/170] 00:00:03, localpref 100
                  AS path: 64600 54591 2831679853 54591 I
                  > to 192.168.0.42 via ge-0/0/5.311
172.31.79.0/24      *[BGP/170] 00:00:03, localpref 100, from
172.30.5.41
                  AS path: 2831679853 64600 I
                  > to 172.30.0.5 via ge-0/0/4.117, label-
switched-path R7-to-R8
                  [BGP/170] 00:00:03, localpref 100
                  AS path: 64600 54591 2831679853 54591 I
                  > 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
                Address: 0x1881278
                Next-hop reference count: 3
                Source: 172.30.5.41
                Next hop type: Router
                Next hop: 172.30.0.5 via ge-0/0/4.117 weight 0x1,
selected
                Label-switched-path R7-to-R8
                Label operation: Push 300064, Push 300256(top)
                Label TTL action: prop-ttl, prop-ttl(top)
                Next hop: 172.30.0.45 via ge-0/0/4.167 weight
0x8001
                Label-switched-path Bypass->172.30.0.5
                Label operation: Push 300064, Push 300256, Push
299920(top)
                Label TTL action: prop-ttl, prop-ttl, prop-
ttl(top)
                Protocol next hop: 172.30.5.8
                Push 300064
                Indirect next hop: 18889f8 -
                State: <Active Int Ext>
                Local AS: 54591 Peer AS: 54591
                Age: 1:07      Metric2: 20
```

9) Configure customer CE3 VPLS service.

Customer CE3 requires a VPLS service between R2 and R3. the site connected to R3 is multi-homed 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
```

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
```

```
[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-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 multi-homing.

```
[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
```

```
[edit]
lab@R3# commit
commit complete
```

```
[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

[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

```

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                    WE -- interface and instance
EM -- encapsulation mismatch     encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not
present                         -> -- only outbound connection is
CM -- control-word mismatch     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: CE3-vpls
  Local site: site-1 (1)
    connection-site
Type  St      Time last up  #
Up trans
  2      rmt Up      Nov 14 22:26:40 2013
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

Connection to site 2 on primary PE, R3, is in Up state.

[edit]
lab@R2# run show route forwarding-table family vpls

```

```

Routing table: CE3-vpls.vpls
VPLS:
Destination          Type RtRef Next hop          Type Index NhRef
Netif
default              perm    0          rjct    768    1
ge-0/0/3.601         user    0          comp    797    2
lsi.1048576          user    0          comp    803    2
00:23:9c:8b:6c:96/48 dynm      0          ucst    793    3
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)
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: CE3-vpls
  Local site: site-2 (2)
    connection-site          Type  St      Time last up      #
Up trans
  1                          rmt   Up      Nov 14 22:26:40 2013
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
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)
* 172.30.5.3:4:2:1/96 (1 entry, 1 announced)
  BGP group ibgp type Internal
    Route Distinguisher: 172.30.5.3:4
    Label-base: 262145, range: 8
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 65535
    AS path: [54591] I
    Communities: target:54591:300 Layer2-info: encaps:VPLS,
control flags:, 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)
* 172.30.5.4:5:2:1/96 (1 entry, 1 announced)
  BGP group ibgp type Internal
    Route Distinguisher: 172.30.5.4:5
    Label-base: 262145, range: 8
    Nexthop: Self
    Flags: Nexthop Change
    Localpref: 1
    AS path: [54591] I
    Communities: target:54591:300 Layer2-info: encaps:VPLS,
control flags:, mtu: 0, site preference: 1
```

Configure the VPLS that if 100 MAC addresses are learned that packets are dropped.

```
[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
```


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.

- 1) Configure code-point aliases.

```
[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

[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 48.

```
[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

[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 45.

```
[edit]
lab@R4# set class-of-service forwarding-classes queue 0 be

[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 47.

```
[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
```

```
[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 48.

```
[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

[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 46.

```
[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

[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 45.

```
[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 interfaces names. All core facing interfaces should be configured in the way shown as previously shown on R4.

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.

```
[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
```

9) Configure a Multi-field classifier for the L3VPN customer interfaces according to Table 44.

```
[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
```

```
[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
```

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 Multi-field classifier on the interfaces.

```
[edit]
lab@R4# set interfaces ge-0/0/5.323 family inet filter input
l3vpn-classifier
```


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 stanza.

The VPLS filter will police all the traffic that enters the interface.

```
[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 Multi-field classifier on the interfaces.

```
[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         13vpn                  low
  001010         12vpn                  low
  001011         12vpn                  high
  101110         13vpn-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             13vpn                  low
  010             12vpn                  low
  011             12vpn                  high
  101             13vpn-priority         low
```

```
[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
  13vpn                 low                    001000
  12vpn                 low                    001010
  12vpn                 high                   001011
  13vpn-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
  13vpn                 low                    001
  12vpn                 low                    010
  12vpn                 high                   011
  13vpn-priority        low                    101
```

```
[edit]
lab@R4# run show class-of-service forwarding-class
Forwarding class          ID          Queue
Policing priority      SPU priority
  be                      0
0      normal            low
  13vpn                  1
1      normal            low
  12vpn                  2
2      normal            low
  13vpn-priority         3
3      normal            low
  nc                      4
4      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

  Scheduler: 13vpn-sc, Forwarding class: 13vpn, 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: 12vpn-sc, Forwarding class: 12vpn, 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: 139
Queues supported: 8, Queues in use: 5
  Scheduler map: core-interfaces, Index: 58651
  Congestion-notification: Disabled

  Logical interface: ge-0/0/4.146, Index: 74
    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-0/0/4.32767, Index: 75

```

```

[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
    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			
	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			

Chapter Ten solutions: Full Day Lab Challenge II


This lab scenario represents a complete 8 hour challenge, aiming to prepare potential candidate for the JNCIE-SP lab exam. It contains all major topics found in the exam blueprint.

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.



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.

Part 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.

Solution - Task 1.1 Service Configuration

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

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.

- a. All devices

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.

- a. All devices

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
```

Solution - Task 1.2: Centralized authentication management

- 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.

- a. All devices

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.

```
[edit]
lab@R1# set system radius-server 10.10.10.1 secret
"$9$uTx31SesYoZDkNdwgJDmPuOlhyK"

[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

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.

- a. All devices

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
```

Solution - Task 1.3: Local user configuration

- 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

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

[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 week days 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

```
[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

JUNOS comes with four predefined login classes: *operator*, *read-only*, *super-user* and *unauthorized*. To accomplish the Task's requirements, you have to 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

The system login message is displayed before the user login. Unlike the system login announcement which is displayed after the user login.

```
[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: Sun 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

```

Solution - Task 1.4: Active configuration archival and logging

- 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

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

By default, JUNOS displays only the emergency messages to the user terminal. Users logged in with the *support* account, have to 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 10MB each.

- a. All devices

```

[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 20MB.

a. All devices

By default, JUNOS rotates the log file when file size reaches 10m. The archive files are stored on the local storage. However, the task requires increasing the size to 20m 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

[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-"
```

Solution - Task 1.5: Advanced Interface configuration and chassis features

- 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 aggregated-ether-options minimum-links 2
```



```
[edit]
lab@R3# set interfaces ae0 aggregated-ether-options lacp active
```

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 aggregated-ether-options minimum-links 2

[edit]
lab@R4# set interfaces ae0 aggregated-ether-options lacp active
```

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
```

```
[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

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 particular 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
    b. Verify the configuration
```

To display the configuration inherit by the apply-group configuration the command below has to 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.

a. All devices

Graceful-restart is disabled by default in JUNOS. To enable it globally, you have to include graceful-restart under the routing-options stanza. This will enable the feature for every routing protocol that supports it. However, you have to 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
```

Solution - Task 1.6: Advanced RE Protection

18) You must ensure that the control plane of all routers are protected. Configure a firewall filter using the following criteria:

Last task is the most time consuming since you have to configure firewall filters with many terms. In addition, this is the first place where you have to 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 subsequent tasks: SSH, BFD, OSPF, RIP, PIM, LDP, RSVP, RADIUS, ping and traceroute.

a. All devices

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

[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

[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
```

```
[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

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]
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

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 have to 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 addresses. Like the configuration groups they also use wildcards to match part of the configuration representing free text.

```
[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 subsequences 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
```

```
[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
```

```
[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
```

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 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

Silently discarding the traffic requires the use of the `discard` action instead of `reject`. To log and store the logs locally you have to 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
```

Appendix - Part 2: Troubleshooting and Configuring IGP

This part is focused on troubleshooting a partially functional IGP network, build with the ISIS and OSPFv3 routing protocols. 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.

Solution - Task 2.1: Troubleshooting

The first task requires fixing the partial working IGP configuration. Using the diagram provided, you have to make sure that all adjacencies are established and routes are exchanged successfully.

- 1) Ensure all IS-IS adjacencies are in 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

On a first glance it seems all adjacencies on router R2 are established. However, the ISIS 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
```

- b. IS-IS adjacency state on R5

Router R5 has 4 interfaces in ISIS which corresponds to the diagram, but only 3 of the adjacency are established successfully. Again, the level of the adjacencies is different then it should be according the diagram. Level 3 suggest 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 ISIS 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

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
ge-0/0/4.35    R5          2 Up          7 2:6:a:e:5:4
ge-0/0/4.37    0172.0020.2507 2 Initializing 24 2:6:a:e:7:4

[edit]
lab@R3# run show isis interface
IS-IS interface database:
Interface      L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
ge-0/0/4.35    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. However, only some of them are successfully established. You may notice also 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
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
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

Router R5 has 3 interfaces configured in ISIS which corresponds to the diagram, but only 2 of the adjacency are established successfully. Similar like 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 ISIS 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
```

Lets 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 ISIS will force all adjacencies to establish Level 1 connections only.

```
[edit]
lab@R2# set protocols isis level 2 disable

[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 ISIS 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**

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 words 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**

Even correcting the missing authentication, the adjacencies are still not established.

[edit]

```
lab@R2# run show isis adjacency
```

Next step is troubleshooting the R2 ISIS adjacencies by enabling traceoptions.

```
[edit]
lab@R2# set protocols isis traceoptions file ISIS-log

[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 ISIS which is defined by the iso address configured under the Loopback interface family iso stanza. 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 ISIS adjacency.

```
[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
```

i. Fixing the adjacency on R5

Star configuring the interfaces with the correct ISIS 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 ISIS 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
commit complete
```

The file captures an error, that ISIS 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

j. Fixing the adjacency on R6

Similar like R5, configure only the required ISIS Levels for each interface on R6 as shown in the diagram.

```
[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# 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 issues with the adjacencies to routers R4 and R7 are solved in the next steps.

```
[edit]
lab@R3# set protocols isis level 1 disable

[edit]
lab@R3# run show isis adjacency
```

Interface	System	L	State	Hold (secs)	SNPA
ae0.0	R4	2	Down	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.

```
[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$0I1OIhrWLNyGJevgJGDmPBIE"; ## SECRET-DATA
        hello-authentication-type md5;
    }
}
interface ge-0/0/4.48;

[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

The adjacency table on router R7 provided an indication of the problem with the ISIS 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
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 ISIS protocol can be configured under the ISIS 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 ISIS 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.fff-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)
  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 Buffer size TLV #14, length: 2
    LSP Buffer size: 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# 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

```

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 ISIS 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 ISIS protocol and are assigned with family iso. Next, you have to check if the Loopback interface is configured with a system ID.

```

[edit]
lab@R8# show interfaces lo0
unit 0 {
    family inet {
        address 172.20.25.8/32;
    }
    family inet6 {
        address 2001:fedc::8/128;
    }
}

```

ISIS 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

- c. Interfaces between area 49.0001 and 49.0005 should not elect DIS.

Similar to OSPF routing protocol, ISIS protocol uses a method to limit the number of adjacencies and the traffic in broadcast segments. ISIS elects one of the routes as Designated Intermediate System (DIS). Configuring point-to-point on the broadcast interface in ISIS, 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
```

- 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
```

- 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
```

- d. R6

```
[edit]
lab@R6# set protocols isis interface ge-0/0/4.46 point-to-point
```

- e. Verify the configuration

The adjacency table on router R5 is missing the SNPA on the interfaces configured as point-to-point. This is 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 ISIS protocol, you still have to explicitly advertise the IPv4 and IPv6 Loopback addresses. If you configure the Loopback interface under ISIS protocol stanza directly, the Loopback address is advertised as internal ISIS route. Rather than using a routing policy to advertise the route, then it will appear as external ISIS route.

a. All IS-IS routers except R2

Passive option prevents the ISIS protocol to try bringing up adjacency over the interface.

```
[edit]
lab@R3# set protocols isis interface lo0.0 passive
```

b. R2

The requirement to advertise R2 Loopback needs the use of routing policy applied to the ISIS protocol. Due to the fact that ISIS support IPv6 routing by default, adding the Loopback interface under ISIS 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 ISIS 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
```

c. Verify the configuration

IPv4 and IPv6 addresses are advertised in ISIS to all routers.

```
[edit]
lab@R2# run show route 172.20.25.0/28 table inet.0

inet.0: 533 destinations, 1008 routes (533 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.0/28      *[IS-IS/18] 00:15:37, metric 30
                   > to 172.20.20.82 via ge-0/0/4.62
                   [BGP/170] 00:57:33, localpref 100, from 172.20.25.5
                   AS path: (65001) I, validation-state: unverified
                   > to 172.20.20.82 via ge-0/0/4.62, Push 299780
                   [BGP/170] 00:55:18, localpref 100, from 172.20.25.6
                   AS path: (65001) I, validation-state: unverified
                   > to 172.20.20.82 via ge-0/0/4.62, Push 299780
172.20.25.1/32     *[IS-IS/160] 00:15:37, metric 11, tag2 10
                   > to 172.20.20.82 via ge-0/0/4.62
172.20.25.2/32     *[Direct/0] 05:27:43
                   > via lo0.0
172.20.25.3/32     *[IS-IS/18] 00:15:37, metric 30
                   > to 172.20.20.82 via ge-0/0/4.62
172.20.25.4/32     *[IS-IS/18] 00:15:37, metric 20
                   > to 172.20.20.82 via ge-0/0/4.62
172.20.25.5/32     *[IS-IS/15] 00:15:37, metric 20
                   > to 172.20.20.82 via ge-0/0/4.62
172.20.25.6/32     *[IS-IS/15] 00:15:37, metric 10
                   > to 172.20.20.82 via ge-0/0/4.62
```

```
172.20.25.7/32      *[IS-IS/18] 00:15:37, metric 30
                   > to 172.20.20.82 via ge-0/0/4.62
172.20.25.8/32      *[IS-IS/18] 00:15:37, metric 30
                   > to 172.20.20.82 via ge-0/0/4.62
```

[edit]

```
lab@R3# run show route 172.20.25.0/28 table inet.0
```

```
inet.0: 536 destinations, 552 routes (535 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
```

```
172.20.25.0/28      *[Aggregate/130] 01:17:20
                   Reject
172.20.25.1/32      *[IS-IS/165] 00:35:00, metric 11, tag2 10
                   > to 172.20.20.14 via ge-0/0/4.35
172.20.25.2/32      *[IS-IS/165] 00:15:24, metric 40
                   > to 172.20.20.14 via ge-0/0/4.35
                   to 172.20.20.2 via ae0.0
172.20.25.3/32      *[Direct/0] 05:23:27
                   > via lo0.0
172.20.25.4/32      *[IS-IS/18] 00:54:35, metric 10
                   > to 172.20.20.2 via ae0.0
172.20.25.5/32      *[IS-IS/18] 00:57:53, metric 10
                   > to 172.20.20.14 via ge-0/0/4.35
172.20.25.6/32      *[IS-IS/165] 00:15:25, metric 10
                   > to 172.20.20.14 via ge-0/0/4.35
172.20.25.7/32      *[IS-IS/18] 00:54:35, metric 20
                   > to 172.20.20.2 via ae0.0
172.20.25.8/32      *[IS-IS/18] 00:54:35, metric 20
                   > to 172.20.20.2 via ae0.0
```

[edit]

```
lab@R2# run show route 2001:fedc::/64
```

```
inet6.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
```

```
2001:fedc::/64      *[IS-IS/165] 00:45:22, metric 30
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2001:fedc::1/128    *[IS-IS/160] 00:45:10, metric 11, tag2 10
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2001:fedc::2/128    *[Direct/0] 00:45:34
                   > via lo0.0
2001:fedc::3/128    *[IS-IS/18] 00:30:03, metric 20
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2001:fedc::4/128    *[IS-IS/18] 00:45:22, metric 20
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2001:fedc::5/128    *[IS-IS/15] 00:45:22, metric 10
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2001:fedc::6/128    *[IS-IS/15] 00:45:22, metric 20
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2001:fedc::7/128    *[IS-IS/18] 00:09:52, metric 30
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2001:fedc::8/128    *[IS-IS/18] 00:45:22, metric 30
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
```

[edit]

```
lab@R3# run show route 2001:fedc::/64
```

```
inet6.0: 148 destinations, 148 routes (148 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
```

```
2001:fedc::/64      *[Aggregate/130] 00:31:27
                   Reject
```

```

2001:fedc::1/128 * [IS-IS/165] 00:31:27, metric 11, tag2 10
                  > to fe80::206:a00:230e:504 via ge-0/0/4.35
2001:fedc::2/128 * [IS-IS/165] 00:31:27, metric 30
                  > to fe80::206:a00:230e:504 via ge-0/0/4.35
2001:fedc::3/128 * [Direct/0] 00:31:27
                  > via lo0.0
2001:fedc::4/128 * [IS-IS/18] 00:02:26, metric 10
                  > to fe80::205:86ff:fea7:dec1 via ae0.0
2001:fedc::5/128 * [IS-IS/18] 00:31:27, metric 10
                  > to fe80::206:a00:230e:504 via ge-0/0/4.35
2001:fedc::6/128 * [IS-IS/18] 00:02:26, metric 20
                  > to fe80::206:a00:230e:504 via ge-0/0/4.35
                  > to fe80::205:86ff:fea7:dec1 via ae0.0
2001:fedc::7/128 * [IS-IS/18] 00:11:30, metric 10
                  > to fe80::206:a00:250e:704 via ge-0/0/4.37
2001:fedc::8/128 * [IS-IS/18] 00:02:26, metric 20
                  > to fe80::206:a00:250e:704 via ge-0/0/4.37
                  > to fe80::205:86ff:fea7:dec1 via ae0.0

```

Solution - Task 2.2: Connectivity to OSPFv3 Area

- 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 level 1 area are mutually reachable.

Router R1 must be connected to the network using the OSPFv3 routing protocol. The Loopback IPv4 and IPv6 addresses of ISIS Level 1 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

```

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 ISIS Level 1 to OSPF.

The routing policy is advertising IPv4 Loopback addresses from ISIS Level 1 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 ISIS 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 ISIS Level 1 Loopback routes to the OSPF domain.

```
[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term isis.level1.routes
from level 1

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv4 term isis.level1.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.level1.routes
then accept

[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.level1.routes
from level 1

[edit]
lab@R5# set policy-options policy-statement to.ospf3.ipv6 term isis.level1.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.level1.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
```

c. R6

Similar configuration is applied to router R6.

```
[edit]
```

```
lab@R6# set protocols ospf3 realm ipv4-unicast export to.ospf3.ipv4
```

```
[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.level1.routes
from level 1
```

```
[edit]
```

```
lab@R6# set policy-options policy-statement to.ospf3.ipv4 term isis.level1.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.level1.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.level1.routes
from level 1
```

```
[edit]
```

```
lab@R6# set policy-options policy-statement to.ospf3.ipv6 term isis.level1.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.level1.routes
then accept
```

```
[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
```

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. Only the ISIS Level 1 routes are being advertised.

```
[edit]
lab@R1# run show route 172.20.25.0/28 table inet.0
```

inet.0: 509 destinations, 1439 routes (504 active, 0 holddown, 55 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

```
172.20.25.0/28      *[BGP/170] 01:11:28, 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-
r4
                   [BGP/170] 01:09:08, 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-
r4
172.20.25.1/32     *[Direct/0] 05:51:35
                   > via lo0.0
172.20.25.2/32     *[OSPF3/150] 01:30:47, metric 20, tag 0
                   > to 172.20.20.70 via ge-0/0/4.16
172.20.25.5/32     *[OSPF3/150] 01:31:48, metric 0, tag 0
                   > to 172.20.20.66 via ge-0/0/4.15
172.20.25.6/32     *[OSPF3/150] 01:30:47, metric 0, tag 0
                   > 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: 547 destinations, 1959 routes (542 active, 0 holddown, 5 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 ISIS 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

inet.0: 547 destinations, 1959 routes (542 active, 0 holddown, 5 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: 541 destinations, 1496 routes (541 active, 0 holddown, 0 hidden)
Restart Complete
+ = 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: 541 destinations, 1496 routes (541 active, 0 holddown, 0 hidden)
Restart Complete
+ = 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: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

2001:fedc::1/128    *[Direct/0] 00:57:04
                   > via lo0.0
2001:fedc::2/128    *[OSPF3/150] 00:56:53, metric 20, tag 0
                   > to fe80::206:a00:f0e:504 via ge-0/0/4.15
2001:fedc::5/128    *[OSPF3/150] 00:56:53, metric 0, tag 0
                   > to fe80::206:a00:f0e:504 via ge-0/0/4.15
2001:fedc::6/128    *[OSPF3/150] 00:56:13, metric 0, tag 0
                   > to fe80::206:a00:100e:604 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 need to 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]
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
```

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
```

```
[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
```

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: 509 destinations, 1439 routes (504 active, 0 holddown, 55 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF3/150] 01:13:17, 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.

```
[edit]
lab@R1# run show route table inet6.0

inet6.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

::/0              *[OSPF3/150] 00:58:26, metric 0, tag 0
                   > to fe80::206:a00:100e:604 via ge-0/0/4.16
                   to fe80::206:a00:f0e:504 via ge-0/0/4.15
```

Solution - Task 2.3: RIP redistribution

- 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* stanza on R7 and R8 to accomplish this task.

a. R3

Next, you have to advertise a summary routes of IPv4 and IPv6 Loopback addresses to the DC-1 router. However, theTask requires those summaries to be advertised only if routers R7 and R8 have connectivity to any of routers R3 and R4. You are not allowed amending the *routing-options* section on routers R7 and R8. This means you have to ensure 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

[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

```

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]
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
```

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
```

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 RIPng 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

[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
```

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

[edit]
lab@R7# set protocols ripng group DC-1 preference 17
```

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]
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 protocols rip group DC-1 preference 17

[edit]
lab@R8# set protocols ripng group DC-1 preference 17

```

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

```

Neighbor	Local State	Source Address	Destination Address	Send Mode	Receive Mode	In 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: 58 destinations, 58 routes (57 active, 1 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
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 RIPng neighbor is connected successfully and routes are received by routers R7 and R8.

```

[edit]
lab@R7# run show route protocol ripng

inet6.0: 146 destinations, 147 routes (146 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
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

```

The IPv4 and IPv6 summary routes are advertised to RIP and RIPng protocols.

```

[edit]
lab@R7# run show route advertising-protocol rip 172.20.201.7

inet.0: 46 destinations, 46 routes (46 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

172.20.25.0/28      *[IS-IS/18] 00:00:11, metric 20
                  > to 172.20.20.9 via ge-0/0/4.47

[edit]
lab@R7# run show route advertising-protocol ripng fe80::206:a00:c90e:705

inet6.0: 146 destinations, 147 routes (146 active, 0 holddown, 0 hidden)
Restart Complete

```

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

```
2001:fedc::/64      *[IS-IS/165] 00:33:51, metric 20
                    > to fe80::206:a00:250e:304 via ge-0/0/4.37
                    [BGP/170] 00:32:36, from 172.20.25.10
                    AS path: I, validation-state: unverified
                    > to fe80::206:a00:250e:304 via ge-0/0/4.37
```

[edit]

```
lab@R7# run show route advertising-protocol ripng fe80::206:a00:c90e:705
```

```
inet6.0: 146 destinations, 147 routes (146 active, 0 holddown, 0 hidden)
```

Restart Complete

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

```
2001:fedc::/64      *[IS-IS/165] 00:33:51, metric 20
                    > to fe80::206:a00:250e:304 via ge-0/0/4.37
                    [BGP/170] 00:32:36, from 172.20.25.10
                    AS path: I, validation-state: unverified
                    > to fe80::206:a00:250e:304 via ge-0/0/4.37
```

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 ISIS routers. However, those routes are missing from the ISIS Level 1 routers.

[edit]

```
lab@R5# run show route 10.15.7.0/24
```

b. R3 and R4

The reason for this is the disabled IPv4 ISIS routing on routers R3 and R4 in the initial configuration. Deleting the command and reverting to the default ISIS operation solves the Task.

[edit]

```
lab@R3# delete protocols isis no-ipv4-routing
```

[edit]

```
lab@R3# commit
commit complete
```

c. Verify the configuration

After committing the change, the IPv4 and IPv6 RIP routes are sent to ISIS Level 1 routers.

[edit]

```
lab@R5# run show route 10.15.7.0/24
```

```
inet.0: 546 destinations, 976 routes (541 active, 0 holddown, 5 hidden)
```

Restart Complete

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

```
10.15.7.0/24      *[IS-IS/18] 00:00:03, metric 22, tag 100
                  > to 172.20.20.17 via ge-0/0/4.45
```

[edit]

```
lab@R5# run show route 2015:4321::/64
```

```
inet6.0: 146 destinations, 173 routes (146 active, 0 holddown, 0 hidden)
```

Restart Complete

<http://www.inetzero.com> - Copyright 2017 iNET ZERO. All rights reserved

For personal non commercial use only – do not distribute

<https://t.me/learningnets>

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

```
2015:4321::/64      *[IS-IS/165] 00:26:44, metric 30
                    > to fe80::206:a00:230e:304 via ge-0/0/4.35
```

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 ISIS. To do that you have to use routing policy for the ISIS protocol.

```
[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 metric
2000
```

```
[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 metric
2000
```

```
[edit]
lab@R8# set policy-options policy-statement to.isis term ripng.routes then accept
```

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: 558 destinations, 575 routes (553 active, 0 holddown, 5 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
```

```
172.20.25.0/28      *[IS-IS/18] 01:31:35, metric 20
                    > to 172.20.20.25 via ge-0/0/4.48
```

Solution - Task 2.4: Multilevel IS-IS configuration

- 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 ISIS Level 1 and Level 2.

An ISIS 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 ISIS 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

[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 ISIS Level 2 routes. Since the RIP routes are advertised in ISIS by routers R7 and R8, they become ISIS 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
```

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]
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
```

```
[edit]
```

```
lab@R6# set protocols isis export to.isis
```

c. R2

You have to 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
```

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: 533 destinations, 1008 routes (533 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

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

```
10.15.0.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.1.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.2.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.3.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.4.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.5.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.6.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.7.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.8.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.9.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.10.0/24     *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.11.0/24     *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.12.0/24     *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.13.0/24     *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.14.0/24     *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.15.0/24     *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > 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] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:1::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:2::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:3::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:4::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:5::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:6::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:7::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:8::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:9::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:a::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:b::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:c::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:d::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:e::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:f::/64   *[IS-IS/165] 01:31:19, metric 40
                   > to fe80::206:a00:340e:504 via ge-0/0/4.52
```

```
[edit]
lab@R2# run show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
R2.00-00              0xb 0xd659 1098 L1
R5.00-00              0x15 0x765b 1101 L1 L2 Attached
R5.02-00              0x8 0xb598 857 L1 L2
R6.00-00              0x13 0x76f1 1155 L1 L2 Attached
R6.00-01              0xf 0xbb2d 368 L1 L2
R6.02-00              0xa 0xf24e 756 L1 L2
R6.03-00              0x9 0xd474 1120 L1 L2
  7 LSPs

IS-IS level 2 link-state database:
  0 LSPs
```

Solution - Task 2.5: Advanced IS-IS configuration

- 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 ISIS provides routing for IPv4 and IPv6 protocols by default, both families use a single SPF algorithm to calculate paths in the network. However, you can configure ISIS to calculate an alternate topology per family. Multi topology 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
```

b. R5

```
[edit]
lab@R5# set protocols isis topologies ipv6-unicast

[edit]
lab@R5# set protocols isis interface ge-0/0/4.52 level 1 metric 1000
```

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
```

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
```

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
```

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
```

g. 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: 533 destinations, 1008 routes (533 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.15.0.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.1.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.2.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > to 172.20.20.82 via ge-0/0/4.62
10.15.3.0/24      *[IS-IS/18] 00:01:33, metric 32, tag 100
                  > 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] 01:31:19, metric 40
                  > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:1::/64  *[IS-IS/165] 01:31:19, metric 40
                  > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:2::/64  *[IS-IS/165] 01:31:19, metric 40
                  > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:3::/64  *[IS-IS/165] 01:31:19, metric 40
                  > to fe80::206:a00:340e:504 via ge-0/0/4.52
2015:4321:4::/64  *[IS-IS/165] 01:31:19, metric 40
                  > to fe80::206:a00:340e:504 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 ISIS metric or cost can have value up to 63 with a combined path cost of 1023. If that is not sufficient, ISIS can be configured to support much wider cost values. All routers that are part of the ISIS Level 2 are configured for wide metrics only.

a. R3

```
[edit]
lab@R3# set protocols isis level 2 wide-metrics-only
```

b. R4

```
[edit]
lab@R4# set protocols isis level 2 wide-metrics-only
```

c. R7

```
[edit]
lab@R7# set protocols isis level 2 wide-metrics-only
```

d. R8

```
[edit]
lab@R8# set protocols isis level 2 wide-metrics-only
```

e. Verify the configuration

The output below shows that router R5 calculates a combined metric value above 1023.

```
[edit]
lab@R5# run show route 10.15.7.0/24

inet.0: 547 destinations, 1958 routes (542 active, 0 holddown, 5 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.15.7.0/24          *[IS-IS/18] 00:00:04, metric 2020, tag 100
                     > 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
```

Part 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.

Solution - Task 3.1: iBGP design and configuration

- 1) Configure the necessary connectivity to the dedicated route-reflector device. It already has partial configuration, however you have to ensure that appropriate IGP configuration is in place to provide connectivity to and from the route-reflector.

a. Route-reflector

Starting with the connectivity towards the route-reflector device from routers R3 and R4.

```
[edit]
lab@route-reflector# set protocols isis level 2 authentication-key
"$9$7LdwgmPQn6AGD.5Qn9CvWLXdsHk.Tz6Yg39ApRELxNV24"

[edit]
lab@route-reflector# set protocols isis level 2 authentication-type simple

[edit]
lab@route-reflector# set protocols isis level 2 wide-metrics-only
```

b. R3

```
[edit]
lab@R3# set protocols isis interface ge-0/0/4.39
```

c. R4

```
[edit]
lab@R4# set protocols isis interface ge-0/0/4.49
```

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-AS'es 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 ISIS 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]
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
```

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]
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
```

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
```

```
[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
```

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

[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

```

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]
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
```

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]
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
```

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]
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
```

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
```

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
```

j. Verify the configuration

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: 3 Peers: 4 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 505 489 0 0 0 0
inet6.0 114 114 0 0 0 0
bgp.l3vpn.0 40 30 0 0 0 0
bgp.l2vpn.0 7 6 0 0 0 0
bgp.mvpn.0 5 3 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.25.5 65002 599 1699 0 1 28:57
Establ
inet.0: 16/16/16/0
inet6.0: 88/88/88/0
bgp.l3vpn.0: 10/10/10/0
bgp.l2vpn.0: 2/2/2/0
bgp.mvpn.0: 2/2/2/0
VPNA.inet.0: 10/10/10/0
VPNA.mvpn.0: 2/2/2/0
172.20.25.6 65002 486 2220 0 1 28:59
Establ
inet.0: 0/16/16/0
inet6.0: 5/5/5/0
bgp.l3vpn.0: 0/9/9/0
bgp.l2vpn.0: 0/1/1/0
bgp.mvpn.0: 0/2/2/0
VPNA.inet.0: 0/9/9/0
VPNA.mvpn.0: 0/2/2/0
172.20.25.10 65001 1426 1636 0 1 28:18
Establ
inet.0: 457/457/457/0
inet6.0: 5/5/5/0
bgp.l3vpn.0: 20/21/21/0
bgp.l2vpn.0: 4/4/4/0
bgp.mvpn.0: 1/1/1/0
VPNA.inet.0: 20/21/21/0
```

VPNA.mvpn.0: 1/1/1/0

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: 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
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...
172.20.25.1      65002      1266    2072      0      0    1:52:03
Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 1/9/9/0
  bgp.l2vpn.0: 1/1/1/0
  bgp.mvpn.0: 1/1/1/0
  VPNA.inet.0: 1/9/9/0
  VPNA.mvpn.0: 1/1/1/0
172.20.25.2      65002      309    2071      0      0    1:51:59
Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 0/0/0/0
  bgp.l2vpn.0: 1/1/1/0
  bgp.mvpn.0: 0/0/0/0
172.20.25.3      65001     1727      617      0      1    31:13
Establ
  inet.0: 473/474/474/0
  bgp.l3vpn.0: 30/31/31/0
  bgp.l2vpn.0: 4/4/4/0
  bgp.mvpn.0: 2/2/2/0
  inet6.0: 21/22/22/0
  VPNA.inet.0: 30/31/31/0
  VPNA.mvpn.0: 2/2/2/0
172.20.25.4      65001      573      940      0      1    30:47
Establ
  inet.0: 0/474/474/0
  bgp.l3vpn.0: 1/31/31/0
  bgp.l2vpn.0: 0/4/4/0
  bgp.mvpn.0: 0/2/2/0
  inet6.0: 0/22/22/0
  VPNA.inet.0: 1/31/31/0
  VPNA.mvpn.0: 0/2/2/0
172.20.25.6      65002     1466    2064      0      0    1:51:14
Establ
  inet.0: 16/490/490/0
  bgp.l3vpn.0: 0/31/31/0
  bgp.l2vpn.0: 0/4/4/0
  bgp.mvpn.0: 0/2/2/0
  VPNA.inet.0: 0/31/31/0
  VPNA.mvpn.0: 0/2/2/0
```

```
[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
```

```

          964          489          0          0          0          0
bgp.13vpn.0
          0          0          0          0          0          0
bgp.12vpn.0
          6          3          0          0          0          0
bgp.mvpn.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      1397      441          0          0          3:18:14
Establ
  inet.0: 0/0/0/0
  bgp.13vpn.0: 0/0/0/0
  bgp.12vpn.0: 0/0/0/0
  bgp.mvpn.0: 0/0/0/0
  VPLS.12vpn.0: 0/0/0/0
172.20.25.5      65002      2174      439          0          0          3:18:11
Establ
  inet.0: 473/474/474/0
  bgp.13vpn.0: 0/0/0/0
  bgp.12vpn.0: 3/3/3/0
  bgp.mvpn.0: 0/0/0/0
  VPLS.12vpn.0: 3/3/3/0
172.20.25.6      65002      1558      437          0          0          3:17:14
Establ
  inet.0: 16/490/490/0
  bgp.13vpn.0: 0/0/0/0
  bgp.12vpn.0: 0/3/3/0
  bgp.mvpn.0: 0/0/0/0
  VPLS.12vpn.0: 0/3/3/0

```

- 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

```

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
```

c. Route-reflector

Because the router-reflector has the IPv4 summary route learned by ISIS in the routing table, 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.

```
[edit]
```

```
lab@route-reflector# set protocols bgp advertise-inactive
```

d. Verify the configuration

The outputs below shows 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: 547 destinations, 1958 routes (542 active, 0 holddown, 5 hidden)
```

```
Restart Complete
```

Prefix	Nexthop	MED	Lclpref	AS path
172.20.25.0/28	172.20.25.3		100	(65001) I

```
inet.3: 39 destinations, 47 routes (6 active, 0 holddown, 39 hidden)
```

```
Restart Complete
```

```
[edit]
```

```
lab@R5# run show route advertising-protocol bgp 172.20.25.2 172.20.25.0/28
```

```
inet.0: 547 destinations, 1958 routes (542 active, 0 holddown, 5 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 ISIS 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
```

```
inet.0: 528 destinations, 530 routes (528 active, 0 holddown, 0 hidden)
```

```
+ = 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.

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. Since the other routers in the network cannot resolve this address, it will stay hidden and not usable. You have to change the next-hop address to the Loopback IPv6 address already advertised by ISIS.

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
```

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

```

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

```

d. R8

```

[edit]
lab@R8# set protocols bgp group internal family inet unicast

[edit]
lab@R8# set protocols bgp group internal family inet6 unicast

```

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

```

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

```

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

```

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: 146 destinations, 172 routes (146 active, 0 holddown, 0 hidden)
Restart Complete
  Prefix                Nexthop          MED      Lclpref    AS path
  2001:fedc::/64        2001:fedc::3      100       100        (65001) I

inet6.3: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
Restart Complete

[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

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.

```
[edit]
lab@R3# set protocols bgp group internal neighbor 172.20.25.10 passive
```

b. Verify the configuration

```
[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
```

Solution - Task 3.2: eBGP peers and configuration

- 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
```

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
```

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]
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

```

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: 4 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0
                510        504         0           0         0         0
inet6.0
                115        114         0           0         0         0
bgp.l3vpn.0
                 31         30         0           0         0         0
bgp.l2vpn.0
                  2          2         0           0         0         0
bgp.mvpn.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      603     259     0       0      1:54:02
Establ
inet.0: 456/461/456/0

```

```

[edit]
lab@R1# run show bgp summary
Groups: 3 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0
                1425        490         0           0         0         0
bgp.l3vpn.0
                 71         30         0           0         0         0
bgp.l2vpn.0
                  2          1         0           0         0         0
bgp.mvpn.0
                  5          3         0           0         0         0
Peer           AS      InPkt   OutPkt   OutQ   Flaps  Last  Up/Dwn
State|#Active/Received/Accepted/Damped...
144.44.0.101   16151      610     517     0       0      2:39:47
Establ
inet.0: 0/461/406/0

[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
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...
134.45.0.1 4849703 604 1347 0 0 1:58:25
Establ
inet.0: 0/461/456/0
2001:134:45::1 4849703 299 263 0 0 1:58:07
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 knob *accept-remote-nexthop*.

Unlike the connection to U3, where two separate eBGP sessions are used to configure R3 to peer with P1, you have to configure a single eBGP session advertising both IPv4 and IPv6 prefixes.

a. R3

Configure the interface to connect with peer P1 according the diagram.

```

[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 have to 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
```

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
```

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.0 521 489 0 0 0 0
inet6.0 114 114 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.30.0.2 11599906 238 343 0 0 1:56:10
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: 148 destinations, 148 routes (148 active, 0 holddown, 0 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:f0f4:f691:d2::/64
* 2001:f0f4:f691:d3::/64
* 2001:f0f4:f691:d4::/64
* 2001:f0f4:f691:d5::/64
*
2001:172:30::2
2001:172:30::2
2001:172:30::2
2001:172:30::2
2001:172:30::2
2001:172:30::2
11599906 I
11599906 I
11599906 I
11599906 I
11599906 I
11599906 I

```

The output below is taken after the `multihop` knob 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: 148 destinations, 148 routes (148 active, 0 holddown, 0 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
* 2001:f0f4:f691:d5::/64	::ffff:172.30.0.2			11599906 I

- 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 have to 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 have to 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

```

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::206:a01:9a0e:605 --> ff02::1
16 bytes from fe80::206:a01:9a0e:908, icmp_seq=0 hlim=64 time=1.833 ms
16 bytes from fe80::206:a01:9a0e:908, icmp_seq=1 hlim=64 time=1.821 ms
^C
--- ff02::1 ping6 statistics ---
2 packets transmitted, 2 packets received, 0% packet loss

```

```
round-trip min/avg/max/std-dev = 1.821/1.827/1.833/0.006 ms
```

```
[edit]
lab@R6# set protocols bgp group C1-v4 peer-as 65120
```

```
[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 have to specify the local interface. This is needed because any IPv6 enabled interface has a link local address automatically configured and BGP needs to know which to use for the particular 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::206:a01:9a0e:908
```

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
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0	1438	489	0	0	0	0	0
bgp.l3vpn.0	111	40	0	0	0	0	0
bgp.l2vpn.0	14	5	0	0	0	0	0
bgp.mvpn.0	8	4	0	0	0	0	0
inet6.0	49	26	0	0	0	0	0

Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last	Up/Dwn
172.20.100.2	65120	257	1473	0	0	1:55:37	Established
fe80::206:a01:9a0e:908	65120	256	257	0	0	1:55:01	Established

```
inet.0: 16/16/16/0
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 have to discover the eBGP AS and configure the eBGP session using md5 authentication password jncie.

To configure the eBGP sessions to peer with C2 you have to first discover the peer AS number since 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
```

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
```

Looking 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)
```

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
```

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: 4 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
inet.0
                510        504         0           0         0         0
inet6.0
                115        114         0           0         0         0
bgp.l3vpn.0
                31         30         0           0         0         0
bgp.l2vpn.0
                 2         2         0           0         0         0
bgp.mvpn.0
                 0         0         0           0         0         0
Peer           AS      InPkt   OutPkt   OutQ   Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.200.2   65123      3     383      0       0         0         6
Establ
  inet.0: 16/16/16/0
  inet6.0: 5/5/5/0
```

Solution - Task 3.3: BGP policy configuration

This section is mostly based on BGP route manipulation according to the given requirements. This is one of the most important part 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 have to 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 then next-hop self
```

To match 4byte AS numbers in regular expressions you need to 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
```

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
```

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 then next-hop
2001:fedc::5

[edit]
lab@R5# set policy-options policy-statement nhs term v6 then accept

[edit]
lab@R5# set policy-options policy-statement nhs term v4 then next-hop self

[edit]
lab@R5# set protocols bgp group internal export nhs

[edit]
lab@R5# set protocols bgp group sub-AS65001 export nhs
```

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
```

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 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

```

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: 1 Peers: 4 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0
    491      489      0      0      0      0
inet6.0
    116      114      0      0      0      0
bgp.l3vpn.0
    43       41      0      0      0      0
bgp.l2vpn.0
    6        6      0      0      0      0
bgp.mvpn.0
    4        4      0      0      0      0
Peer      AS      InPkt  OutPkt  OutQ  Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.20.25.3 65001 1880    1549    0      1    1:17:09
Establ
inet.0: 32/33/33/0
inet6.0: 109/110/110/0
bgp.l3vpn.0: 20/21/21/0
bgp.l2vpn.0: 2/2/2/0
bgp.mvpn.0: 3/3/3/0
172.20.25.4 65001 261     664     0      1    1:17:07
Establ
inet.0: 0/1/1/0
inet6.0: 0/1/1/0
bgp.l3vpn.0: 0/0/0/0
bgp.l2vpn.0: 1/1/1/0
bgp.mvpn.0: 0/0/0/0
172.20.25.7 65001 167     1644    0      1    1:17:14
Establ
inet.0: 0/0/0/0
inet6.0: 0/0/0/0
bgp.l3vpn.0: 11/12/12/0
bgp.l2vpn.0: 1/1/1/0
bgp.mvpn.0: 1/1/1/0
172.20.25.8 65001 550     1262    0      1    1:17:18
Establ
inet.0: 457/457/457/0
inet6.0: 5/5/5/0
bgp.l3vpn.0: 10/10/10/0
bgp.l2vpn.0: 2/2/2/0
bgp.mvpn.0: 0/0/0/0

```

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.0 964 489 0 0 0 0
bgp.l3vpn.0 0 0 0 0 0 0
bgp.l2vpn.0 6 3 0 0 0 0
bgp.mvpn.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 1397 441 0 0 3:18:14
Establ
inet.0: 0/0/0/0
bgp.l3vpn.0: 0/0/0/0
bgp.l2vpn.0: 0/0/0/0
bgp.mvpn.0: 0/0/0/0
VPLS.l2vpn.0: 0/0/0/0
172.20.25.5 65002 2174 439 0 0 3:18:11
Establ
inet.0: 473/474/474/0
bgp.l3vpn.0: 0/0/0/0
bgp.l2vpn.0: 3/3/3/0
bgp.mvpn.0: 0/0/0/0
VPLS.l2vpn.0: 3/3/3/0
172.20.25.6 65002 1558 437 0 0 3:17:14
Establ
inet.0: 16/490/490/0
bgp.l3vpn.0: 0/0/0/0
bgp.l2vpn.0: 0/3/3/0
bgp.mvpn.0: 0/0/0/0
VPLS.l2vpn.0: 0/3/3/0
```

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. This is because the same route is advertised by the ISIS protocol also. To force BGP advertising the inactive routes you have to use `advertise-inactive` knob.

a. R3

```
[edit]
lab@R3# set protocols bgp advertise-inactive
```

b. R5

```
[edit]
lab@R5# set protocols bgp advertise-inactive
```

c. R6

```
[edit]
lab@R6# set protocols bgp advertise-inactive
```

d. R8

```
[edit]
lab@R8# set protocols bgp advertise-inactive
```

e. 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, 5 hidden)
Restart Complete
  Prefix                Nexthop          MED      Lclpref    AS path
  172.20.25.0/28        Self              0

```

```
[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              0

```

```
[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, 5 hidden)
Restart Complete
  Prefix                Nexthop          MED      Lclpref    AS path
  172.20.25.0/28        Self              0

```

```
[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, 5 hidden)
Restart Complete
  Prefix                Nexthop          MED      Lclpref    AS path
  172.20.25.0/28        Self              0

```

```
[edit]
lab@R1# run show route advertising-protocol bgp 144.44.0.101 172.20.25.0/28

inet.0: 509 destinations, 1439 routes (504 active, 0 holddown, 55 hidden)
Restart Complete
  Prefix                Nexthop          MED      Lclpref    AS path
* 172.20.25.0/28        Self              0

```

13) Router R1 is chosen as backup gateway to AS 16151 therefor 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

```

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

inet.0: 509 destinations, 1439 routes (504 active, 0 holddown, 55 hidden)
Restart Complete
  Prefix                Nexthop          MED      Lclpref   AS path
* 172.20.25.0/28        Self            100              I
```

14) Make sure that router R8 is the preferred exit gateway for the same prefixes received from all Upstream peers. This requirement should not effect R8 routing table. In case the connectivity to U1 is lost, R1 and R5 should prefer their local connections to the upstreams.

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
```

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
```

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: 559 destinations, 576 routes (554 active, 0 holddown, 5 hidden)
Restart Complete
  Prefix                Nexthop          MED      Lclpref   AS path
* 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
* 1.176.0.0/12         Self          200          16151 33112
49129 16320 52954 I

[edit]
lab@R8# run show route

inet.0: 559 destinations, 576 routes (554 active, 0 holddown, 5 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

1.64.0.0/10            *[BGP/170] 02:46:10, localpref 100
                        AS path: 16151 61671 I, validation-state: unverified
                        > to 144.44.0.1 via ge-0/0/5.511
1.84.160.0/20          *[BGP/170] 02:46:10, localpref 100
                        AS path: 16151 33112 I, validation-state: unverified
                        > to 144.44.0.1 via ge-0/0/5.511
1.96.0.0/11            *[BGP/170] 02:46:10, localpref 100
                        AS path: 16151 33112 63164 40776 51777 I, validation-state:
unverified
                        > to 144.44.0.1 via ge-0/0/5.511
1.161.192.0/21         *[BGP/170] 02:46:10, localpref 100
                        AS path: 16151 33112 30404 32138 45045 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

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
```

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@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@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 with using a 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|4849703) .*"

```

The routing policies are applied to the BGP neighbors they are intended to.

```
[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

[edit]
lab@R6# set protocols bgp group sub-AS65001 export to.AS65001

[edit]
lab@R1# set protocols bgp group U2 export to.U2

```

b. R5

You have to 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 have to 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

[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

```

c. Verify the configuration

After applying the routing policies, only the 2 summary routes are advertised to peer C1.

```
[edit]
lab@R6# run show route advertising-protocol bgp 172.20.100.2

inet.0: 541 destinations, 1496 routes (541 active, 0 holddown, 0 hidden)
Restart Complete
   Prefix                Nexthop          MED      Lclpref   AS path
* 101.15.128.0/20        Self              65123 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: 559 destinations, 576 routes (554 active, 0 holddown, 5 hidden)
Restart Complete
101.10.32.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 101.10.32.0/24 -> {indirect(1048576)}
Page 0 idx 0, (group C2 type External) Type 1 val 0xaf40a68 (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: 0
  *BGP      Preference: 170/-101
            Next hop type: Indirect, Next hop index: 0
            Address: 0xb1d9840
            Next-hop reference count: 32
            Source: 172.20.25.10
            Next hop type: Router, Next hop index: 682
            Next hop: 172.20.20.25 via ge-0/0/4.48, selected
            Label operation: Push 299888
            Label TTL action: no-prop-ttl
            Load balance label: Label 299888: None;
            Label element ptr: 0xaf15ac0
            Label parent element ptr: 0x0
            Label element references: 1
            Label element child references: 0
            Label element lsp id: 0
            Session Id: 0x141
            Protocol next hop: 172.20.25.6
            Indirect next hop: 0xa8fe8b0 1048576 INH Session ID: 0x156
            State: <Active Int Ext>
            Local AS: 65001 Peer AS: 65001
            Age: 1:32:34 Metric2: 1
            Validation State: unverified
            Task: BGP_65001.172.20.25.10
            Announcement bits (3): 0-KRT 7-BGP_RT_Background 8-Resolve tree 4
            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: 172.20.25.6 Metric: 1
              Indirect next hop: 0xa8fe8b0 1048576 INH Session ID:
0x156
              Indirect path forwarding next hops: 1
                Next hop type: Router
                Next hop: 172.20.20.25 via ge-0/0/4.48
```

```

Session Id: 0x141
172.20.25.6/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 1
Nexthop: 172.20.20.25 via ge-0/0/4.48

```

- 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.

```

[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]
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

```

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

```

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

```

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              0         0          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      0         0          65123 I
* 101.15.129.0/24      172.20.200.2      0         0          65123 I
* 101.15.130.0/24      172.20.200.2      0         0          65123 I
* 101.15.131.0/24      172.20.200.2      0         0          65123 I
* 101.15.132.0/24      172.20.200.2      0         0          65123 I
* 101.15.133.0/24      172.20.200.2      0         0          65123 I
* 101.15.134.0/24      172.20.200.2      0         0          65123 I
* 101.15.135.0/24      172.20.200.2      0         0          65123 I
* 101.15.136.0/24      172.20.200.2      0         0          65123 I
* 101.15.137.0/24      172.20.200.2      0         0          65123 I
* 101.15.138.0/24      172.20.200.2      0         0          65123 I
* 101.15.139.0/24      172.20.200.2      0         0          65123 I
* 101.15.140.0/24      172.20.200.2      0         0          65123 I
* 101.15.141.0/24      172.20.200.2      0         0          65123 I
* 101.15.142.0/24      172.20.200.2      0         0          65123 I
* 101.15.143.0/24      172.20.200.2      0         0          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  0         0          65123 I
* 2001:10:15:133::/64   ::ffff:172.20.200.2  0         0          65123 I
* 2001:10:15:134::/64   ::ffff:172.20.200.2  0         0          65123 I
* 2001:10:15:135::/64   ::ffff:172.20.200.2  0         0          65123 I
* 2001:10:15:136::/64   ::ffff:172.20.200.2  0         0          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.

```
[edit]
lab@R7# set policy-options policy-statement from.rip term accept.rest then tag 100
```

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

[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 `policy-option` stanza 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
```

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: 535 destinations, 551 routes (535 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
    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: 535 destinations, 551 routes (535 active, 0 holddown, 0 hidden)
Restart Complete
```

Prefix	Nexthop	MED	Localpref	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]
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
```

b. R5

Similar configuration is applied to router R5.

```
[edit]
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 protocols bgp group U3-v4 import from.U3-v4
```

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 protocols bgp group U1 import from.U1
```

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: 559 destinations, 576 routes (554 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 then 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:

- 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
```

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
```

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

[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
```

d. Verify the configuration

The damping profile applied to a particular route can be seen with extensive view of the routing table. After the configuration is committed, depending on the subnet mask size different damping parameters are used.

```
[edit]
lab@R8# run show route 1.84.160.0/20 extensive

1.84.160.0/20 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.84.160.0/20 -> {144.44.0.1}
Page 0 idx 0, (group C2 type External) Type 1 val 0xaf3ab78 (adv_entry)
  Advertised metrics:
    Nexthop: Self
    AS path: [65001] 16151 33112 I
    Communities:
Page 0 idx 2, (group internal type Internal) Type 1 val 0xaf42a04 (adv_entry)
  Advertised metrics:
    Flags: Nexthop Change
```

```

    Nexthop: Self
    Localpref: 200
    AS path: [65001] 16151 33112 I
    Communities:
Path 1.84.160.0 from 144.44.0.1 Vector len 4. Val: 0 2
    *BGP      Preference: 170/-101
              Next hop type: Router, Next hop index: 682
              Address: 0xa6e417c
              Next-hop reference count: 917
              Source: 144.44.0.1
              Next hop: 144.44.0.1 via ge-0/0/5.511, selected
              Session Id: 0x148
              State: <Active Ext>
              Local AS: 65001 Peer AS: 16151
              Age: 17:02
              Validation State: unverified
              Task: BGP_16151.144.44.0.1
              Announcement bits (3): 0-KRT 7-BGP_RT_Background 8-Resolve tree 4
              AS path: 16151 33112 I
              Accepted
              Localpref: 100
              Router ID: 144.44.0.1
              Merit (last update/now): 733/372
              damping-parameters: aggressive
              Last update: 00000000:14:40 First update: 00000000:36:44
              Flaps: 2

```

```

[edit]
lab@R8# run show route 1.96.0.0/11 extensive

1.96.0.0/11 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.96.0.0/11 -> {144.44.0.1}
Page 0 idx 0, (group C2 type External) Type 1 val 0xaf3ad70 (adv_entry)
    Advertised metrics:
        Nexthop: Self
        AS path: [65001] 16151 33112 63164 40776 51777 I
        Communities:
Page 0 idx 2, (group internal type Internal) Type 1 val 0xaf42a20 (adv_entry)
    Advertised metrics:
        Flags: Nexthop Change
        Nexthop: Self
        Localpref: 200
        AS path: [65001] 16151 33112 63164 40776 51777 I
        Communities:
Path 1.96.0.0 from 144.44.0.1 Vector len 4. Val: 0 2
    *BGP      Preference: 170/-101
              Next hop type: Router, Next hop index: 682
              Address: 0xa6e417c
              Next-hop reference count: 917
              Source: 144.44.0.1
              Next hop: 144.44.0.1 via ge-0/0/5.511, selected
              Session Id: 0x148
              State: <Active Ext>
              Local AS: 65001 Peer AS: 16151
              Age: 17:02
              Validation State: unverified
              Task: BGP_16151.144.44.0.1
              Announcement bits (3): 0-KRT 7-BGP_RT_Background 8-Resolve tree 4
              AS path: 16151 33112 63164 40776 51777 I
              Accepted
              Localpref: 100
              Router ID: 144.44.0.1
              Merit (last update/now): 733/372
              Default damping parameters used
              Last update: 00000000:14:40 First update: 00000000:36:44
              Flaps: 2

```

The output of route with damping disable is shown below.

```
[edit]
lab@R8# run show route 2.0.0.0/7 extensive

inet.0: 559 destinations, 576 routes (554 active, 0 holddown, 5 hidden)
Restart Complete
2.0.0.0/7 (1 entry, 1 announced)
TSI:
KRT in-kernel 2.0.0.0/7 -> {144.44.0.1}
Page 0 idx 0, (group C2 type External) Type 1 val 0xaf3adc4 (adv_entry)
  Advertised metrics:
    Nexthop: Self
    AS path: [65001] 16151 61671 60177 3091 5721 1770 35283 6098 41407 I
    Communities:
Page 0 idx 2, (group internal type Internal) Type 1 val 0xaf42a74 (adv_entry)
  Advertised metrics:
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 200
    AS path: [65001] 16151 61671 60177 3091 5721 1770 35283 6098 41407 I
    Communities:
Path 2.0.0.0 from 144.44.0.1 Vector len 4. Val: 0 2
  *BGP Preference: 170/-101
    Next hop type: Router, Next hop index: 682
    Address: 0xa6e417c
    Next-hop reference count: 917
    Source: 144.44.0.1
    Next hop: 144.44.0.1 via ge-0/0/5.511, selected
    Session Id: 0x148
    State: <Active Ext>
    Local AS: 65001 Peer AS: 16151
    Age: 19:40
    Validation State: unverified
    Task: BGP_16151.144.44.0.1
    Announcement bits (3): 0-KRT 7-BGP_RT_Background 8-Resolve tree 4
    AS path: 16151 61671 60177 3091 5721 1770 35283 6098 41407 I
    Accepted
    Localpref: 100
    Router ID: 144.44.0.1
```

Solution - Task 3.4: BGP general requirements

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 explicit permitted by a particular task.

22) 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: 3 Peers: 5 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0							
	969	490	0	0	0	0	0
bgp.l3vpn.0							
	71	30	0	0	0	0	0
bgp.l2vpn.0							
	2	1	0	0	0	0	0
bgp.mvpn.0							
	5	3	0	0	0	0	0
Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last	Up/Dwn
State #Active/Received/Accepted/Damped...							
10.10.0.14	64888	36	62	0	1	22:43	
Establ							
VPNA.inet.0: 8/10/8/0							
144.44.0.101	16151	412	181	0	1	22:51	
Establ							
inet.0: 406/461/406/0							
172.20.25.2	65002	49	757	0	1	22:32	
Establ							
inet.0: 0/0/0/0							
bgp.l3vpn.0: 0/0/0/0							
bgp.l2vpn.0: 0/0/0/0							
bgp.mvpn.0: 0/0/0/0							
L2VPN.l2vpn.0: 0/0/0/0							
172.20.25.5	65002	1216	758	0	1	22:47	
Establ							
inet.0: 68/474/474/0							
bgp.l3vpn.0: 30/40/40/0							
bgp.l2vpn.0: 1/1/1/0							
bgp.mvpn.0: 3/3/3/0							
VPNA.inet.0: 30/40/40/0							
L2VPN.l2vpn.0: 1/1/1/0							
VPNA.mvpn.0: 3/3/3/0							
172.20.25.6	65002	2153	2309	0	0	4:57:14	
Establ							
inet.0: 16/34/34/0							
bgp.l3vpn.0: 0/31/31/0							
bgp.l2vpn.0: 0/1/1/0							
bgp.mvpn.0: 0/2/2/0							
VPNA.inet.0: 0/31/31/0							
L2VPN.l2vpn.0: 0/1/1/0							
VPNA.mvpn.0: 0/2/2/0							

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: 559 destinations, 576 routes (554 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
```

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: 3 Peers: 5 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 1425 490 0 0 0 0
bgp.l3vpn.0 71 30 0 0 0 0
bgp.l2vpn.0 2 1 0 0 0 0
bgp.mvpn.0 5 3 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.10.0.14 64888 231 377 0 0 2:39:46
Establ
  VPNA.inet.0: 8/10/8/0
144.44.0.101 16151 610 517 0 0 2:39:47
Establ
  inet.0: 0/461/406/0
172.20.25.2 65002 324 1281 0 0 2:32:21
Establ
  inet.0: 0/0/0/0
  bgp.l3vpn.0: 0/0/0/0
  bgp.l2vpn.0: 0/0/0/0
  bgp.mvpn.0: 0/0/0/0
  L2VPN.l2vpn.0: 0/0/0/0
172.20.25.5 65002 2081 1282 0 0 2:32:22
Establ
  inet.0: 474/474/474/0
  bgp.l3vpn.0: 30/40/40/0
  bgp.l2vpn.0: 1/1/1/0
  bgp.mvpn.0: 3/3/3/0
  VPNA.inet.0: 30/40/40/0
  L2VPN.l2vpn.0: 1/1/1/0
  VPNA.mvpn.0: 3/3/3/0
```

```

172.20.25.6          65002          1475          1278          0          0          2:31:12
Establ
  inet.0: 16/490/490/0
  bgp.l3vpn.0: 0/31/31/0
  bgp.l2vpn.0: 0/1/1/0
  bgp.mvpn.0: 0/2/2/0
  VPNA.inet.0: 0/31/31/0
  L2VPN.l2vpn.0: 0/1/1/0
  VPNA.mvpn.0: 0/2/2/0

```

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 changing 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

inet.0: 535 destinations, 551 routes (535 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(1048585)}
Page 0 idx 1, (group sub-AS65002 type External) Type 1 val 0xb38f3e8 (adv_entry)
  Advertised metrics:
    Nexthop: 172.20.25.10
    Localpref: 200
    AS path: [65001] 16151 61671 I
    Communities:
Path 1.64.0.0 from 172.20.25.10 Vector len 4. Val: 1
  *BGP      Preference: 170/-201
            Next hop type: Indirect, Next hop index: 0
            Address: 0xa6ef354
            Next-hop reference count: 914
            Source: 172.20.25.10
            Next hop type: Router, Next hop index: 694
            Next hop: 172.20.20.34 via ge-0/0/4.39, selected
            Session Id: 0x14e
            Protocol next hop: 172.20.25.10
            Indirect next hop: 0xa901000 1048585 INH Session ID: 0x16a
            State: <Active Int Ext>
            Local AS: 65001 Peer AS: 65001
            Age: 20 Metric2: 10
            Validation State: unverified
            ORR Generation-ID: 0
            Task: BGP_65001.172.20.25.10
            Announcement bits (3): 0-KRT 6-BGP_RT_Background 7-Resolve tree 4
            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: 0xa901000 1048585 INH Session ID:
0x16a

            Indirect path forwarding next hops: 1
              Next hop type: Router
              Next hop: 172.20.20.34 via ge-0/0/4.39
              Session Id: 0x14e
              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
```

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.

```
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: 535 destinations, 551 routes (535 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(1048586)}
Page 0 idx 1, (group sub-AS65002 type External) Type 1 val 0xb38f5a8 (adv_entry)
  Advertised metrics:
    Nexthop: 172.20.25.8
    Localpref: 200
    AS path: [65001] 16151 61671 I
    Communities:
Path 1.64.0.0 from 172.20.25.10 Vector len 4. Val: 1
  *BGP      Preference: 170/-201
            Next hop type: Indirect, Next hop index: 0
            Address: 0xa6eccf0
            Next-hop reference count: 914
            Source: 172.20.25.10
            Next hop type: Router, Next hop index: 704
            Next hop: 172.20.20.2 via ae0.0, selected
            Label-switched-path r3-to-r4
            Label operation: Push 300305
            Label TTL action: no-prop-ttl
            Load balance label: Label 300305: None;
            Label element ptr: 0xaf15be0
            Label parent element ptr: 0xaf159a0
```

```

Label element references: 2
Label element child references: 1
Label element lsp id: 0
Session Id: 0x147
Protocol next hop: 172.20.25.8
Indirect next hop: 0xa901000 1048586 INH Session ID: 0x16c
State: <Active Int Ext>
Local AS: 65001 Peer AS: 65001
Age: 11 Metric2: 1
Validation State: unverified
Task: BGP_65001.172.20.25.10
Announcement bits (3): 0-KRT 6-BGP_RT_Background 7-Resolve tree 4
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: 1
    Indirect next hop: 0xa901000 1048586 INH Session ID:
        Indirect path forwarding next hops: 1
            Next hop type: Router
            Next hop: 172.20.20.2 via ae0.0
            Session Id: 0x147
172.20.25.8/32 Originating RIB: inet.3
    Metric: 1 Node path count: 1
    Forwarding nexthops: 1
        Nexthop: 172.20.20.2 via ae0.0

```

0x16c

Part 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.

Solution - Task 4.1: MPLS and RSVP 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.
 - a. All devices except R2 and R8

Enabling RSVP protocol is straightforward, simply configuring the interfaces under the `protocols rsvp` stanza. In addition, you have to enable the family `mpls` for the interface you have add the interface under the `protocol mpls` stanza 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
```

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
```

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
```

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

```

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
  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%
  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

```

RSVP neighbor table is displayed on routers R1 and R5. However, the neighbor sessions will be established when a LSP is signaled.

```

[edit]
lab@R1# run show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd

```

```

172.20.20.70      0 1/0      2:27:14      9   975/974   461
172.20.20.66      0 1/0      2:27:54      9   979/979   1765

```

```

[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 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/8xdDjqf5F24UHqfQzcycle87"

```

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"

```

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 ISIS and stored in the TED table.

a. R3

First, you have to 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

```

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

[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
```

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
```

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
```

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
```

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
               Up         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>
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, ISIS 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   194     2     0 IS-IS(1)
R3.00(172.20.25.3) Rtr   468     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: 334, Remote interface index: 333
  To: R4.02, Local: 172.20.20.1, Remote: 0.0.0.0
    Local interface index: 341, Remote interface index: 0
  To: R7.03, Local: 172.20.20.5, Remote: 0.0.0.0
    Local interface index: 335, Remote interface index: 0
ID              Type Age(s) LnkIn LnkOut Protocol
R3.02           Net    12     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    4     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: 334, Remote interface index: 334
  To: R4.02, Local: 172.20.20.2, Remote: 0.0.0.0
    Local interface index: 341, Remote interface index: 0
```

```

To: R6.00(172.20.25.6), Local: 172.20.20.21, Remote: 172.20.20.22
  Local interface index: 335, Remote interface index: 333
To: R7.02, Local: 172.20.20.9, Remote: 0.0.0.0
  Local interface index: 336, Remote interface index: 0
ID                                     Type Age(s) LnkIn LnkOut Protocol
R4.02                                 Net   368     2     2 IS-IS(2)
  To: R4.00(172.20.25.4), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  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
ID                                     Type Age(s) LnkIn LnkOut Protocol
R4.03                                 Net   368     0     2 IS-IS(2)
  To: R4.00(172.20.25.4), 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
R5.00(172.20.25.5)                   Rtr   195     4     2 IS-IS(2)
  To: R4.00(172.20.25.4), Local: 172.20.20.18, Remote: 172.20.20.17
  Local interface index: 334, Remote interface index: 334
  To: R3.00(172.20.25.3), Local: 172.20.20.14, Remote: 172.20.20.13

```

4) Configure the following LSPs using the criteria below:

- a. 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.

a. R1

First, you have to 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.

```

[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

```

b. 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.

```
[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
```

c. 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

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: Sun 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
  Explt 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
    Detour RESV rcvfrom: 172.20.20.17 (ge-0/0/4.45) 133 pkts, Entropy label: Yes
    Detour Explt 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: Sun 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
  Explt 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 Explt 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 back 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
```

Again, the 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 path 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 back to the previous path. First configure the path that has to 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
```

- b. R7

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
```

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
  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: Sun 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

[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

```

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.

```
[edit]
lab@R5# set protocols mpls path pri-to-r3
```

```
[edit]
lab@R5# set protocols mpls path sec-to-r3
```

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
        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: Sun 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)-
>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]

```

- 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.

d. 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 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
```

e. 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
```

f. 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
  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: Sun 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
  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: Sun 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

```

- 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.

```

[edit]
lab@R1# set protocols mpls optimize-timer 7200

```

- 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
  LSPTtype: 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

```

Solution - Task 4.2: MPLS and LDP configuration

- 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.

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 have to configure a 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 need to 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
```

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
```

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 stanza on all routers

```
[edit]
lab@R5# set protocols ldp interface lo0.0
```

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
```

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
```

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
```

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
```

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

Since some of the required LSPs are already configured earlier, you only need to enable LDP tunneling over them. Add simply ldp-tunneling to achieve this under the particular 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

[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

```

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

```

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

```

d. R5

```

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r3 ldp-tunneling

[edit]
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
```

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
```

- 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
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
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.

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.

48)

- a. All devices

Since, 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
```

Part 5: IPv6 configuration

Solution - Task 5.1: Native 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 need to 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
```

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
```

Local IPv6 summary routes is advertised to peer C2 also.

```
[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
```

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]
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"
```

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         0         0          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              0         0          I

[edit]
lab@R8# run show route advertising-protocol bgp 144.44.0.1 2001:101:10:32::/64

[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              0         0          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              0         0          I

```


Solution - Task 5.2: IPv6 tunneling configuration

- 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 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 can not use this approach.

The ISIS 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 ISIS 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 ISIS topology for IPv6 earlier is useful for ISIS IPv6 shortcuts also. This is required to ensure the IPv6 routes are following the route of the LSP between the 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 ISIS 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 ISIS IPv6 operation, but it is required by the task in this case. It will keep the transport label on the egress PE. Since 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 have to 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 ."
```

b. R5

Router R5 requires only the configuration of ISIS 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
```

c. Verify the configuration

After the ISIS 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
    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
                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
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

```

Part 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.

Solution - Task 6.1: VPNA configuration

- 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 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 need to 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
```

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 to accept also BGP neighbors under routing-instances you have to 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
```

- 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 has to 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.

```
[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. Since it is 4byte AS number, you have to 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
```

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
```

vrf-import policy accepts the routes with VPNA route-target received by MP-BGP from the other PEs.

```
[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
```

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
```

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 routing-instances VPNA protocols ospf export to.VPNA-CE2
```

Configure the sham-link between the two PEs:

```
[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
```

d. Verify the configuration

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

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: 46 destinations, 66 routes (46 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.0.0/24      *[OSPF/10] 00:08:29, metric 1
                  > to 10.10.0.2 via ge-0/0/5.900
10.100.1.0/24      *[OSPF/10] 00:08:29, metric 1
                  > to 10.10.0.2 via ge-0/0/5.900
10.100.2.0/24      *[OSPF/10] 00:08:29, metric 1
                  > to 10.10.0.2 via ge-0/0/5.900
10.100.3.0/24      *[OSPF/10] 00:08:29, metric 1
                  > to 10.10.0.2 via ge-0/0/5.900
10.100.4.0/24      *[OSPF/10] 00:08:29, metric 1
                  > to 10.10.0.2 via ge-0/0/5.900
10.100.5.0/24      *[OSPF/10] 00:08:29, metric 1
                  > to 10.10.0.2 via ge-0/0/5.900
10.100.6.0/24      *[OSPF/10] 00:08:29, metric 1
                  > to 10.10.0.2 via ge-0/0/5.900
10.100.7.0/24      *[OSPF/10] 00:08:29, metric 1
                  > to 10.10.0.2 via ge-0/0/5.900
10.100.8.0/24      *[BGP/170] 00:06:14, 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
10.100.9.0/24      *[BGP/170] 00:06:14, 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
10.100.10.0/24     *[BGP/170] 00:06:14, 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
10.100.11.0/24     *[BGP/170] 00:06:14, 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
10.100.12.0/24     *[BGP/170] 00:06:14, 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
10.100.13.0/24     *[BGP/170] 00:06:14, 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
10.100.14.0/24     *[BGP/170] 00:06:14, 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
10.100.15.0/24     *[BGP/170] 00:06:14, 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
```

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.

```
[edit]
lab@R7# run show route table VPNA

VPNA.inet.0: 46 destinations, 58 routes (46 active, 0 holddown, 10 hidden)
+ = Active Route, - = Last Active, * = Both

10.100.0.0/24      *[BGP/170] 00:07:47, 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 172.20.20.5 via ge-0/0/4.37, label-switched-path Bypass-
>172.20.20.9
10.100.1.0/24      *[BGP/170] 00:07:47, 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 172.20.20.5 via ge-0/0/4.37, label-switched-path Bypass-
>172.20.20.9
10.100.2.0/24      *[BGP/170] 00:07:47, 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 172.20.20.5 via ge-0/0/4.37, label-switched-path Bypass-
>172.20.20.9
10.100.3.0/24      *[BGP/170] 00:07:47, 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 172.20.20.5 via ge-0/0/4.37, label-switched-path Bypass-
>172.20.20.9
10.100.4.0/24      *[BGP/170] 00:07:47, 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 172.20.20.5 via ge-0/0/4.37, label-switched-path Bypass-
>172.20.20.9
10.100.5.0/24      *[BGP/170] 00:07:47, 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 172.20.20.5 via ge-0/0/4.37, label-switched-path Bypass-
>172.20.20.9
10.100.6.0/24      *[BGP/170] 00:07:47, 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 172.20.20.5 via ge-0/0/4.37, label-switched-path Bypass-
>172.20.20.9
10.100.7.0/24      *[BGP/170] 00:07:47, 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 172.20.20.5 via ge-0/0/4.37, label-switched-path Bypass-
>172.20.20.9
10.100.8.0/24      *[OSPF/10] 00:06:05, metric 1
                  > to 10.10.0.6 via ge-0/0/5.901
10.100.9.0/24      *[OSPF/10] 00:06:05, metric 1
                  > to 10.10.0.6 via ge-0/0/5.901
10.100.10.0/24     *[OSPF/10] 00:06:05, metric 1
                  > to 10.10.0.6 via ge-0/0/5.901
10.100.11.0/24     *[OSPF/10] 00:06:05, metric 1
                  > to 10.10.0.6 via ge-0/0/5.901
```

```

10.100.12.0/24      *[OSPF/10] 00:06:05, metric 1
                   > to 10.10.0.6 via ge-0/0/5.901
10.100.13.0/24      *[OSPF/10] 00:06:05, metric 1
                   > to 10.10.0.6 via ge-0/0/5.901
10.100.14.0/24      *[OSPF/10] 00:06:05, metric 1
                   > to 10.10.0.6 via ge-0/0/5.901
10.100.15.0/24      *[OSPF/10] 00:06:05, metric 1
                   > to 10.10.0.6 via ge-0/0/5.901

```

- 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 have to 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 have to make sure that R8 is the preferred PE for sending and receiving traffic to the site. To achieve this you have to 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 have to 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

[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 stanza.

```
[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
```

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

[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

```

c. 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: 559 destinations, 576 routes (554 active, 0 holddown, 5 hidden)
Restart Complete
```

```
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
```

```
VPNA.inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 10.10.0.0/30          10.10.0.10          64888 I
* 10.10.0.4/30          10.10.0.10          64888 I
* 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: 43 destinations, 43 routes (42 active, 0 holddown, 1 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.

```
[edit]
lab@R1# run show route receive-protocol bgp 10.10.0.14
```

```
inet.0: 509 destinations, 1439 routes (504 active, 0 holddown, 55 hidden)
Restart Complete
```

```
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```

```
VPNA.inet.0: 45 destinations, 86 routes (43 active, 0 holddown, 22 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: 45 destinations, 86 routes (43 active, 0 holddown, 22 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 10.10.0.16/30          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

- 4) Site C4 is configured to peer with AS 23.890 by mistake. You have to 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.

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 to pretend 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.

```
[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 have to ensure 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

[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
```

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
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 (42 active, 0 holddown, 1 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
* 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 (42 active, 0 holddown, 1 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
* 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. One way to disable that behavior is using the `as-override` knob. It will overwrite the originating AS with the MPLS core AS and the routes will be accepted by other sites part from the same VPN. However, it is not allowed to use this approach to solve the task as it's 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 sent to the same AS, is using the `advertise-peer-as` knob.

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
```

b. R5

```
[edit]
lab@R5# set routing-instances VPNA protocols bgp group VPNA-CE4 advertise-peer-as
```

c. R8

```
[edit]
lab@R8# set routing-instances VPNA protocols bgp group VPNA-CE3 advertise-peer-as
```

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: 45 destinations, 115 routes (45 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 10.10.0.0/30           Self              0         0          65002 64888 I
* 10.10.0.4/30           Self              0         0          65002 64888 I
* 10.100.16.0/24         Self              0         0          65002 64888 I
* 10.100.17.0/24         Self              0         0          65002 64888 I
* 10.100.18.0/24         Self              0         0          65002 64888 I
* 10.100.19.0/24         Self              0         0          65002 64888 I
* 10.100.20.0/24         Self              0         0          65002 64888 I
* 10.100.21.0/24         Self              0         0          65002 64888 I
* 10.100.22.0/24         Self              0         0          65002 64888 I
* 10.100.23.0/24         Self              0         0          65002 64888 I
```

6) Use static routing to resolve any hidden routes in RIB inet.3 on the route-reflector.

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.

a. Route-reflector

```
[edit]
lab@route-reflector# set routing-options rib inet.3 static route 0.0.0.0/0
```

```
[edit]
lab@route-reflector# set routing-options rib inet6.3 static route
0:0:0:0:ffff::/96
```

b. Verify the configuration

```
[edit]
lab@route-reflector# run show route summary
Autonomous system number: 65001
Router ID: 172.20.25.10

inet.0: 528 destinations, 530 routes (528 active, 0 holddown, 0 hidden)
  Direct:      4 routes,      4 active
  Local:       3 routes,      3 active
  BGP:        490 routes,    488 active
  IS-IS:      33 routes,     33 active

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

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

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

bgp.l3vpn.0: 41 destinations, 42 routes (41 active, 0 holddown, 0 hidden)
  BGP:        42 routes,    41 active

inet6.0: 143 destinations, 145 routes (143 active, 0 holddown, 0 hidden)
  Direct:      2 routes,     2 active
  Local:       2 routes,     2 active
  BGP:       116 routes,   114 active
  IS-IS:      25 routes,   25 active

inet6.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Static:      1 routes,     1 active

bgp.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
  BGP:        6 routes,     6 active

bgp.mvpn.0: 6 destinations, 8 routes (6 active, 0 holddown, 0 hidden)
  BGP:        8 routes,     6 active
```

Solution - Task 6.2: VPNA multicast configuration

- 7) 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.
- 8) 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 have to enable VPNA for multicast support using NG-MVPN. Unlike the Draft-Rosen, 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. Draft-Rosen supports only PIM/GRE. NG-MVPN on the other hand provides a higher flexibility and in addition to PIM/GRE supports a few more options for the multicast dataplane. This includes 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.

```
[edit]
lab@R1# set protocols bgp group internal family inet-mvpn signaling
```

- 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 have to add the group to the range of SSM groups.

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

- 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@R1# 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
```

d. R1

The configuration of routers R1 and R5, where the receivers are located, is much simpler. Since, they are dedicated as receiver-only sites, only the mvpn protocol has to 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]
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
```

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
```

- 9) Manually designate the receiving sites and ensure that they cannot be 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
```

Since the source is multihomed to routers R3 and R7, the rest of the PEs have to select single forwarded to avoid duplicating the multicast traffic. By default, the PE with the

numerical 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 mvpn stanza to routing-instance.

```
[edit]
lab@R1# set routing-instances VPNA protocols mvpn unicast-umh-election
```

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
```

h. R3

```
[edit]
lab@R3# set routing-instances VPNA protocols mvpn unicast-umh-election
```

i. R7

```
[edit]
lab@R7# set routing-instances VPNA protocols mvpn unicast-umh-election
```

10) 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

Configuring dedicated virtual-tunnels interface to serve only the multicast traffic.

```
[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
```

k. R5

```
[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]
lab@R5# set routing-instances VPNA protocols mvpn route-target export-target
target target:1507417L:200
```

l. R3

```
[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
```

m. R7

```
[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

[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
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
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
  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

```

Solution - Task 6.3: VPLS configuration

- 11) Configure 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.

The Task requires configuring VPLS connectivity between the three VPLS sites. Although VPLS provides point-to-multipoint communications, you have to 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.

```

[edit]
lab@R1# set protocols bgp group internal family l2vpn signaling

```

- b. R2

Configure a new physical interface to connect site VPLS-HUB. You have to 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 have to 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

[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
```

c. R7

Configure a new physical interface to connect site VPLS-S1. You have to 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]
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
```

Lowering the site-range to 2 from the default value of 65534. This way this site could 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
```

d. R8

Configure a new physical interface to connect site VPLS-S2. You have to 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.

```
[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
```

- 13) According to the diagram, the two VPLS spoke sites are configured with different VLANs than the HUB. You have to 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 have to 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
```

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
```

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(swaps .1500) Out(swaps
.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
```

- 14) 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

Limiting 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
```

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
```

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
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
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   OR

```

```

[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

```

```

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   OR

```

The ping tests are successful from the Spoke sites to the Hub. However, the ping between the Spokes is not possible.

```

[edit]
root@vr-device# run 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
```

```
[edit]
root@vr-device# run 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
```

Solution - Task 6.4: L2VPN configuration

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 *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 have to 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
```

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
```

b. R4

The 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 has to be enabled. The configuration below is the only 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
```

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
```

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
```

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
```

- 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.

JUNOS allows for steering L2VPN traffic over a particular LSP. Use a community assigned to the L2circuit and then match on it using routing policy, which is then applied to the forwarding table

R4

Configuring a dedicated LSP avoiding links with GOLD color.

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r5-l2circuit to 172.20.25.5

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r5-l2circuit 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.

```
[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-l2circuit
```

Defining the community under the policy-option stanza.

```
[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
```

d. 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-l2circuit to 172.20.25.4

[edit]
```

```

lab@R5# set protocols mpls label-switched-path r5-to-r4-l2circuit 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

[edit]
lab@R5# set policy-options policy-statement mapping-l2circuit-lsp term 1 then
install-nexthop lsp r5-to-r4-l2circuit

[edit]
lab@R5# set policy-options community l2circuit-map members 1700:1700

[edit]
lab@R5# set routing-options forwarding-table export mapping-l2circuit-lsp

```

e. 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

```

```

[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

+ = 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
l2circuit       > 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
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
  
```

Using ping to test the connectivity between the sites.

[edit]

root@vr-device# **run 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
  
```


Chapter eleven solutions: Full Day Lab Challenge III

This lab scenario represents a complete 8 hour challenge, aiming to prepare potential candidate for the JNCIE-SP lab exam. It contains all major topics found in the exam blueprint.

The external devices (CE and peers) are virtualized on the VR-device. Use only the interfaces and VLANs on the diagram for connectivity between the routers.



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



Part 1: System Features

This 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.

Solution - Task 1.1 Service Configuration

NOTE: It is highly recommended to read the whole lab and verify 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 on how the final network should look like.

- 1) SSH access is preconfigured as part of the initial configuration on all routers. Limit the SSH connections to 2 at a time and no more than 2 connection attempts per minute.
 - a. All devices.

RACK RENTAL NOTE: SSH is already configured as part of the initial configuration with user root allowed to login. Please, do not change it, as it is required by iNETZERO rack rental service.

Below two configuration lines are required to successfully complete the first task.

```
[edit]
lab@R1# set system services ssh connection-limit 2
```

```
[edit]
lab@R1# set system services ssh rate-limit 2
```

- 2) Enable NETCONF protocol over SSH using the standard port – TCP 830.
 - a. All devices.

NETCONF provides mechanisms to install, manipulate, and delete the configuration of network devices. After enabling SSH on JUNOS devices, NETCONF is enabled automatically accepting connections on the default SSH port 22. The task asks for enabling NETCONF over the standard port 830 defined by RFC4742. In JUNOS, this is done when you add `ssh` option under the `netconf` service.

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

Below is excerpt from the system connection table on router R1 after committing the change.

```
[edit]
lab@R1# run show system connections
```

tcp4

0

0

*.830

.

LISTEN

- 3) Configure SNMPv2c with community *superlab* on all routers, polled from management system with address 10.10.10.1/32. Since SNMPv2c is using clear text transmissions, make sure that SNMP traffic is accepted only over the management interface *fxp0.0*. Ensure that the reachability to 10.10.10.1/32 is provided even if the *rpd* process is not yet running. Assume the IP address 10.10.1.200 is used as gateway in the management segment.

Next, you have to configure the SNMPv2c management protocol. As the task suggests, you need to secure the SNMP communication to the devices also. You have to explicitly specify the exact interface where the request can come from and the exact IP address of the network management system that can access it. Along with that, you have to configure a single static route reaching the NMS via the management interface *fxp0*. Next part requires that the route to the management system is reachable when the RPD process is not running. This can happen if the process crashes or during the system boot process. To accomplish this, *backup-router* command under the *system* stanza has to be configured. When RDP process starts the backup route is removed from the local routing and forwarding tables.

- a. All devices.

```
[edit]
lab@R1# set snmp interface fxp0.0

[edit]
lab@R1# set snmp community superlab clients 10.10.10.1/32

[edit]
lab@R1# set routing-options static route 10.10.10.1/32 next-hop
10.10.1.200

[edit]
lab@R1# set system backup-router 10.10.1.200

[edit]
lab@R1# set system backup-router destination 10.10.10.1/32
```

Solution - Task 1.2: Centralized authentication management

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. This Task requires configuring the devices to communicate successfully with a TACACS+ server.

- 4) Configure all devices to use TACACS server located at 10.10.10.1 for authentication. The communication with the TACACS server should be encrypted with password *jncie123*.

- a. All devices.

```
[edit]
lab@R1# set system tacplus-server 10.10.10.1 secret
"$9$ZPDHmtpBEhrz3A0B1hcbsYoUj"
```

- 5) Ensure that local user database is used for authentication only when the TACACS server is unreachable.

Configure only `tacplus` in the authentication order will successfully accomplish the requirement. This way the local database is only used if the TACACS sever is unreachable.

- a. All devices.

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

Solution - Task 1.3: Local user configuration

- 6) Configure user `noc1` and `noc2` using local passwords `noc123` and `noc456` respectively. Any session of users `noc1` and `noc2` should be disconnected after 5 minutes of inactivity. The users should have full access to the device but without the possibility to manipulate or list files on the local storage. Entering the shell mode should be restricted as well.

The task asks for configuring two local user accounts having access to all CLI levels. The user's sessions must be disconnected automatically after 5 minutes of inactivity. The Idle timeout can be configured under a `class` definition only. It cannot be configured under the user section. Because of that and because both users share the same attributes it is better if you configure a single login class. To configure the correct permissions, you need to use a regular expression explicitly denying what should not be available to the users.

- a. All devices.

```
[edit]
lab@R1# set system login class noc idle-timeout 5

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

[edit]
lab@R1# set system login class noc deny-commands "file *|start shell"

[edit]
lab@R1# set system login user noc1 class noc

[edit]
lab@R1# set system login user noc1 authentication encrypted-password
"$1$mpieBnO/$Bxv4vufCWgpSHePZWmBcv1"

[edit]
lab@R1# set system login user noc2 class noc

[edit]
lab@R1# set system login user noc2 authentication encrypted-password
"$1$xJPr9/Hu$41PGFN11scAhWItEkyx/0"
```

NOTE: The user password is displayed as it is printed on the terminal. The \$1\$ in the beginning suggests that a MD5 hash algorithm was used to hash the password. By default, JUNOS uses MD5 to securely store the local user passwords.

To verify that the correct permissions are applied to a particular user. Login with the user credentials and issue the command below.

```
noc1@R1> show cli authorization
Current user: 'noc1' class 'noc'
Permissions:
  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
    unified-edge-- Can view unified edge configuration
    unified-edge-control-- Can modify unified edge configuration
    storage       -- Can view fibre channel storage protocol configuration
    storage-control-- Can modify fibre channel storage protocol configuration
    all-control  -- Can modify any configuration
Individual command authorization:
  Allow regular expression: none
  Deny regular expression: file *|start shell
  Allow configuration regular expression: none
  Deny configuration regular expression: none
```

- 7) Users must be locked out of the device for 3 minutes if they fail to provide correct password 3 times.

The next task requirements are straightforward. All you have to do is configure the two lines below.

- a. All devices.

```
[edit]
lab@R1# set system login retry-options tries-before-disconnect 3
```

```
[edit]
lab@R1# set system login retry-options lockout-period 3
```

- 8) Make sure that system alarms are displayed during login of any user not part of default login classes. You have to ensure that users of any current or future custom login classes would have this requirement.

By default, only the critical log messages are displayed to the user console. As with the idle timeout, this can be configured only in a login class. The task requires the change to have effect for already configured classes, but also for any new one configured in the future. To accomplish this you have to use apply-groups. The feature allows you to create configuration containers, which can be applied to different levels in the configuration. The same configuration CLI structures are used to define the apply-groups. One significant advantage of configuring apply-groups is to use of wildcard expressions that can inherit any statement that matches the expression.

The configuration lines below show how you can accomplish the task using configuration groups.

- a. All devices

```
[edit]
lab@R1# set groups user_alarm system login class <*> login-alarms
```

```
[edit]
lab@R1# set system login apply-groups user_alarm
```

Solution - Task 1.4: Active configuration archival and logging

- 9) You should ensure that all devices transfer the configuration every 2 hours to a SCP server located on address 10.10.10.1. Using username *lab* and password *workbook*. The configuration files must be copied to destination folder */home/archive* on the server.

- a. All devices.

JUNOS has embedded features to automatically send the configuration to remote FTP or secure SCP servers. The administrator can configure the transfer to happen once per configured interval or every time when a commit of the configuration is performed.

```
[edit]
lab@R1# set system archival configuration transfer-interval 120
```

```
[edit]
lab@R1# set system archival configuration archive-sites
"scp://lab@10.10.10.1/home/archive" password
"$9$tfQNOhrbwgaGixNVYoGq.tuORcl"
```


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.

10) Configure all routers to send all log messages related to the *rpd* process and the kernel to syslog server 10.10.10.1.

a. All devices.

JUNOS offers a rich set of methods for collecting and storing logs. One method is sending the logs to an external server using the syslog protocol. Using regular expressions provide a powerful way to selectively match on certain log messages.

Configure the devices to send logs to syslog server with IP address 10.10.10.1 and match only the messages related to the *rpd* process and the kernel.

```
[edit]
lab@R1# set system syslog host 10.10.10.1 any any

[edit]
lab@R1# set system syslog host 10.10.10.1 match "rpd|kernel"
```

11) Amend the default syslog configuration in the way that all commands run by any user should be saved in file *interactive-commands* stored locally. To collect enough data for audit trail, but in the same to conserve space, make sure that no more than 5 files each 3MB are kept in the storage.

a. All devices.

The task requires configuring a log file stored locally on the routers. You have to change the archive settings as requested. Moreover you have to make the file world-readable making it accessible by any user on the system.

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

[edit]
lab@R1# set system syslog file interactive-commands archive size 3m

[edit]
lab@R1# set system syslog file interactive-commands archive files 5

[edit]
lab@R1# set system syslog file interactive-commands archive world-
readable
```

Solution - Task 1.5: Advanced Interface configuration and chassis functions

12) Configure an aggregated Ethernet interface connecting R3 and R4. Assume that a configuration template requires the interface to have the name *ae3*. Ensure the configuration template continues to be used successfully without adjusting it.

a. R3

Next section is related to advanced interface configuration. You have to configure an aggregated Ethernet interface based on the requirements in the Task. Before you can create an aggregated interface in JUNOS, you need to define the number of the ae-interface available on the system. This will effectively create an empty aggregated interface with a globally unique MAC address assigned. The aggregated interfaces are numbered in a sequenced order starting from ae0. The requirement is to configure a single interface with number ae3. To do that you have to set the aggregate devices to 4 under the chassis configuration which creates aggregated interfaces in the range from ae0 to ae3.

```
[edit]
lab@R3# set chassis aggregated-devices ethernet device-count 4

[edit]
lab@R3# set interfaces ge-0/0/1 gigether-options 802.3ad ae3

[edit]
lab@R3# set interfaces ge-0/0/2 gigether-options 802.3ad ae3

[edit]
lab@R3# set interfaces ae3 aggregated-ether-options lacp active

[edit]
lab@R3# set interfaces ae3 unit 0 description "=== connection to R4 ==="

[edit]
lab@R3# set interfaces ae3 unit 0 family inet address 10.50.0.81/30
```

b. R4

```
[edit]
lab@R4# set chassis aggregated-devices ethernet device-count 4

[edit]
lab@R4# set interfaces ge-0/0/1 gigether-options 802.3ad ae3

[edit]
lab@R4# set interfaces ge-0/0/2 gigether-options 802.3ad ae3

[edit]
lab@R4# set interfaces ae3 aggregated-ether-options lacp active

[edit]
lab@R4# set interfaces ae3 unit 0 description "=== connection to R3 ==="

[edit]
lab@R4# set interfaces ae3 unit 0 family inet address 10.50.0.82/30
```

c. Verify the configuration.

Aggregated interface with name ae3 is successfully created.

```
[edit]
lab@R3# run show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Up
Interface index: 140, SNMP ifIndex: 518
```

```

Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Disabled
Pad to minimum frame size: Disabled
Device flags      : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags       : None
CoS queues      : 8 supported, 8 maximum usable queues
Current address: 00:05:86:49:ba:c0, Hardware address: 02:06:0a:0e:03:01
Last flapped    : 2016-12-13 08:21:40 UTC (00:50:40 ago)
Input rate      : 0 bps (0 pps)
Output rate     : 24 bps (0 pps)
Active alarms   : None
Active defects   : None
Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 332) (SNMP ifIndex 527)
Flags: Up SNMP-Traps 0x4004000 Encapsulation: ENET2
Input packets : 52
Output packets: 3
Protocol aenet, AE bundle: ae3.0

```

- 13) Configure the aggregated Ethernet interface to monitor the member links. Due to the fact, that vMX is a software based device, increase the monitoring interval to 30s.

Next you have to configure the LACP protocol running over the aggregated Ethernet interface to monitor the operation of the interface. Unlike unmonitored ae- interface, LACP requires at least one of the sides to be configured as active in the bundle. By default the interfaces are created as passive, hence you need to configure one of ends as active otherwise the interfaces cannot be bundled.

a. R3

```

[edit]
lab@R3# set interfaces ae3 aggregated-ether-options lacp periodic slow

```

b. R4

```

[edit]
lab@R4# set interfaces ae3 aggregated-ether-options lacp periodic slow

```

Verifying the status of the aggregated Ethernet interface.

```

[edit]
lab@R3# run show lacp interfaces ae3
Aggregated interface: ae3

```

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
ge-0/0/1	Actor	No	No	Yes	Yes	Yes	Yes	Slow	Active
ge-0/0/1	Partner	No	No	Yes	Yes	Yes	Yes	Slow	Active
ge-0/0/2	Actor	No	No	Yes	Yes	Yes	Yes	Slow	Active
ge-0/0/2	Partner	No	No	Yes	Yes	Yes	Yes	Slow	Active

```

LACP protocol:
ge-0/0/1      Receive State  Current  Slow periodic  Collecting distributing
ge-0/0/2      Receive State  Current  Slow periodic  Collecting distributing

```

- 14) Enable support of switching label packets on all logical interfaces on all routers except the RR. You must use the most efficient way configuration-wise to accomplish this. This change should not affect the management interface.

To solve the task you have to utilize apply-groups to configure the mpls family to all interfaces. However, you have to exclude the fxp0 from the list of interface. Apply-groups are inherit by the configuration file when applied to the CLI hierarchy they are intended to configure. Similarly, they can be excluded when they are applied deeper in the hierarchy. The use of wildcard expressions inherit any statement that matches the expression.

a. All devices

```
[edit]
lab@R1# set groups enable_mpls interfaces <*> unit <*> family mpls

[edit]
lab@R1# set interfaces apply-groups enable_mpls

[edit]
lab@R1# set interfaces fxp0 unit 0 apply-groups-except enable_mpls

[edit]
lab@R1# show interfaces | display inheritance
ge-0/0/4 {
  vlan-tagging;
  unit 12 {
    description "=== connection to R2 ===";
    vlan-id 12;
    family inet {
      address 10.50.0.1/30;
    }
    family inet6;
    ##
    ## 'mpls' was inherited from group 'enable_mpls'
    ##
    family mpls;
  }
  unit 14 {
    description "=== connection to R4 ===";
    vlan-id 14;
    family inet {
      address 10.50.0.5/30;
    }
    family inet6;
    ##
    ## 'mpls' was inherited from group 'enable_mpls'
    ##
    family mpls;
  }
}
```

Solution - Task 1.6: Advanced RE Protection

15) You must protect the control plane of all routers, including the route-reflector.

Configure input firewall filter using the following criteria:

Last task is the most time consuming since you have to configure firewall filters with many terms. In addition, this is the first place where you have to 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 the main IPv4 protocols: SSH, SNMP, BFD, OSPF, RIP, PIM, LDP, RSVP, TACACS+, ping and traceroute. If needed, also allow protocols required for tasks

<http://www.inetzero.com> - Copyright 2017 iNET ZERO. All rights reserved

For personal non commercial use only – do not distribute

<https://t.me/learningnets>

a. All devices including the router-reflector

The SSH term is configured with a policer assigned to it to limit the traffic to 5Mbps. The actual policer configuration is shown in the next step of the Task.

```
[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 policer
5m-limit

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ssh then accept

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-tacacs from
source-port tacacs

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-tacacs then
policer 5m-limit

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-tacacs then
accept
```

The SNMP term must allow only traffic from configured sources under the SNMP configuration stanza. This is accomplished with an apply-path prefix-list that can extract the addresses using wildcard masks. The exact prefix-list configuration is shown in the next steps.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-snmp from
destination-prefix-list snmp-host

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-snmp from
protocol udp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-snmp from
destination-port snmp

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-snmp then policer
5m-limit

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-snmp then accept
```

The BFD protocol is a simple hello mechanism used for detecting failures in the network. Usually used with routing protocols like BGP, OSPF, ISIS etc. for faster neighbor detection. It is supported over direct links and multihop as well. This is usually needed for iBGP

session that can be established across several hops. Depending on the type of BFD session, different ports are used. Port 3784 is used for direct sessions whereas 4784 for multihop sessions. You have to enable both. At a later stage BFD will be required for monitoring iBGP sessions.

```
[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
```

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bfd from port
4784
```

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-bfd then accept
```

```
[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
```

Here again apply-path prefix lists are used because the task requires that all BGP neighbours must be automatically allowed by the filter.

```
[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
destination-prefix-list ipv4-local-addresses
```

```
[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# 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
```

```
[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

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ldp then accept

[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

```

The task requires to only accept ping traffic rather than all ICMP protocol messages. This means that you have to narrow down the match criteria to only echo-reply and echo-request. Like the SSH term, ping traffic must be policed up to 5Mbps.

```

[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 policer
5m-limit

[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-ping then accept

```

Traceroute tool uses UDP or ICMP packets for detecting the hops to the destination. You have to configure the filter to accept both options. UDP ports by default defined in the RFC are within the range from 33434 to 33534.

```

[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
policer 5m-limit
```

```
[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
```

```
[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 policer 5m-limit
```

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term accept-traceroute-reply
then accept
```

- Allow main IPv6 protocols required by the subsequences tasks: OSPFv3, BGP, RIPng, ping and traceroute.
 - a. All devices including the router-reflector

Unlike the IPv4 filters, matching IPv6 traffic 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
```

```
[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-bgp
from destination-prefix-list ipv6-local-addresses
```

```
[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
```

```
[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-
ripng from next-header udp
```

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 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 for 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
```

```
[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 policer 5m-limit
```

```
[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
```

As the IPv4 traceroute tool, the IPv6 counterpart is using the same port range.

```
[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 policer 5m-limit
```

```
[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 policer 5m-limit
```

```
[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term accept-
traceroute-reply then accept
```

- Ensure that SNMP packets are accepted only from already configured hosts, as well as any added additionally without requiring a change in the firewall filter configuration.

- a. All devices including the router-reflector

Apply-path prefix-lists are similar to the configuration apply-groups but instead to populate the configuration, they extract configuration statements and create a dynamic list of IP addresses. They also leverage the use wildcards to match part of the configuration representing free text.

```
[edit]
lab@R1# set policy-options prefix-list snmp-host apply-path "snmp
community <*> clients <*>"
```

You can verify the matched statements added to the prefix list dynamically with the `display inheritance` option. An example is shown below.

```
[edit]
lab@R1# show policy-options prefix-list snmp-host | display inheritance
##
## apply-path was expanded to:
## 10.10.10.1/32;
##
##
apply-path "snmp community <*> clients <*>";
```

- You must accept BGP sessions only from configured neighbors. Any additional neighbors must be automatically accepted.
 - a. All devices including the router-reflector

Apply-path prefix-list are also used for matching dynamically the configured BGP neighbors.

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

[edit]
lab@R1# show policy-options prefix-list bgp-neighbors | display inheritance
##
## apply-path was expanded to:
## 10.50.250.10/32;
## 172.17.0.10/32;
##
##
apply-path "protocols bgp group <*> neighbor <*>";
```

- Limit SSH, SNMP, TACACS+, ping and traceroute protocol traffic to not more than 5Mbps.
 - a. All devices including the router-reflector

For the policer the bandwidth limit and the burst size limit are mandatory configuration statements. The bandwidth limit is the actual limit allowed by the policer, whereas the burst size is the amount of the traffic that can pass not policed in case of a burst. Usually it is calculated to cover 10 times the MTU size.

```
[edit]
lab@R1# set firewall policer 5m-limit if-exceeding bandwidth-limit 5m

[edit]
lab@R1# set firewall policer 5m-limit if-exceeding burst-size-limit 15k

[edit]
lab@R1# set firewall policer 5m-limit then discard
```

- All other traffic must be silently discarded. Make sure that the last 400 packets are logged in the local volatile memory.

a. All devices including the router-reflector

The last part of the filter must silently discard the rest of the packets not matched by any term above. There are two ways to drop packet in firewall filter. The reject action will send back icmp unreachable message to the source and scrap the packet. Using the discard option will silently discard the packet. Besides that, you have to log the last 400 discarded packets locally on the devices. Important point here is the fact that the local volatile memory has to be used for storage the logs. Using the action log will log the last 400 packets locally in FPC memory of the interface cards.

```
[edit]
lab@R1# set firewall filter protect-re-ipv4 term discard-all then log

[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 log

[edit]
lab@R1# set firewall family inet6 filter protect-re-ipv6 term discard-all
then discard
```

Applying the filters to the Loopback interface. This will ensure that traffic received through any physical interfaces to the RE is inspected by the filters. You must apply the filter in the appropriate family.

```
[edit]
lab@R1# set interfaces lo0 unit 0 family inet filter input protect-re-
ipv4

[edit]
lab@R1# set interfaces lo0 unit 0 family inet6 filter input protect-re-
ipv6
```

The output below displays the content of the local volatile memory. A few discarded packets are dropped by the filter and logged in the non-volatile memory.

```
[edit]
lab@R1# run show firewall log
Log :
Time      Filter  Action Interface  Protocol  Src Addr
Dest Addr
09:50:43  pfe      D      ge-0/0/4.16  PIM      fe80::206:a00:100e:604
ff02::d
09:50:31  pfe      D      ge-0/0/4.12  PIM      fe80::206:a00:c0e:204
ff02::d
09:50:22  pfe      D      ge-0/0/4.14  PIM      fe80::206:a00:e0e:404
ff02::d
09:50:08  pfe      D      ge-0/0/4.16  PIM      fe80::206:a00:100e:604
ff02::d
09:49:58  pfe      D      ge-0/0/4.12  PIM      fe80::206:a00:c0e:204
ff02::d
09:49:46  pfe      D      ge-0/0/4.14  PIM      fe80::206:a00:e0e:404
ff02::d
09:49:33  pfe      D      ge-0/0/4.14  IGMP     10.50.0.6
224.0.0.1
```

Part 2: Troubleshooting and Configuring IGP

This part focuses on troubleshooting a partially functional OSPFv2 and OSPFv3 network. You have to fix the problems according the given list of requirements and also configure additional requirements. It is highly recommended to read the whole part before starting to solve the tasks.

Solution - Task 2.1: Troubleshooting

First Task requires fixing the partial working IGP. Using the diagram provided, you have to make sure that all adjacencies are established and routes are successfully exchanged.

- 1) Ensure all of the OSPF neighbors are successfully established (Full state). Make sure the following criteria are met:

Start with the examination the current state of the network. List the neighbor table on each device to determine what is not working as expected. Since the network must support IPv6 addressing you have to troubleshoot both OSPF and OSPFv3 routing protocol.

- All neighbors that require authentication must authenticate successfully to Full state.
 - a. OSPF and OSPFv3 states on R1

The OSPF neighbor table on R1 looks like everything is working properly. All of the three neighbors are successfully reached Full state. However, if you carefully examine the content of the table below you may notice that the neighbor IDs of the device on interfaces ge-0/0/4.14 and ge-0/0/4.16 are the same.

```
[edit]
lab@R1# run show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.50.0.2	ge-0/0/4.12	Full	10.50.250.2	128	28
10.50.0.6	ge-0/0/4.14	Full	10.50.250.6	128	33
10.50.0.10	ge-0/0/4.16	Full	10.50.250.6	128	39

```
[edit]
lab@R1# run show ospf interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
ge-0/0/4.12	BDR	0.0.0.0	10.50.250.2	10.50.250.1	1
ge-0/0/4.14	BDR	0.0.0.0	10.50.250.6	10.50.250.1	1
ge-0/0/4.16	BDR	0.0.0.0	10.50.250.6	10.50.250.1	1

The table below displays the neighbor table of the OSPFv3 protocol. Unlike the OSPF table, the neighbors in Full state are only two. The expected number of OSPFv3 neighbors in Full state should be three as well. Moreover, the OSPFv3 interface table has only two interfaces enabled in OSPFv3.

```
[edit]
lab@R1# run show ospf3 neighbor
```

ID	Interface	State	Pri	Dead
10.50.250.2	ge-0/0/4.12	Full	128	26
Neighbor-address fe80::206:a00:c0e:204				
10.50.250.6	ge-0/0/4.14	Full	128	35
Neighbor-address fe80::206:a00:e0e:404				

```
[edit]
lab@R1# run show ospf3 interface
Interface      State   Area      DR ID      BDR ID      Nbrs
ge-0/0/4.12    BDR     0.0.0.0    10.50.250.2 10.50.250.1 1
ge-0/0/4.14    BDR     0.0.0.0    10.50.250.6 10.50.250.1 1
```

b. OSPF and OSPFv3 states on R2

The OSPF neighbor table on R2 contains one neighbor less than the expected number of three neighbors. There are three OSPF enabled interfaces, with one without a neighbor.

```
[edit]
lab@R2# run show ospf neighbor
Address        Interface      State   ID          Pri   Dead
10.50.0.1      ge-0/0/4.12   Full    10.50.250.1 128   38
10.50.0.14     ge-0/0/4.24   Full    10.50.250.6 128   32

[edit]
lab@R2# run show ospf interface
Interface      State   Area      DR ID      BDR ID      Nbrs
ge-0/0/4.12    DR      0.0.0.0    10.50.250.2 10.50.250.1 1
ge-0/0/4.24    BDR     0.0.0.0    10.50.250.6 10.50.250.2 1
ge-0/0/4.26    DR      0.0.0.0    10.50.250.2 0.0.0.0     0
```

The OSPFv3 neighbor table has all the expected neighbors in Full state. Similar to router R1, the fact that two of the neighbors have the same router IDs suggest that something is wrong. As with the OSPF protocol, OSPFv3 protocol is using the same 32-bit router ID to uniquely identify the nodes in the network.

```
[edit]
lab@R2# run show ospf3 neighbor
ID             Interface      State   Pri   Dead
10.50.250.1    ge-0/0/4.12   Full    128   37
  Neighbor-address fe80::206:a00:c0e:104
10.50.250.6    ge-0/0/4.24   Full    128   36
  Neighbor-address fe80::206:a00:180e:404
10.50.250.6    ge-0/0/4.26   Full    128   34
  Neighbor-address fe80::206:a00:1a0e:604

[edit]
lab@R2# run show ospf3 interface
Interface      State   Area      DR ID      BDR ID      Nbrs
ge-0/0/4.12    DR      0.0.0.0    10.50.250.2 10.50.250.1 1
ge-0/0/4.24    BDR     0.0.0.0    10.50.250.6 10.50.250.2 1
ge-0/0/4.26    BDR     0.0.0.0    10.50.250.6 10.50.250.2 1
```

c. OSPF and OSPFv3 states on R3

According to the diagram, R3 is internal router in Area 50 and must have three OSPF neighbors. Nevertheless, the OSPF neighbor table is empty, whereas two interfaces are enabled for OSPF. If you carefully compare the interface tables for OSPF and OSPFv3, you may notice that in fact the Area number is different. The State of the interface to R7 is different from DR or BDR. On the broadcast Interfaces with only two routers it is expected to have a DR or BDR state for the neighbors. A State different than DR and BDR should attract your attention for further examination.

```
[edit]
lab@R3# run show ospf interface
Interface      State   Area      DR ID      BDR ID      Nbrs
ge-0/0/4.35    DR      0.0.5.0    10.50.250.3 0.0.0.0     0
ge-0/0/4.37    DRouter 0.0.5.0    0.0.0.0     0.0.0.0     0
```

The OSPFv3 neighbor table has one neighbor successfully reached Full state with router R5.

```
[edit]
lab@R3# run show ospf3 interface
Interface          State   Area      DR ID          BDR ID          Nbrs
ge-0/0/4.35        BDR    0.0.0.50  10.50.250.5    10.50.250.3     1
ge-0/0/4.37        DR     0.0.0.50  10.50.250.3    0.0.0.0         0

[edit]
lab@R3# run show ospf3 neighbor
ID          Interface      State   Pri   Dead
10.50.250.5 ge-0/0/4.35    Full    128   37
Neighbor-address fe80::206:a00:230e:504
```

d. OSPF and OSPFv3 states on R4

R4 should have five neighbors, but currently only four are in Full state. The OSPF interface table contains all five interfaces.

```
[edit]
lab@R4# run show ospf neighbor
Address      Interface      State   ID          Pri   Dead
10.50.0.5    ge-0/0/4.14    Full    10.50.250.1  128   39
10.50.0.13   ge-0/0/4.24    Full    10.50.250.2  128   39
10.50.0.34   ge-0/0/4.45    Full    10.50.250.5  128   39
10.50.0.94   ge-0/0/4.47    Full    10.50.250.7  128   31

[edit]
lab@R4# run show ospf interface
Interface      State   Area      DR ID          BDR ID          Nbrs
ge-0/0/4.14    DR     0.0.0.0   10.50.250.6    10.50.250.1     1
ge-0/0/4.24    DR     0.0.0.0   10.50.250.6    10.50.250.2     1
ge-0/0/4.45    DR     0.0.0.0   10.50.250.6    10.50.250.5     1
ge-0/0/4.46    DR     0.0.0.0   10.50.250.6    0.0.0.0         0
ge-0/0/4.47    BDR    0.0.0.50  10.50.250.7    10.50.250.6     1
```

The OSPFv3 neighbor table has even less neighbors in Full state. Two of the five OSPFv3 neighbors are totally missing.

```
[edit]
lab@R4# run show ospf3 neighbor
ID          Interface      State   Pri   Dead
10.50.250.1 ge-0/0/4.14    Full    128   38
Neighbor-address fe80::206:a00:e0e:104
10.50.250.2 ge-0/0/4.24    Full    128   32
Neighbor-address fe80::206:a00:180e:204
10.50.250.5 ge-0/0/4.45    Full    128   38
Neighbor-address fe80::206:a00:2d0e:504

[edit]
lab@R4# run show ospf3 interface
Interface      State   Area      DR ID          BDR ID          Nbrs
ge-0/0/4.14    DR     0.0.0.0   10.50.250.6    10.50.250.1     1
ge-0/0/4.24    DR     0.0.0.0   10.50.250.6    10.50.250.2     1
ge-0/0/4.45    DR     0.0.0.0   10.50.250.6    10.50.250.5     1
ge-0/0/4.46    DR     0.0.0.0   10.50.250.6    0.0.0.0         0
ge-0/0/4.47    DR     0.0.0.50  10.50.250.6    0.0.0.0         0
```

e. OSPF and OSPFv3 states on R5

The OSPF neighbor table on R5 contains only one neighbor in state the Full, however 3 interfaces are enabled for OSPF.

```
[edit]
lab@R5# run show ospf neighbor
Address      Interface      State      ID              Pri    Dead
10.50.0.33    ge-0/0/4.45    Full       10.50.250.6     128    34

[edit]
lab@R5# run show ospf interface
Interface      State      Area      DR ID           BDR ID          Nbrs
ge-0/0/4.45    BDR       0.0.0.0   10.50.250.6     10.50.250.5     1
ge-0/0/4.56    DR        0.0.0.0   10.50.250.5     0.0.0.0         0
ge-0/0/4.35    DR        0.0.0.50  10.50.250.5     0.0.0.0         0
```

The OSPFv3 table displays all three neighbors in state Full, that is the expected number of neighbors without considering R8. Again the router ID of two neighbors is the same.

```
[edit]
lab@R5# run show ospf3 neighbor
ID            Interface      State      Pri    Dead
10.50.250.6    ge-0/0/4.45    Full       128    38
Neighbor-address fe80::206:a00:2d0e:404
10.50.250.6    ge-0/0/4.56    Full       128    33
Neighbor-address fe80::206:a00:380e:604
10.50.250.3    ge-0/0/4.35    Full       128    35
Neighbor-address fe80::206:a00:230e:304

[edit]
lab@R5# run show ospf3 interface
Interface      State      Area      DR ID           BDR ID          Nbrs
ge-0/0/4.45    BDR       0.0.0.0   10.50.250.6     10.50.250.5     1
ge-0/0/4.56    BDR       0.0.0.0   10.50.250.6     10.50.250.5     1
ge-0/0/4.35    DR        0.0.0.50  10.50.250.5     10.50.250.3     1
```

f. OSPF and OSPFv3 states on R6

The OSPF neighbor table on router R6 has only one neighbor in Full state. Although, the OSPF interface table contains four interfaces listed.

```
[edit]
lab@R6# run show ospf neighbor
Address      Interface      State      ID              Pri    Dead
10.50.0.9    ge-0/0/4.16    Full       10.50.250.1     128    34

[edit]
lab@R6# run show ospf interface
Interface      State      Area      DR ID           BDR ID          Nbrs
ge-0/0/4.16    DR        0.0.0.0   10.50.250.6     10.50.250.1     1
ge-0/0/4.26    DR        0.0.0.0   10.50.250.6     0.0.0.0         0
ge-0/0/4.46    DR        0.0.0.0   10.50.250.6     0.0.0.0         0
ge-0/0/4.56    DR        0.0.0.0   10.50.250.6     0.0.0.0         0
```

The OSPFv3 table has two interfaces in Full state. Whereas the interface table has four interfaces.

```
[edit]
lab@R6# run show ospf3 neighbor
ID            Interface      State      Pri    Dead
10.50.250.2    ge-0/0/4.26    Full       128    36
Neighbor-address fe80::206:a00:1a0e:204
10.50.250.5    ge-0/0/4.56    Full       128    34
Neighbor-address fe80::206:a00:380e:504

[edit]
lab@R6# run show ospf3 interface
Interface      State      Area      DR ID           BDR ID          Nbrs
ge-0/0/4.16    DR        0.0.0.0   10.50.250.6     0.0.0.0         0
```


ge-0/0/4.26	DR	0.0.0.0	10.50.250.6	10.50.250.2	1
ge-0/0/4.46	DR	0.0.0.0	10.50.250.6	0.0.0.0	0
ge-0/0/4.56	DR	0.0.0.0	10.50.250.6	10.50.250.5	1

g. OSPF and OSPFv3 states on R7

Similar to router R3, R7 is also an internal router in Area 50. Although, router R7 is expected to have two OSPF neighbors, currently only one is in Full state. The OSPF interface table contains two interfaces. The State of the interface to R3 could be a clue why the adjacency is not established. You might expect to have either DR or BDR state on broadcast interface with only two routers. A state different than DR and BDR should attract your attention for further examination.

```
[edit]
lab@R7# run show ospf neighbor
Address          Interface      State      ID          Pri    Dead
10.50.0.93       ge-0/0/4.47   Full       10.50.250.6 128    33
```

```
[edit]
lab@R7# run show ospf interface
Interface      State  Area      DR ID          BDR ID          Nbrs
ge-0/0/4.37    DOther 0.0.0.50   0.0.0.0        0.0.0.0         0
ge-0/0/4.47    DR     0.0.0.50   10.50.250.7    10.50.250.6     1
```

The OSPFv3 neighbor table on R7 is completely empty and only one interface is listed in the interface table.

```
[edit]
lab@R7# run show ospf3 neighbor
```

```
[edit]
lab@R7# run show ospf3 interface
Interface      State  Area      DR ID          BDR ID          Nbrs
ge-0/0/4.37    DR     0.0.0.50   10.50.250.7    0.0.0.0         0
```

h. OSPF and OSPFv3 states on R8

R8 has not enabled neither of the two protocols in the initial configuration.

```
[edit]
lab@R8# run show ospf neighbor
OSPF instance is not running
```

```
[edit]
lab@R8# run show ospf3 neighbor
OSPF instance is not running
```

i. Fixing neighbor adjacencies on R1

Since the interface to R6 is missing from the OSPFv3 interface table it is worth checking the configuration concerning OSPFv3.

```
[edit]
lab@R1# show protocols ospf3
ospf3 {
    area 0.0.0.0 {
        interface ge-0/0/4.12;
        interface ge-0/0/4.14;
        interface ge-0/0/4.16;
    }
}
```

It looks like interface `ge-0/0/4.16` is present in the configuration. It is worth verifying if the interface is configured for IPv6 traffic. The output below confirms that family `inet6` is missing from the interface configuration.

```
[edit]
lab@R1# run show interfaces ge-0/0/4.16
Logical interface ge-0/0/4.16 (Index 334) (SNMP ifIndex 533)
  Description: === connection to R6 ===
  Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.16 ] Encapsulation: ENET2
  Input packets : 87
  Output packets: 193
  Protocol inet, MTU: 1500
    Flags: Sendbroadcast-pkt-to-re
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.50.0.8/30, Local: 10.50.0.9, Broadcast: 10.50.0.11
  Protocol multiservice, MTU: Unlimited

[edit]
lab@R1# show interfaces ge-0/0/4.16
description "=== connection to R6 ===";
vlan-id 16;
family inet {
    address 10.50.0.9/30;
}
```

Adding the missing family `inet6` enables the IPv6 protocol on the interface. This will assign a link-local IPv6 address on the interface, which is enough for successful OSPFv3 communication. OSPFv3 is using the link-local addresses to establish adjacencies between the neighbors. After the configuration change, all three neighbors reached Full state.

```
[edit]
lab@R1# set interfaces ge-0/0/4 unit 16 family inet6

[edit]
lab@R1# run show ospf3 neighbor
ID                Interface          State    Pri    Dead
10.50.250.2       ge-0/0/4.12         Full     128    39
  Neighbor-address fe80::206:a00:c0e:204
10.50.250.6       ge-0/0/4.14         Full     128    31
  Neighbor-address fe80::206:a00:e0e:404
10.50.250.6       ge-0/0/4.16         Full     128    35
  Neighbor-address fe80::206:a00:100e:604
```

j. Fixing neighbor adjacencies on R2

The OSPF neighbor table on R2 contains the interface to router R6 but the adjacency is missing from the neighbor table. Enabling the traceoptions for OSPF protocol is helpful in most of the cases to provide a clue where might be the cause of the problem.

```
[edit]
lab@R2# set protocols ospf traceoptions file OSPF_log

[edit]
lab@R2# set protocols ospf traceoptions flag hello

[edit]
lab@R2# set protocols ospf traceoptions flag error

[edit]
lab@R2# commit
commit complete
```

The output of the log file clearly reveals that the subnet mask on the router R2 interface does not match the mask configured on router R6. IP addresses from the same subnet is a mandatory requirement for successful OSPF communication.

```
[Edit]
lab@R2# run show log OSPF_log
Nov 30 11:07:36 trace_on: Tracing to "/var/log/OSPF_log" started

Nov 30 11:07:36.833737 OSPF sent Hello 10.50.0.17 -> 224.0.0.5 (ge-0/0/4.26 IFL
334 area 0.0.0.0)
Nov 30 11:07:36.833746 Version 2, length 44, ID 10.50.250.2, area 0.0.0.0
Nov 30 11:07:36.833751 mask 255.255.255.254, hello_ivl 10, opts 0x12, prio 128
Nov 30 11:07:36.833755 dead_ivl 40, DR 10.50.0.17, BDR 0.0.0.0
Nov 30 11:07:41.017774 OSPF packet ignored: subnet mismatch from 10.50.0.18 on
intf ge-0/0/4.26 area 0.0.0.0
Nov 30 11:07:41.017819 OSPF rcvd Hello 10.50.0.18 -> 224.0.0.5 (ge-0/0/4.26 IFL
334 area 0.0.0.0)
Nov 30 11:07:41.017825 Version 2, length 44, ID 10.50.250.6, area 0.0.0.0
Nov 30 11:07:41.017830 checksum 0x0, authtype 0
Nov 30 11:07:41.017835 mask 255.255.255.252, hello_ivl 10, opts 0x12, prio 128
Nov 30 11:07:41.017840 dead_ivl 40, DR 10.50.0.18, BDR 0.0.0.0
```

Listing the interface configuration shows that the subnet mask of the interface on R2 to R6 is indeed configured to /31 instead of /30.

```
[edit]
lab@R2# show interfaces ge-0/0/4.26
description "=== connection to R6 ===";
vlan-id 26;
family inet {
    address 10.50.0.17/31;
}
family inet6;
```

To amend the configuration with correct interface subnet mask you can use the rename command. It will help to do the change with single command which eventually could save valuable time on the real exam.

```
[edit]
lab@R2# rename interfaces ge-0/0/4.26 family inet address 10.50.0.17/31 to
address 10.50.0.17/30

[edit]
lab@R2# commit
commit complete
```

The change is enough to bring up the adjacency between R2 and R6.

```
[edit]
lab@R2# run show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.50.0.1	ge-0/0/4.12	Full	10.50.250.1	128	37
10.50.0.14	ge-0/0/4.24	Full	10.50.250.6	128	34
10.50.0.18	ge-0/0/4.26	Full	10.50.250.6	128	32

k. Fixing neighbor adjacencies on R3

The fact that R3 has at least two interfaces enabled for OSPF but not even one working adjacency suggests a problem with the OSPF configuration. Again, use traceoptions to find where the problem might be.

```
[edit]
lab@R3# set protocols ospf traceoptions file OSPF_log
```

```
[edit]
lab@R3# set protocols ospf traceoptions flag hello
```

```
[edit]
lab@R3# set protocols ospf traceoptions flag error
```

```
[edit]
lab@R3# commit
commit complete
```

The output of the file shows that R3 is rejecting the hello packets from routers R5 and R7 because of an Area mismatch error. The log message also reveals which node is configured incorrectly. In this case, the local device R3 is configured with Area ID number of 0.0.5.0 instead of 0.0.0.50.

```
[edit]
lab@R3# run show log OSPF_log
Nov 30 11:39:06 trace on: Tracing to "/var/log/OSPF_log" started
Nov 30 11:39:13.011008 OSPF packet ignored: area mismatch (0.0.0.50) from
10.50.0.90 on intf ge-0/0/4.37 area 0.0.5.0
Nov 30 11:39:13.494019 OSPF packet ignored: area mismatch (0.0.0.50) from
10.50.0.86 on intf ge-0/0/4.35 area 0.0.5.0
Nov 30 11:39:13.494045 OSPF rcvd Hello 10.50.0.86 -> 224.0.0.5 (ge-0/0/4.35 IFL
332 area 0.0.5.0)
Nov 30 11:39:13.494050 Version 2, length 44, ID 10.50.250.5, area 0.0.0.50
Nov 30 11:39:13.494080 checksum 0xdc31, authtype 0
Nov 30 11:39:13.494089 mask 255.255.255.252, hello_ivl 10, opts 0x12, prio 128
Nov 30 11:39:13.494095 dead_ivl 40, DR 10.50.0.86, BDR 0.0.0.0
```

```
[edit]
lab@R3# rename protocols ospf area 0.0.5.0 to area 50
```

```
[edit]
lab@R3# commit
commit complete
```

After fixing the Area mismatch on R3, the neighbors with router R5 and R7 appear in the OSPF neighbor table.

```
[edit]
lab@R3# run show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.50.0.86	ge-0/0/4.35	Full	10.50.250.5	128	29
10.50.0.90	ge-0/0/4.37	2Way	10.50.250.7	0	37

Checking carefully the output above you should notice that neighbor with R7 is not in Full state. You may also notice that the Priority on the interface to R7 is 0, which means that R3 cannot be Designated or Backup designated router for the particular segment. Nevertheless, checking the table on the other end reveals that the Priority of the interface to R3 is also 0. In that case neither one of both neighbors could be Designated or Backup Designated router. Therefore, the State they reached is 2Way. This confirms the two routers see each other's Hello message, however they cannot exchange updates in that State.

```
[edit]
lab@R7# run show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.50.0.89	ge-0/0/4.37	2Way	10.50.250.3	0	39
10.50.0.93	ge-0/0/4.47	Full	10.50.250.6	128	33

```
[edit]
```

```

lab@R3# delete protocols ospf area 50 interface ge-0/0/4.37 priority
[edit]
lab@R7# delete protocols ospf area 50 interface ge-0/0/4.37 priority
[edit]
lab@R7# commit
commit complete

```

After deleting the interface priority configuration on routers R3 and R7, the two neighbors reach Full State.

```

[edit]
lab@R3# run show ospf neighbor

```

Address	Interface	State	ID	Pri	Dead
10.50.0.86	ge-0/0/4.35	Full	10.50.250.5	128	25
10.50.0.90	ge-0/0/4.37	Full	10.50.250.7	128	33

Because the aggregated ethernet interface was not part of the initial configuration, you have to add it in the appropriate Area under the OSPF configuration on both routers R3 and R4.

```

[edit]
lab@R3# set protocols ospf area 50 interface ae3.0
[edit]
lab@R3# commit
commit complete

[edit]
lab@R4# set protocols ospf area 50 interface ae3.0
[edit]
lab@R4# commit
commit complete

```

Full State for all neighbors on R3.

```

[edit]
lab@R3# run show ospf neighbor

```

Address	Interface	State	ID	Pri	Dead
10.50.0.82	ae3.0	Full	10.50.250.4	128	33
10.50.0.86	ge-0/0/4.35	Full	10.50.250.5	128	33
10.50.0.90	ge-0/0/4.37	Full	10.50.250.7	128	34

Initially, the OSPFv3 neighbor table on router R3 has only single neighbor in Full state with router R5.

Here again enabling traceoptions might help find where the problem could be.

```

[edit]
lab@R3# set protocols ospf3 traceoptions file OSPF3_log
[edit]
lab@R3# set protocols ospf3 traceoptions flag hello detail
[edit]
lab@R3# set protocols ospf3 traceoptions flag error detail
[edit]
lab@R3# commit
commit complete

```

The log file suggest that hello mismatch prevents the R3 to accept the hello packets from R7.

```
[edit]
lab@R3# run show log OSPF3_log
Nov 30 12:00:19 trace_on: Tracing to "/var/log/OSPF3_log" started
Nov 30 12:00:44.325798 OSPF rcvd Hello fe80::206:a00:250e:704 -> ff02::5 (ge-
0/0/4.37 IFL 333 area 0.0.0.50)
Nov 30 12:00:44.325882 Version 3, length 36, ID 10.50.250.7, area 0.0.0.50
Nov 30 12:00:44.325887 checksum 0x0, instance-id 0
Nov 30 12:00:44.325894 intf_id 0.0.0.1, prio 128, opts 0x13, hello_ivl 10
Nov 30 12:00:44.325899 dead_ivl 40, DR 10.50.250.7, BDR 0.0.0.0
Nov 30 12:00:44.325916 OSPF packet ignored: hello interval mismatch 10 from
fe80::206:a00:250e:704 on intf ge-0/0/4.37 area 0.0.0.50
Nov 30 12:00:45.141567 OSPF periodic xmit from fe80::206:a00:230e:304 to ff02::5
(IFL 332 area 0.0.0.50)
Nov 30 12:00:51.918673 Received OSPF packet of type and wire_length 1, 36
Nov 30 12:00:51.918832 OSPF rcvd Hello fe80::206:a00:250e:704 -> ff02::5 (ge-
0/0/4.37 IFL 333 area 0.0.0.50)
Nov 30 12:00:51.918840 Version 3, length 36, ID 10.50.250.7, area 0.0.0.50
Nov 30 12:00:51.918844 checksum 0x0, instance-id 0
Nov 30 12:00:51.918851 intf_id 0.0.0.1, prio 128, opts 0x13, hello_ivl 10
Nov 30 12:00:51.918855 dead_ivl 40, DR 10.50.250.7, BDR 0.0.0.0
```

Verify the configuration on router R3, clearly shows that the hello and dead intervals are set manually to value different than default.

```
[edit]
lab@R3# show protocols ospf3
area 0.0.0.50 {
    interface ge-0/0/4.35;
    interface ge-0/0/4.37 {
        hello-interval 40;
        dead-interval 200;
    }
}
```

Simply deleting the initially configured hello and dead intervals will allow successful neighbor hello packet exchange.

```
[edit]
lab@R3# delete protocols ospf3 area 50 interface ge-0/0/4.37 hello-interval

[edit]
lab@R3# delete protocols ospf3 area 50 interface ge-0/0/4.37 dead-interval

[edit]
lab@R3# commit
commit complete
```

The OSPFv3 neighbor with R7 successfully reaches Full State.

```
[edit]
lab@R3# run show ospf3 neighbor
```

ID	Interface	State	Pri	Dead
10.50.250.5	ge-0/0/4.35	Full	128	29
Neighbor-address fe80::206:a00:230e:504				
10.50.250.7	ge-0/0/4.37	Full	128	34
Neighbor-address fe80::206:a00:250e:704				

Although the aggregated ethernet interface is added to the OSPFv3 configuration, the neighbor do not appear in the neighbor table. Looking again in the traceoptions log file reveals that the OSPFv3 hello packets are ignored because of Area stubness mismatch.

```
[edit]
lab@R3# run show log OSPF3_log
Nov 30 12:32:41.657654 OSPF rcvd Hello fe80::205:86ff:fea7:dec0 -> ff02::5 (ae3.0
IFL 338 area 0.0.0.50)
Nov 30 12:32:41.657659 Version 3, length 36, ID 10.50.250.4, area 0.0.0.50
Nov 30 12:32:41.657663 checksum 0x0, instance-id 0
Nov 30 12:32:41.657668 intf_id 0.0.0.6, prio 128, opts 0x11, hello_ivl 10
Nov 30 12:32:41.657672 dead_ivl 40, DR 10.50.250.4, BDR 0.0.0.0
Nov 30 12:32:41.657690 OSPF packet ignored: area stubness mismatch from
fe80::205:86ff:fea7:dec0 on intf ae3.0 area 0.0.0.50
```

```
[edit]
lab@R3# show protocols ospf3
area 0.0.0.0 {
    interface ge-0/0/4.14;
    interface ge-0/0/4.24;
    interface ge-0/0/4.45;
    interface ge-0/0/4.46;
}
area 0.0.0.50 {
    stub;
    interface ge-0/0/4.47;
}
```

Listing the initial configuration on both routers clearly shows that Area 50 on router R4 is configured as a Stub Area. Because at this point it is not known if Area 50 should be configured as Stub Area, deleting the Stub configuration on R4 is the right approach clearing the mismatch.

```
[edit]
lab@R4# delete protocols ospf3 area 0.0.0.50 stub
```

This helps to achieve Full State to all three neighbors of R3.

```
[edit]
lab@R3# run show ospf3 neighbor
ID          Interface          State    Pri    Dead
10.50.250.4 ae3.0                Full     128    34
Neighbor-address fe80::205:86ff:fea7:dec0
10.50.250.5 ge-0/0/4.35          Full     128    28
Neighbor-address fe80::206:a00:230e:504
10.50.250.7 ge-0/0/4.37          Full     128    33
Neighbor-address fe80::206:a00:250e:704
```

I. Fixing neighbor adjacencies on R4

Router R4 OSPF neighbor table has four neighbors in the Full State. After adding the missing ae3 interface between R3 and R4 in the previous section, only one more neighbor is missing. Again enabling traceoptions will help you to find the issue.

```
[edit]
lab@R4# set protocols ospf traceoptions file OSPF_log

[edit]
lab@R4# set protocols ospf traceoptions flag hello detail

[edit]
lab@R4# set protocols ospf traceoptions flag error detail

[edit]
lab@R4# commit
commit complete

[edit]
```

```

lab@R4# run show log OSPF_log
Nov 30 11:35:02 trace_on: Tracing to "/var/log/OSPF_log" started
Nov 30 11:35:22.649287 OSPF rcvd Hello 10.50.0.38 -> 224.0.0.5 (ge-0/0/4.46 IFL
335 area 0.0.0.0)
Nov 30 11:35:22.649296 Version 2, length 44, ID 10.50.250.6, area 0.0.0.0
Nov 30 11:35:22.649300 checksum 0x0, authtype 0
Nov 30 11:35:22.649305 mask 255.255.255.252, hello_ivl 10, opts 0x12, prio 128
Nov 30 11:35:22.649309 dead_ivl 40, DR 10.50.0.38, BDR 0.0.0.0
Nov 30 11:35:22.649324 OSPF restart signaling: Received hello with LLS data from
nbr ip=10.50.0.38 id=10.50.250.6.
Nov 30 11:35:22.649330 OSPF packet ignored: our router ID received from
10.50.0.38 on intf ge-0/0/4.46 area 0.0.0.0

```

The log file shows that the OSPF packets from R6 are rejected due to duplicated router IDs. This can also be the reason why R1 and R2 have the R6's ID instead of the R4 ID in their OSPF neighbor table.

```

[edit]
lab@R4# show interfaces lo0.0
family inet {
    address 10.50.250.6/32;
}
family inet6 {
    address 2001:abcd::4/128;
}

```

Indeed, initially R4 is configured by mistake with the R6 router ID.

```

[edit]
lab@R4# rename interfaces lo0.0 family inet address 10.50.250.6/32 to address
10.50.250.4/32

```

```

[edit]
lab@R4# commit
commit complete

```

Renaming the correct IP address on lo0.0 interface brings up the adjacencies between R4 and R6.

```

[edit]
lab@R4# run show ospf neighbor

```

Address	Interface	State	ID	Pri	Dead
10.50.0.5	ge-0/0/4.14	Full	10.50.250.1	128	35
10.50.0.13	ge-0/0/4.24	Full	10.50.250.2	128	36
10.50.0.34	ge-0/0/4.45	Full	10.50.250.5	128	32
10.50.0.38	ge-0/0/4.46	Full	10.50.250.6	128	39
10.50.0.81	ae3.0	Full	10.50.250.3	128	38
10.50.0.94	ge-0/0/4.47	Full	10.50.250.7	128	37

The OSPFv3 neighbor table on R4 is missing router R7. This will be fixed later when router R7 configuration is analyzed.

```

[edit]
lab@R4# run show ospf3 neighbor

```

ID	Interface	State	Pri	Dead
10.50.250.1	ge-0/0/4.14	Full	128	38
Neighbor-address fe80::206:a00:e0e:104				
10.50.250.2	ge-0/0/4.24	Full	128	38
Neighbor-address fe80::206:a00:180e:204				
10.50.250.5	ge-0/0/4.45	Full	128	29
Neighbor-address fe80::206:a00:2d0e:504				
10.50.250.6	ge-0/0/4.46	Full	128	39
Neighbor-address fe80::206:a00:2e0e:604				
10.50.250.3	ae3.0	Full	128	38


```
Neighbor-address fe80::205:86ff:fe49:bac0
```

m. Fixing neighbor adjacencies on R5

At this stage, R5 should have R3 and R4 in OSPF Full State. If you recall earlier the table showed only the neighbor to R4 but with the router ID of R6. This is already fixed after the correct router R4 ID was set to the correct value. Verifying the current OSPF neighbor table, reveals that only the neighbor with router R6 is missing.

```
[edit]
lab@R5# run show ospf neighbor
Address          Interface      State      ID           Pri    Dead
10.50.0.33       ge-0/0/4.45   Full       10.50.250.4  128    39
10.50.0.85       ge-0/0/4.35   Full       10.50.250.3  128    37
```

Enabling the traceoptions will give a clue why neighbor to R6 is not present in the OSPF neighbor table.

```
[edit]
lab@R5# set protocols ospf traceoptions file OSPF_log

[edit]
lab@R5# set protocols ospf traceoptions flag hello detail

[edit]
lab@R5# set protocols ospf traceoptions flag error detail

[edit]
lab@R5# commit
commit complete

[edit]
lab@R5# run show log OSPF_log
Nov 30 11:20:48 trace_on: Tracing to "/var/log/OSPF_log" started
Nov 30 11:20:48.210935 OSPF sent Hello 10.50.0.45 -> 224.0.0.5 (ge-0/0/4.56 IFL
334 area 0.0.0.0)
Nov 30 11:20:48.210940 Version 2, length 44, ID 10.50.250.5, area 0.0.0.0
Nov 30 11:20:48.210945 mask 255.255.255.252, hello_ivl 10, opts 0x12, prio 128
Nov 30 11:20:48.210949 dead_ivl 40, DR 10.50.0.45, BDR 0.0.0.0
Nov 30 11:20:48.210952 OSPF restart signaling: Add LLS data for Hello packet on
interface ge-0/0/4.56.
Nov 30 11:20:48.957121 Received OSPF packet of type and wire_length 1, 60
Nov 30 11:20:48.957179 OSPF rcvd Hello 10.50.0.85 -> 224.0.0.5 (ge-0/0/4.35 IFL
336 area 0.0.0.50)
Nov 30 11:20:48.957184 Version 2, length 48, ID 10.50.250.3, area 0.0.0.50
Nov 30 11:20:48.957187 checksum 0x0, authtype 0
Nov 30 11:20:48.957190 mask 255.255.255.252, hello_ivl 10, opts 0x12, prio 128
Nov 30 11:20:48.957195 dead_ivl 40, DR 10.50.0.86, BDR 10.50.0.85
Nov 30 11:20:48.957208 OSPF restart signaling: Received hello with LLS data from
nbr ip=10.50.0.85 id=10.50.250.3.
Nov 30 11:20:48.957215 OSPF restart signaling: Received hello with LR bit set
from nbr ip=10.50.0.85 id=10.50.250.3. Set oob-resync capability 1.
Nov 30 11:20:48.957279 OSPF restart signaling:Save packet length 60 :
Nov 30 11:20:52.292890 OSPF packet ignored: authentication failure from
10.50.0.46
Nov 30 11:20:53.448169 Received OSPF packet of type and wire_length 1, 60
```

It looks like the authentication configured between R5 and R6 is failing. Checking the initial configuration reveals, that simple password is used for authentication between the routers. However, JUNOS for security reason is displaying the password key hashed when the configuration is listed.

One way to find the configured passwords is looking at the hello packet headers monitoring the traffic to and from the router. During the actual JNCIE-SP exam this is the only available way since Internet access is not allowed to use an online tool to reversing the hash.

```
[edit]
lab@R5# run monitor traffic interface ge-0/0/4.56 detail no-resolve size 1500
Address resolution is OFF.
Listening on ge-0/0/4.56, capture size 1500 bytes

11:22:30.022443 In IP6 (class 0xc0, hlim 1, next-header: OSPF (89), length: 40)
fe80::206:a00:380e:604 > ff02::5: OSPFv3, Hello, length 40
  Router-ID 10.50.250.6, Backbone Area
  Options [V6, External, Router]
    Hello Timer 10s, Dead Timer 40s, Interface-ID 0.0.0.4, Priority 128
    Designated Router 10.50.250.6, Backup Designated Router 10.50.250.5
    Neighbor List:
      10.50.250.5
11:22:31.755798 In IP (tos 0xc0, ttl 1, id 13253, offset 0, flags [none],
proto: OSPF (89), length: 76) 10.50.0.46 > 224.0.0.5: OSPFv2, Hello, length 56
[ len 44]
  Router-ID 10.50.250.6, Backbone Area, Authentication Type: simple (1)
  Simple text password: JNCIE-SP
  Options [External, LLS]
    Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.252, Priority 128
    Designated Router 10.50.0.46
    LLS: checksum: 0xffff6, length: 3
    Extended Options (1), length: 4
      Options: 0x00000001 [LSDB resync]
11:22:33.283024 Out IP (tos 0xc0, ttl 1, id 5552, offset 0, flags [none],
proto: OSPF (89), length: 76) 10.50.0.45 > 224.0.0.5: OSPFv2, Hello, length 56
[ len 44]
  Router-ID 10.50.250.5, Backbone Area, Authentication Type: simple (1)
  Simple text password: jncie-sp
  Options [External, LLS]
    Hello Timer 10s, Dead Timer 40s, Mask 255.255.255.252, Priority 128
    Designated Router 10.50.0.45
    LLS: checksum: 0xffff6, length: 3
    Extended Options (1), length: 4
      Options: 0x00000001 [LSDB resync]
```

The output proves that the password keys configured on R5 and R6 are different. Setting one of the two passwords on both routers brings the adjacency to Full State.

```
[edit]
lab@R5# set protocols ospf area 0 interface ge-0/0/4.56 authentication simple-
password JNCIE-SP

[edit]
lab@R5# commit
commit complete

[edit]
lab@R6# set protocols ospf area 0 interface ge-0/0/4.56 authentication simple-
password JNCIE-SP

[edit]
lab@R6# commit
commit complete
```

```
[edit]
lab@R5# run show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.50.0.33	ge-0/0/4.45	Full	10.50.250.4	128	39
10.50.0.46	ge-0/0/4.56	Full	10.50.250.6	128	38
10.50.0.85	ge-0/0/4.35	Full	10.50.250.3	128	37

The OSPFv3 neighbor table has all neighbors in Full State.

```
[edit]
lab@R5# run show ospf3 neighbor
ID          Interface          State    Pri    Dead
10.50.250.4 ge-0/0/4.45             Full     128    39
  Neighbor-address fe80::206:a00:2d0e:404
10.50.250.6 ge-0/0/4.56             Full     128    39
  Neighbor-address fe80::206:a00:380e:604
10.50.250.3 ge-0/0/4.35             Full     128    35
  Neighbor-address fe80::206:a00:230e:304
```

n. Fixing neighbor adjacencies on R6

At this point all R6 adjacencies are in Full state in the OSPF and OSPFv3 neighbor tables.

```
[edit]
lab@R6# run show ospf neighbor
Address      Interface          State    ID          Pri    Dead
10.50.0.9    ge-0/0/4.16       Full     10.50.250.1  128    39
10.50.0.17   ge-0/0/4.26       Full     10.50.250.2  128    35
10.50.0.37   ge-0/0/4.46       Full     10.50.250.4  128    37
10.50.0.45   ge-0/0/4.56       Full     10.50.250.5  128    35
```

```
[edit]
lab@R6# run show ospf3 neighbor
ID          Interface          State    Pri    Dead
10.50.250.1 ge-0/0/4.16       Full     128    34
  Neighbor-address fe80::206:a00:100e:104
10.50.250.2 ge-0/0/4.26       Full     128    34
  Neighbor-address fe80::206:a00:1a0e:204
10.50.250.4 ge-0/0/4.46       Full     128    32
  Neighbor-address fe80::206:a00:2e0e:404
10.50.250.5 ge-0/0/4.56       Full     128    38
  Neighbor-address fe80::206:a00:380e:504
```

o. Fixing neighbor adjacencies on R7

As found early on during troubleshooting the R4 adjacencies, the OSPFv3 neighbor between R4 and R7 was completely missing from the neighbor table on R4. Therefore, router R7 neighbor table does not have router R4 as a neighbor also. Simply verifying the initial OSPFv3 configuration on R7 reveals that the interface to R4 is missing from the ospf3 configuration.

```
[edit]
lab@R7# run show ospf3 neighbor
ID          Interface          State    Pri    Dead
10.50.250.3 ge-0/0/4.37       Full     128    39
  Neighbor-address fe80::206:a00:250e:304
```

```
[edit]
lab@R7# show protocols ospf3
area 0.0.0.50 {
  interface ge-0/0/4.37;
}
```

```
[edit]
lab@R7# set protocols ospf3 area 50 interface ge-0/0/4.47
```

```
[edit]
lab@R7# commit
commit complete
```

Adding the interface under Area 50 in protocols ospf3 stanza fixes the adjacency.

```
[edit]
lab@R7# run show ospf3 neighbor
ID          Interface          State    Pri    Dead
10.50.250.3 ge-0/0/4.37             Full     128    34
  Neighbor-address fe80::206:a00:250e:304
10.50.250.4 ge-0/0/4.47             Full     128    33
  Neighbor-address fe80::206:a00:2f0e:404
```

After this step, the adjacencies on all routers, except on R8, should be in Full State.

- Area 0 routers should not exchange any Network LSAs.

The Network LSA or LSA Type 2 is advertised only by the Designated nodes in a broadcast and NBMA segments. The task requires eliminating the Network LSAs from Area 0. Since it is not explicitly specified, you have to assume that both OSPF and OSPFv3 protocol should be configured. Because the OSPF database is the same on all routers across the same Area, listing the content of the OSPF database on one routers is enough to check for Network LSAs.

```
[edit]
lab@R1# run show ospf database

    OSPF database, Area 0.0.0.0
Type  ID          Adv Rtr          Seq          Age    Opt  Cksum  Len
Router *10.50.250.1  10.50.250.1     0x800000011  962    0x22  0xbc4e  60
Router 10.50.250.2  10.50.250.2     0x800000011  2491   0x22  0xdedb  60
Router 10.50.250.4  10.50.250.4     0x800000008  484    0x22  0xfcfc  72
Router 10.50.250.5  10.50.250.5     0x80000000d  1985   0x22  0x33d9  48
Router 10.50.250.6  10.50.250.6     0x80002481   2613   0x22  0x5a9b  72
Network *10.50.0.1  10.50.250.1     0x80000000a  1712   0x22  0x1e2f  32
Network *10.50.0.5  10.50.250.1     0x80000000c  212    0x22  0xe37   32
Network *10.50.0.9  10.50.250.1     0x800000009  2462   0x22  0x83a   32
Network 10.50.0.13  10.50.250.2     0x80000000a  2491   0x22  0xc577  32
Network 10.50.0.18  10.50.250.6     0x80000519   13     0x22  0x5ac3  32
Network 10.50.0.34  10.50.250.5     0x800000005  1443   0x22  0x91e   32
Network 10.50.0.38  10.50.250.6     0x800000005  880    0x22  0xe43c  32
Network 10.50.0.46  10.50.250.6     0x800000003  1751   0x22  0xa673  32
Summary 10.50.0.80     10.50.250.4     0x800000006  3551   0x22  0x96db  28
Summary 10.50.0.80     10.50.250.5     0x800000005  3524   0x22  0x9cd4  28
Summary 10.50.0.84     10.50.250.4     0x800000006  747    0x22  0x78f4  28
Summary 10.50.0.84     10.50.250.5     0x80000000a  2909   0x22  0x6009  28
Summary 10.50.0.88     10.50.250.4     0x800000006  220    0x22  0x5019  28
Summary 10.50.0.88     10.50.250.5     0x800000006  2293   0x22  0x4a1e  28
Summary 10.50.0.92     10.50.250.4     0x800000006  2326   0x22  0x1e48  28
Summary 10.50.0.92     10.50.250.5     0x800000007  1150   0x22  0x2a38  28
```

```
[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    10.50.250.1     0x80000001f  1913   0x6a5e  72
Router 0.0.0.0     10.50.250.2     0x800000020  1957   0x4e4d  72
Router 0.0.0.0     10.50.250.4     0x80000001e  1713   0x83e6  88
Router 0.0.0.0     10.50.250.5     0x80000001b  3150   0xa36e  56
Router 0.0.0.0     10.50.250.6     0x80002481   1477   0x6d44  88
Network *0.0.0.1    10.50.250.1     0x800000010  713    0x52fc  32
Network 0.0.0.3     10.50.250.2     0x80000000f  1358   0x7ccb  32
Network 0.0.0.1     10.50.250.4     0x800000016  662    0x44ff  32
Network 0.0.0.2     10.50.250.4     0x800000015  3425   0x4af8  32
Network 0.0.0.3     10.50.250.4     0x80000001a  2283   0x60d9  32
Network 0.0.0.1     10.50.250.6     0x800011b9   1905   0xd1b9  32
Network 0.0.0.3     10.50.250.6     0x8000134b   1049   0xbe34  32
```

<http://www.inetzero.com> - Copyright 2017 iNET ZERO. All rights reserved

For personal non commercial use only – do not distribute

<https://t.me/learningnets>

```

Network      0.0.0.4          10.50.250.6          0x80001458  2553  0xa53d  32

    OSPF3 Link-Local database, interface ge-0/0/4.12 Area 0.0.0.0
Type      ID          Adv Rtr          Seq          Age  Cksum  Len
Link      *0.0.0.1      10.50.250.1      0x80000011    1313  0xb0eb  44
Link      0.0.0.1      10.50.250.2      0x80000010    759   0xc1d9  44

    OSPF3 Link-Local database, interface ge-0/0/4.14 Area 0.0.0.0
Type      ID          Adv Rtr          Seq          Age  Cksum  Len
Link      *0.0.0.2      10.50.250.1      0x80000012    2513  0xcacd  44
Link      0.0.0.1      10.50.250.4      0x8000000b    856   0x108a  44

    OSPF3 Link-Local database, interface ge-0/0/4.16 Area 0.0.0.0
Type      ID          Adv Rtr          Seq          Age  Cksum  Len
Link      *0.0.0.3      10.50.250.1      0x8000000f    113   0xecab  44
Link      0.0.0.1      10.50.250.6      0x8000000f    2382  0x4c44  44

```

As highlighted in the output of the OSPF database, Network LSA are exchanged by the routers in Area0. This is due the fact that all links are by default treated as broadcast on the Ethernet interfaces.

NOTE: The LSAs marked with ‘*’ in front are advertised by the local device where the output is taken.

a. All devices with interface in Area 0

To complete the task, you have make sure that LSA type 2 are not sent in Area 0. This means that you should not have DR nodes in any of the segments in Area 0. The only way to accomplish that is to change the type of the OSPF interface to point-to-point link, this way a DR is not elected. Because you have only two devices per segment on all links, you are safe to do that without the risk of isolating any of the routes from forming OSPF adjacencies. Do not forget to make the change also for OSPFv3 as well.

```

[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.12 interface-type p2p

[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.14 interface-type p2p

[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.16 interface-type p2p

[edit]
lab@R1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.12 interface-type p2p

[edit]
lab@R1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.14 interface-type p2p

[edit]
lab@R1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.16 interface-type p2p

```

b. Verify the configuration

After committing the change and verifying the OSPF database again, the Network LSAs are not advertised by the routers in Area 0.

```

[edit]
lab@R1# run show ospf database

    OSPF database, Area 0.0.0.0
Type      ID          Adv Rtr          Seq          Age  Opt  Cksum  Len

```

```

Router *10.50.250.1      10.50.250.1      0x8000001c      15  0x22  0x8b56  108
Router 10.50.250.2      10.50.250.2      0x8000001e      16  0x22  0x6bf8  108
Router 10.50.250.4      10.50.250.4      0x80000014      16  0x22  0x136a  132
Router 10.50.250.5      10.50.250.5      0x80000018      17  0x22  0xeb3b   84
Router 10.50.250.6      10.50.250.6      0x8000248d      16  0x22  0x95c4  132
Summary 10.50.0.80      10.50.250.4      0x8000000b     2957 0x22  0x8ce0   28
Summary 10.50.0.80      10.50.250.5      0x8000000a     2936 0x22  0x92d9   28
Summary 10.50.0.84      10.50.250.4      0x8000000b      346 0x22  0x6ef9   28
Summary 10.50.0.84      10.50.250.5      0x8000000f     2320 0x22  0x560e   28
Summary 10.50.0.88      10.50.250.4      0x8000000b      130 0x22  0x461e   28
Summary 10.50.0.88      10.50.250.5      0x8000000b     1704 0x22  0x4023   28
Summary 10.50.0.92      10.50.250.4      0x8000000b     1733 0x22  0x144d   28
Summary 10.50.0.92      10.50.250.5      0x8000000c      169 0x22  0x203d   28

```

[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      10.50.250.1      0x80000022    141  0x2bf   72
Router    0.0.0.0      10.50.250.2      0x80000025    143  0x721e   72
Router    0.0.0.0      10.50.250.4      0x80000024    142  0xaca4   88
Router    0.0.0.0      10.50.250.5      0x8000001f    143  0x6e95   56
Router    0.0.0.0      10.50.250.6      0x80002486    142  0x4966   88

```

```

OSPF3 Link-Local database, interface ge-0/0/4.12 Area 0.0.0.0
Type      ID          Adv Rtr          Seq          Age  Cksum  Len
Link      *0.0.0.1      10.50.250.1      0x80000001    342  0xd0db   44
Link      0.0.0.1      10.50.250.2      0x80000001    249  0xdfca   44

```

```

OSPF3 Link-Local database, interface ge-0/0/4.14 Area 0.0.0.0
Type      ID          Adv Rtr          Seq          Age  Cksum  Len
Link      *0.0.0.2      10.50.250.1      0x80000001    342  0xecbc   44
Link      0.0.0.7      10.50.250.4      0x80000001    202  0xe7b6   44

```

```

OSPF3 Link-Local database, interface ge-0/0/4.16 Area 0.0.0.0
Type      ID          Adv Rtr          Seq          Age  Cksum  Len
Link      *0.0.0.3      10.50.250.1      0x80000001    342  0x99d    44
Link      0.0.0.1      10.50.250.6      0x80000001    142  0x6836   44

```

- Ensure that all routers in Area 0 can reach the loopback address of all other routers in the Area. Make sure that they are advertised as internal OSPF routes.

There are two ways in JUNOS for advertising prefixes in OSPF. You can export the routes using routing policy or you can enable the OSPF protocol over the interface in the Area where the routes should be advertised. The first method advertises the prefixes as external in the OSPF database. Whereas the second method inserts the routes as internal in the particular Area. Since, the way routes are advertised in OSPF plays a significant role in the algorithm for selecting the best routes, it is very important to know how the routes enter the OSPF database. This particular task requires Loopback IP addresses to be advertised as internal. This means you have to add the Loopback interfaces under the Area 0 configuration. It is worth mentioning that in this case OSPF will try to establish adjacencies over the Loopback interfaces. The task does not require that but on the actual exam, you might be asked to disable sending hello message on interfaces enabled for OSPF. Configuring the interfaces as 'passive' will prevent the interfaces configured for OSPF to establish adjacencies.

- a. All devices in Area0

[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 lo0.0

[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 lo0.0

```

b. Verify the configuration

Configuring the Loopback interfaces for OSPF and OSPFv3 on all routers in Area 0 will advertise them as internal OSPF routes.

```

[edit]
lab@R1# run show route 10.50.250.0/24

inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.1/32      *[Direct/0] 12:28:27
                   > via lo0.0
10.50.250.2/32      *[OSPF/10] 00:42:40, metric 1
                   > to 10.50.0.2 via ge-0/0/4.12
10.50.250.4/32      *[OSPF/10] 00:40:12, metric 1
                   > to 10.50.0.6 via ge-0/0/4.14
10.50.250.5/32      *[OSPF/10] 00:29:01, metric 2
                   > to 10.50.0.6 via ge-0/0/4.14
                   > to 10.50.0.10 via ge-0/0/4.16
10.50.250.6/32      *[OSPF/10] 00:29:01, metric 1
                   > to 10.50.0.10 via ge-0/0/4.16

[edit]
lab@R1# run show route table inet6.0 2001:abcd::/64

inet6.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:abcd::1/128    *[Direct/0] 00:26:05
                   > via lo0.0
2001:abcd::2/128    *[OSPF3/10] 00:00:24, metric 1
                   > to fe80::206:a00:c0e:204 via ge-0/0/4.12
2001:abcd::4/128    *[OSPF3/10] 00:00:21, metric 1
                   > to fe80::206:a00:e0e:404 via ge-0/0/4.14
2001:abcd::5/128    *[OSPF3/10] 00:00:16, metric 2
                   > to fe80::206:a00:e0e:404 via ge-0/0/4.14
                   > to fe80::206:a00:100e:604 via ge-0/0/4.16
2001:abcd::6/128    *[OSPF3/10] 00:00:14, metric 1
                   > to fe80::206:a00:100e:604 via ge-0/0/4.16

```

Solution - Task 2.2: RIP redistribution

- 2) Make sure that routers R4 and R7 are able to receive RIP and RIPng routes from DC-1. Advertise the most specific summary route representing these routes to the OSPF domain.

In this Task, you are asked to configure RIP and RIPng on R4 and R7 exchanging routes with DC router. The received routes should be further advertised into OSPF and OSPFv3 as summary routes.

- a. R4 – IPv4 addressing

First, you have to configure the interface to DC according to the diagram. Do not forget to enable the interface for vlan-tagging otherwise the commit will fail.

```
[edit]
lab@R4# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R4# set interfaces ge-0/0/5 unit 200 description "=== connection to DC ==="

[edit]
lab@R4# set interfaces ge-0/0/5 unit 200 vlan-id 200

[edit]
lab@R4# set interfaces ge-0/0/5 unit 200 family inet address 10.50.200.4/24
```

Since RIPng is using link-local IPv6 addresses to exchange routing information, enabling the interface for IPv6 addressing is enough for successful RIPng communication.

```
[edit]
lab@R4# set interfaces ge-0/0/5 unit 200 family inet6
```

RIP neighbor configuration must be configured under the group stanza.

```
[edit]
lab@R4# set protocols rip group DC-1 neighbor ge-0/0/5.200
```

An aggregate route is needed to summarize the routes received by RIP and advertised further into OSPF.

```
[edit]
lab@R4# set routing-options aggregate route 172.31.0.0/20
```

Routing policies are used in JUNOS for manipulation routing information to and from the routing table. The policy used for exporting the IPv4 aggregate route is configured and applied to OSPF.

```
[edit]
lab@R4# set policy-options policy-statement to.ospf term 1 from protocol
aggregate

[edit]
lab@R4# set policy-options policy-statement to.ospf term 1 from route-filter
172.31.0.0/20 exact

[edit]
lab@R4# set policy-options policy-statement to.ospf term 1 then accept

[edit]
```



```
lab@R4# set policy-options policy-statement to.ospf term explicit.reject then
reject
```

```
[edit]
lab@R4# set protocols ospf export to.ospf
```

```
[edit]
lab@R4# commit
commit complete
```

b. R7 – IPv4 addressing

The same configuration should be applied to router R7 also.

```
[edit]
lab@R7# set interfaces ge-0/0/5 unit 200 description "=== connection to DC ==="
```

```
[edit]
lab@R7# set interfaces ge-0/0/5 unit 200 vlan-id 200
```

```
[edit]
lab@R7# set interfaces ge-0/0/5 unit 200 family inet address 10.50.200.7/24
```

```
[edit]
lab@R7# set interfaces ge-0/0/5 unit 200 family inet6
```

```
[edit]
lab@R7# set protocols rip group DC-1 neighbor ge-0/0/5.200
```

```
[edit]
lab@R7# set routing-options aggregate route 172.31.0.0/20
```

```
[edit]
lab@R7# set policy-options policy-statement to.ospf term 1 from protocol
aggregate
```

```
[edit]
lab@R7# set policy-options policy-statement to.ospf term 1 from route-filter
172.31.0.0/20 exact
```

```
[edit]
lab@R7# set policy-options policy-statement to.ospf term 1 then accept
```

```
[edit]
lab@R7# set policy-options policy-statement to.ospf term explicit.reject then
reject
```

```
[edit]
lab@R7# set protocols ospf export to.ospf
```

```
[edit]
lab@R7# commit
commit complete
```

c. C4 – IPv6 addressing

```
[edit]
lab@R4# set routing-options rib inet6.0 aggregate route 2007:ab12::/36
```

```
[edit]
lab@R4# set policy-options policy-statement to.ospf3.ipv6 term 1 from protocol
aggregate
```

```
[edit]
lab@R4# set policy-options policy-statement to.ospf3.ipv6 term 1 from route-
filter 2007:ab12::/36 exact
```

```
[edit]
lab@R4# set policy-options policy-statement to.ospf3.ipv6 term 1 then accept

[edit]
lab@R4# set policy-options policy-statement to.ospf3.ipv6 term explicit.reject
then reject

[edit]
lab@R4# set protocols ospf3 export to.ospf3.ipv6
```

d. R7 – IPv6 addressing

Similar configuration should be applied for the IPv6 addressing. A care should be taken with the aggregate IPv6 address. The correct summary prefix must be configured to cover all IPv6 prefixes received by DC via RIPv6.

```
[edit]
lab@R7# set routing-options rib inet6.0 aggregate route 2007:ab12::/36

[edit]
lab@R7# set policy-options policy-statement to.ospf3.ipv6 term 1 from protocol
aggregate
[edit]
lab@R7# set policy-options policy-statement to.ospf3.ipv6 term 1 from route-
filter 2007:ab12::/36 exact

[edit]
lab@R7# set policy-options policy-statement to.ospf3.ipv6 term 1 then accept

[edit]
lab@R7# set policy-options policy-statement to.ospf3.ipv6 term explicit.reject
then reject

[edit]
lab@R7# set protocols ospf3 export to.ospf3.ipv6
```

e. Verifying IPv4 configuration

RIP neighbors are successfully established and routes are received on R4 and R7 from DC. JUNOS by default accepts RIP routes without the need of routing policy explicitly configured.

```
[edit]
lab@R4# run show rip neighbor
```

Neighbor	Local State	Source Address	Destination Address	Send Mode	Receive Mode	In Met
ge-0/0/5.200	Up	10.50.200.4	224.0.0.9	mcast	both	1

```
[edit]
lab@R7# run show rip neighbor
```

Neighbor	Local State	Source Address	Destination Address	Send Mode	Receive Mode	In Met
ge-0/0/5.200	Up	10.50.200.7	224.0.0.9	mcast	both	1

```
[edit]
lab@R4# run show route protocol rip
```

```
inet.0: 45 destinations, 45 routes (45 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/24      *[RIP/100] 00:00:16, metric 2, tag 0
                  > to 10.50.200.200 via ge-0/0/5.200
172.31.1.0/24      *[RIP/100] 00:00:16, metric 2, tag 0
```

```

> to 10.50.200.200 via ge-0/0/5.200
172.31.2.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.3.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.4.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.5.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.6.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.7.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.8.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.9.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.10.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.11.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.12.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.13.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.14.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
172.31.15.0/24 * [RIP/100] 00:00:16, metric 2, tag 0
> to 10.50.200.200 via ge-0/0/5.200
224.0.0.9/32 * [RIP/100] 00:00:16, metric 1
MultiRecv

```

Verify that the summary route is advertised in the OSPF domain and received by the other devices.

```

[edit]
lab@R1# run show route 172.31.0.0/20

inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/20 * [OSPF/150] 00:02:22, metric 0, tag 0
> to 10.50.0.6 via ge-0/0/4.14

```

f. Verifying the IPv6 configuration

The IPv6 summary route is successfully advertised to the OSPFv3 protocol.

```

[edit]
lab@R7# run show ripng neighbor

Neighbor      State      Source      Dest      In
-----      -
Address      Address    Send Recv  Met
-----      -
ge-0/0/5.200  Up         fe80::206:a00:c80e:705  ff02::9  yes yes  1

[edit]
lab@R4# run show route protocol ripng

inet.0: 46 destinations, 51 routes (46 active, 0 holddown, 0 hidden)

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

inet6.0: 31 destinations, 31 routes (31 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2007:ab12::/64 * [RIPng/100] 00:00:16, metric 2, tag 0
> to fe80::206:a00:c80e:908 via ge-0/0/5.200

```

```

2007:ab12:1::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:2::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:3::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:4::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:5::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:6::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:7::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:8::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:9::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:a::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:b::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:c::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:d::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:e::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
2007:ab12:f::/64    *[RIPng/100] 00:00:16, metric 2, tag 0
                   > to fe80::206:a00:c80e:908 via ge-0/0/5.200
ff02::9/128        *[RIPng/100] 00:00:24, metric 1
                   MultiRecv

[edit]
lab@R3# run show route 2007:ab12::/36

inet6.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2007:ab12::/36      *[OSPF3/150] 00:00:08, metric 0, tag 0
                   to fe80::206:a00:250e:704 via ge-0/0/4.37
                   > to fe80::205:86ff:fea7:dec0 via ae3.0

[edit]
lab@R1# run show route 2007:ab12::/36

inet6.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2007:ab12::/36      *[OSPF3/150] 01:17:08, metric 100, tag 0
                   > to fe80::206:a00:e0e:404 via ge-0/0/4.14

```

3) DC-1 must have reachability to the loopback addresses of all routers in the OSPF domain.

a. R4 – IPv4 addressing

In the reverse direction to the DC router only the Loopback IP addresses should be advertised. To export routes into RIP, you must use explicit routing-policy.

Below, the use of the 'orlonger' option in the route-filter statement ensures that only the more specific routes in the range of the configured subnet are matched. In this particular case, this includes all of Loopback IP addresses individually.

```
[edit]
lab@R4# set policy-options policy-statement to.rip term 1 from route-filter
10.50.250.0/28 orlonger
```

```
[edit]
lab@R4# set policy-options policy-statement to.rip term 1 then accept
```

Although not explicitly required, a reject term is configured to prevent accidental leakage of additional prefixes.

```
[edit]
lab@R4# set policy-options policy-statement to.rip term explicit.reject then
reject
```

```
[edit]
lab@R4# set protocols rip group DC-1 export to.rip
```

b. R7 – IPv4 addressing

The same configuration is applied to router R7 as well.

```
[edit]
lab@R7# set protocols rip group DC-1 export to.rip
```

```
[edit]
lab@R7# set policy-options policy-statement to.rip term 1 from route-filter
10.50.250.0/28 orlonger
```

```
[edit]
lab@R7# set policy-options policy-statement to.rip term 1 then accept
```

```
[edit]
lab@R7# set policy-options policy-statement to.rip term explicit.reject then
reject
```

c. R4 – IPv6 addressing

```
[edit]
lab@R4# set protocols ripng group DC-1 export to.rip.ipv6
```

```
[edit]
lab@R4# set policy-options policy-statement to.rip.ipv6 term 1 from route-filter
2001:abcd::/64 orlonger
```

```
[edit]
lab@R4# set policy-options policy-statement to.rip.ipv6 term 1 then accept
```

```
[edit]
lab@R4# set policy-options policy-statement to.rip.ipv6 term explicit.reject then
reject
```

d. R7 – IPv6 addressing

```
[edit]
lab@R7# set protocols ripng group DC-1 export to.rip.ipv6
```

```
[edit]
lab@R7# set policy-options policy-statement to.rip.ipv6 term 1 from route-filter
2001:abcd::/64 orlonger
```

```
[edit]
lab@R7# set policy-options policy-statement to.rip.ipv6 term 1 then accept
```

```
[edit]
lab@R7# set policy-options policy-statement to.rip.ipv6 term explicit.reject then
reject
```

e. Verify IPv4 configuration

The use of the command below, lists the routes advertised by router R4 to DC router. Unlike with BGP, you must use the local IP address to list the advertised routes by the local router to a RIP neighbor.

```
[edit]
lab@R4# run show route advertising-protocol rip 10.50.200.4

inet.0: 45 destinations, 45 routes (45 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.1/32    *[OSPF/10] 00:24:36, metric 1
                  > to 10.50.0.5 via ge-0/0/4.14
10.50.250.2/32    *[OSPF/10] 00:24:31, metric 1
                  > to 10.50.0.13 via ge-0/0/4.24
10.50.250.4/32    *[Direct/0] 00:24:46
                  > via lo0.0
10.50.250.5/32    *[OSPF/10] 00:23:37, metric 1
                  > to 10.50.0.34 via ge-0/0/4.45
10.50.250.6/32    *[OSPF/10] 00:23:32, metric 1
                  > to 10.50.0.38 via ge-0/0/4.46

[edit]
lab@R7# run show route advertising-protocol rip 10.50.200.7

inet.0: 42 destinations, 47 routes (42 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.1/32    *[OSPF/10] 00:40:52, metric 2
                  > to 10.50.0.93 via ge-0/0/4.47
                  [RIP/100] 00:05:54, metric 2, tag 0
                  > to 10.50.200.4 via ge-0/0/5.200
10.50.250.2/32    *[OSPF/10] 00:40:52, metric 2
                  > to 10.50.0.93 via ge-0/0/4.47
                  [RIP/100] 00:05:54, metric 2, tag 0
                  > to 10.50.200.4 via ge-0/0/5.200
10.50.250.4/32    *[OSPF/10] 00:40:52, metric 1
                  > to 10.50.0.93 via ge-0/0/4.47
                  [RIP/100] 00:05:54, metric 2, tag 0
                  > to 10.50.200.4 via ge-0/0/5.200
10.50.250.5/32    *[OSPF/10] 00:40:52, metric 2
                  > to 10.50.0.89 via ge-0/0/4.37
                  > to 10.50.0.93 via ge-0/0/4.47
                  [RIP/100] 00:05:54, metric 2, tag 0
                  > to 10.50.200.4 via ge-0/0/5.200
10.50.250.6/32    *[OSPF/10] 00:40:52, metric 2
                  > to 10.50.0.93 via ge-0/0/4.47
                  [RIP/100] 00:05:54, metric 2, tag 0
                  > to 10.50.200.4 via ge-0/0/5.200
10.50.250.7/32    *[Direct/0] 00:41:42
                  > via lo0.0
```

f. Verify IPv6 configuration

The RIPng uses the same approach but because link-local addresses are used by OSPFv3, you need to use the 'show interfaces' output to find them.

```
[edit]
lab@R4# run show route advertising-protocol ripng fe80::206:a00:c80e:405

inet6.0: 31 destinations, 31 routes (31 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```

2001:abcd::1/128    *[OSPF3/10] 00:34:48, metric 1
                   > to fe80::206:a00:e0e:104 via ge-0/0/4.14
2001:abcd::2/128    *[OSPF3/10] 00:34:45, metric 1
                   > to fe80::206:a00:180e:204 via ge-0/0/4.24
2001:abcd::4/128    *[Direct/0] 00:48:49
                   > via lo0.0
2001:abcd::5/128    *[OSPF3/10] 00:34:40, metric 1
                   > to fe80::206:a00:2d0e:504 via ge-0/0/4.45
2001:abcd::6/128    *[OSPF3/10] 00:34:37, metric 1
                   > to fe80::206:a00:2e0e:604 via ge-0/0/4.46

[edit]
lab@R7# run show route advertising-protocol ripng fe80::206:a00:c80e:705

inet6.0: 28 destinations, 33 routes (28 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:abcd::1/128    *[OSPF3/10] 00:38:50, metric 2
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
                   [RIPng/100] 00:00:10, metric 2, tag 0
                   > to fe80::206:a00:c80e:405 via ge-0/0/5.200
2001:abcd::2/128    *[OSPF3/10] 00:38:47, metric 2
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
                   [RIPng/100] 00:00:10, metric 2, tag 0
                   > to fe80::206:a00:c80e:405 via ge-0/0/5.200
2001:abcd::4/128    *[OSPF3/10] 00:38:37, metric 1
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
                   [RIPng/100] 00:00:10, metric 2, tag 0
                   > to fe80::206:a00:c80e:405 via ge-0/0/5.200
2001:abcd::5/128    *[OSPF3/10] 00:38:37, metric 2
                   > to fe80::206:a00:250e:304 via ge-0/0/4.37
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
                   [RIPng/100] 00:00:10, metric 2, tag 0
                   > to fe80::206:a00:c80e:405 via ge-0/0/5.200
2001:abcd::6/128    *[OSPF3/10] 00:38:37, metric 2
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
                   [RIPng/100] 00:00:10, metric 2, tag 0
                   > to fe80::206:a00:c80e:405 via ge-0/0/5.200
2001:abcd::7/128    *[Direct/0] 00:52:25
                   > via lo0.0

```

- 4) Avoid changing the default preference value of any routing protocol if you need to resolve a routing loop or suboptimal routing.

a. R4 – IPv4 addressing

To prevent non-optimal routing in the network you should reject the OSPF routes received back from the DC through RIP. At this point, it is not obvious that such problem exists but later when R8's Loopback IPv4 address is advertised it will appear as external in the routing table on routers R7 and R4 with preference of 150. When advertised back from the DC as RIP route will make that path preferred. The task restricts you to manipulate the route preference, so the other option you have in JUNOS involves assigning a routing tag to the routes advertised to RIP.

```

[edit]
lab@R4# set policy-options policy-statement to.rip term 1 then tag 200

```

To reject routes based on route tags you have to configure an import policy.

```

[edit]
lab@R4# set policy-options policy-statement from.rip term 1 from protocol rip

[edit]

```

```

lab@R4# set policy-options policy-statement from.rip term 1 from tag 200

[edit]
lab@R4# set policy-options policy-statement from.rip term 1 then reject

[edit]
lab@R4# set policy-options policy-statement from.rip term 2 from protocol rip

[edit]
lab@R4# set policy-options policy-statement from.rip term 2 then accept

[edit]
lab@R4# set protocols rip group DC-1 import from.rip

[edit]
lab@R4# commit
commit complete

```

b. R7 – IPv4 addressing

The same configuration is applied to R7 router.

```

[edit]
lab@R7# set policy-options policy-statement to.rip term 1 then tag 200

[edit]
lab@R7# set policy-options policy-statement from.rip term 1 from protocol rip

[edit]
lab@R7# set policy-options policy-statement from.rip term 1 from tag 200

[edit]
lab@R7# set policy-options policy-statement from.rip term 1 then reject

[edit]
lab@R7# set policy-options policy-statement from.rip term 2 from protocol rip

[edit]
lab@R7# set policy-options policy-statement from.rip term 2 then accept

[edit]
lab@R7# set protocols rip group DC-1 import from.rip

[edit]
lab@R7# commit
commit complete

```

c. R4 – IPv6 addressing

Similar concept is used for IPv6 addressing.

```

[edit]
lab@R4# set policy-options policy-statement to.rip.ipv6 term 1 then tag 200

[edit]
lab@R4# set policy-options policy-statement from.rip.ipv6 term 1 from protocol
ripng

[edit]
lab@R4# set policy-options policy-statement from.rip.ipv6 term 1 from tag 200

[edit]
lab@R4# set policy-options policy-statement from.rip.ipv6 term 1 then reject

[edit]
lab@R4# set policy-options policy-statement from.rip.ipv6 term 2 from protocol
ripng

```



```
[edit]
lab@R4# set policy-options policy-statement from.rip.ipv6 term 2 then accept

[edit]
lab@R4# set protocols ripng group DC-1 import from.rip.ipv6

[edit]
lab@R4# commit
commit complete
```

d. R7 – IPv6 addressing

```
[edit]
lab@R7# set policy-options policy-statement to.rip.ipv6 term 1 then tag 200

[edit]
lab@R7# set policy-options policy-statement from.rip.ipv6 term 1 from protocol
ripng

[edit]
lab@R7# set policy-options policy-statement from.rip.ipv6 term 1 from tag 200

[edit]
lab@R7# set policy-options policy-statement from.rip.ipv6 term 1 then reject

[edit]
lab@R7# set policy-options policy-statement from.rip.ipv6 term 2 from protocol
ripng

[edit]
lab@R7# set policy-options policy-statement from.rip.ipv6 term 2 then accept

[edit]
lab@R7# commit
commit complete
```

e. Verify the configuration

```
[edit]
lab@R7# run show route 10.50.250.0/28

inet.0: 37 destinations, 39 routes (37 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.1/32    *[RIP/100] 03:17:47, metric 2, tag 0
                 > to 10.50.200.4 via ge-0/0/5.200
10.50.250.2/32    *[RIP/100] 03:17:47, metric 2, tag 0
                 > to 10.50.200.4 via ge-0/0/5.200
10.50.250.3/32    *[OSPF/10] 02:05:25, metric 1
                 > to 10.50.0.89 via ge-0/0/4.37
                 [RIP/100] 02:02:58, metric 2, tag 0
                 > to 10.50.200.4 via ge-0/0/5.200
10.50.250.4/32    *[RIP/100] 03:17:47, metric 2, tag 0
                 > to 10.50.200.4 via ge-0/0/5.200
10.50.250.5/32    *[RIP/100] 03:17:47, metric 2, tag 0
                 > to 10.50.200.4 via ge-0/0/5.200
10.50.250.6/32    *[RIP/100] 03:17:47, metric 2, tag 0
                 > to 10.50.200.4 via ge-0/0/5.200
10.50.250.7/32    *[Direct/0] 03:53:35
                 > via lo0.0
                 [RIP/100] 02:02:58, metric 2, tag 0
                 > to 10.50.200.4 via ge-0/0/5.200

[edit]
lab@R7# run show route 10.50.250.0/28

inet.0: 33 destinations, 33 routes (33 active, 0 holddown, 0 hidden)
```

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

```
10.50.250.0/28      *[Aggregate/130] 00:01:40
                   > to 10.50.0.89 via ge-0/0/4.37
10.50.250.3/32     *[OSPF/10] 02:17:45, metric 1
                   > to 10.50.0.89 via ge-0/0/4.37
10.50.250.7/32     *[Direct/0] 04:05:55
                   > via lo0.0
```

[edit]

```
lab@R7# run show route table inet6.0 2001:abcd::/64
```

```
inet6.0: 55 destinations, 55 routes (55 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
2001:abcd::/64     *[Aggregate/130] 00:09:12
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::1/128   *[OSPF3/10] 00:09:03, metric 200
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::2/128   *[OSPF3/10] 00:09:03, metric 200
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::3/128   *[OSPF3/10] 02:15:24, metric 100
                   > to fe80::206:a00:250e:304 via ge-0/0/4.37
2001:abcd::4/128   *[OSPF3/10] 00:09:12, metric 100
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::5/128   *[OSPF3/10] 00:09:12, metric 200
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::6/128   *[OSPF3/10] 00:09:05, metric 200
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::7/128   *[Direct/0] 02:21:24
                   > via lo0.0
```

Solution - Task 2.3: Multi-area OSPF configuration

5) Configure a multi-area OSPF network as described in the diagram. Ensure that the following criteria are met:

- No Type 5 LSAs are allowed in Area 50. Ensure that only R7 is advertising a summary of the DC-1 routes as external in Area 50. The loopback addresses of R3 and R7 routers must appear as internal in Area 50.

The task requires to configure Area 50 in such a way that no LSA type 5 are exchanged within the Area. Type 5 LSAs are carrying external information. Because Area 50 has two ASBR routers which are ASBR routers, you have to configure the area as Not-so-stubby (NSSA) area. Inside the NSSA Area, the external routes are advertised as Type 7 LSAs.

a. R3 – IPv4 addressing

All routers internal to Area 50 are configured for NSSA.

```
[edit]
lab@R3# set protocols ospf area 0.0.0.50 interface lo0.0

[edit]
lab@R3# set protocols ospf area 0.0.0.50 nssa

[edit]
lab@R3# commit
commit complete
```

b. R7 – IPv4 addressing

```
[edit]
lab@R7# set protocols ospf area 0.0.0.50 interface lo0.0

[edit]
lab@R7# set protocols ospf area 0.0.0.50 nssa

[edit]
lab@R7# commit
commit complete
```

c. R4 – IPv4 addressing

```
[edit]
lab@R4# set protocols ospf area 0.0.0.50 nssa
```

The command below will prevent router R4 to advertise the summary route in Area 50, effectively leaving only R7 advertising it.

```
[edit]
lab@R4# set protocols ospf no-nssa-abr

[edit]
lab@R4# commit
commit complete
```

d. R5 – IPv4 addressing

```
[edit]
lab@R5# set protocols ospf area 0.0.0.50 nssa

[edit]
lab@R5# commit
commit complete
```

e. R3 – IPv6 addressing

The same configuration is applied to OSPFv3 regarding IPv6 addressing.

```
[edit]
lab@R3# set protocols ospf3 area 0.0.0.50 interface lo0.0
```

```
[edit]
lab@R3# set protocols ospf3 area 0.0.0.50 nssa
```

f. R7 – IPv6 addressing

```
[edit]
lab@R7# set protocols ospf3 area 0.0.0.50 interface lo0.0
```

```
[edit]
lab@R7# set protocols ospf3 area 0.0.0.50 nssa
```

g. R4 – IPv6 addressing

```
[edit]
lab@R4# set protocols ospf3 area 0.0.0.50 nssa
```

```
[edit]
lab@R4# set protocols ospf3 no-nssa-abr
```

h. R5 – IPv6 addressing

```
[edit]
lab@R5# set protocols ospf3 area 0.0.0.50 nssa
```

i. Verifying IPv4 configuration

```
[edit]
lab@R7# run show route 172.31.0.0/20

inet.0: 36 destinations, 38 routes (36 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/20      *[Aggregate/130] 00:47:59
                   Reject
```

```
[edit]
lab@R7# run show route 10.50.250.0/28

inet.0: 32 destinations, 32 routes (32 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.3/32    *[OSPF/10] 02:11:26, metric 1
                  > to 10.50.0.89 via ge-0/0/4.37
10.50.250.7/32    *[Direct/0] 03:59:36
                  > via lo0.0
```

```
[edit]
lab@R3# run show route 172.31.0.0/20

inet.0: 558 destinations, 1221 routes (554 active, 0 holddown, 8 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/20    *[OSPF/150] 01:13:46, metric 100, tag 0
                  > to 10.50.0.90 via ge-0/0/4.37
```

j. Verifying IPv6 configuration

```
[edit]
lab@R7# run show route 2007:ab12:f::/36

inet6.0: 30 destinations, 32 routes (30 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2007:ab12::/36      *[Aggregate/130] 00:35:55
                   Reject
```

```
[edit]
lab@R3# run show route 2007:ab12:f::/36

inet6.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2007:ab12::/36      *[OSPF3/150] 01:14:54, metric 0, tag 0
                   > to fe80::206:a00:250e:704 via ge-0/0/4.37
```

- To spread the traffic between both ABRs of Area 50, make sure that R5 is the preferred transit node for IPv4 traffic to the Backbone area, whereas IPv6 traffic prefers the path through R4. However, you must preserve full reachability in case one of the ABRs fail. Do not advertise anything apart from the IPv4 and IPv6 RIP prefixes summary routes and the R3 and R7 loopback addresses out of Area 50.

By default, the translation from Type 7 to Type 5 LSA is performed by the ABR with the lowest router ID. In this particular case, that role is played by router R4 for both IPv4 and IPv6. However, the task requires that R5 should be preferred gateway to and from Area 50. The steps below describe how this can be accomplished.

a. R4 – IPv4 addressing

The command below restricts the subnets on the interfaces in Area 50 to be advertised to Area 0.

```
[edit]
lab@R4# set protocols ospf area 0.0.0.50 area-range 10.50.0.80/28 restrict
```

The two commands below advertise R3 and R7 loopback IP addresses with higher metric to Area 0.

```
[edit]
lab@R4# set protocols ospf area 50 area-range 10.50.250.7/32 override-metric 500
```

```
[edit]
lab@R4# set protocols ospf area 50 area-range 10.50.250.3/32 override-metric 500
```

The next two commands are needed to make R4 less preferred than R5 for the IPv4 OSPF routes.

```
[edit]
lab@R4# set protocols ospf area 0 interface ge-0/0/4.14 metric 500
```

```
[edit]
lab@R4# set protocols ospf area 0 interface ge-0/0/4.24 metric 500
```

```
[edit]
lab@R4# commit
commit complete
```

b. R5 – IPv4 addressing

The two commands below advertise the external summary route advertised by router R5 as Type 1 route. This way the routes advertised by router R5 are more preferred than the route advertised by router R4 as default Type 2.

```
[edit]
lab@R5# set protocols ospf area 0.0.0.50 nssa area-range 172.31.0.0/20 override-
metric 1
```

```
[edit]
lab@R5# set protocols ospf area 0.0.0.50 nssa area-range 172.31.0.0/20 override-
metric metric-type 1
```

The command below restricts the subnets on the interfaces in Area 50 to be advertised to Area 0.

```
[edit]
lab@R5# set protocols ospf area 0.0.0.50 area-range 10.50.0.80/28 restrict
```

```
[edit]
lab@R5# commit
commit complete
```

c. R7 – IPv4 addressing

This command is required to make sure that R5 will prefer the RIP summary route from R7 instead the same route from R4.

```
[edit]
lab@R7# set policy-options policy-statement to.ospf term 1 then external type 1
```

d. R1 – IPv4 addressing

The command below is needed to make R4 less preferred than R5 for the IPv4 routes advertised in OSPF.

```
[edit]
lab@R1# set protocols ospf area 0 interface ge-0/0/4.14 metric 500
```

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

e. R2 – IPv4 addressing

The command below is needed to make R4 less preferred than R5 for the IPv4 routes advertised in OSPF.

```
[edit]
lab@R2# set protocols ospf area 0 interface ge-0/0/4.24 metric 500
```

```
[edit]
lab@R2# commit
commit complete
```

f. R4 – IPv6 addressing

The IPv6 traffic must prefer the path through router R4. Since R4 has the lower router ID compared to R5, you have to make sure only that locally advertised IPv6 routes are always preferred. Again making them Type 1 external is enough to accomplish this.

```
[edit]
lab@R4# set policy-options policy-statement to.ospf3.ipv6 term 1 then external
type 1
```

```
[edit]
lab@R4# commit
commit complete
```

g. Verifying the IPv4 configuration

Below outputs show the separation of the traffic as required by the Task. A few prefixes are selected to prove that.

```
[edit]
lab@R6# run show route 10.50.250.7
```

```
inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.50.250.7/32      *[OSPF/10] 00:00:14, metric 3
                   > to 10.50.0.45 via ge-0/0/4.56
```

```
[edit]
lab@R6# run show route 172.31.0.0/20
```

```
inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
172.31.0.0/20      *[OSPF/150] 00:00:18, metric 5, tag 0
                   > to 10.50.0.45 via ge-0/0/4.56
```

```
[edit]
lab@R6# run show route 10.50.0.88
```

```
inet.0: 555 destinations, 943 routes (551 active, 0 holddown, 5 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
[edit]
lab@R6# run show route 10.50.0.92
```

```
inet.0: 555 destinations, 943 routes (551 active, 0 holddown, 5 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
[edit]
lab@R1# run show route 10.50.250.7 extensive
```

```
inet.0: 547 destinations, 549 routes (547 active, 0 holddown, 1 hidden)
10.50.250.7/32 (1 entry, 1 announced)
    State: <FlashAll>
TSI:
KRT in-kernel 10.50.250.7/32 -> {10.50.0.10, 10.50.0.6}
    *OSPF    Preference: 10
             Next hop type: Router, Next hop index: 1048597
             Address: 0xa896410
             Next-hop reference count: 39
             Next hop: 10.50.0.10 via ge-0/0/4.16 weight 0x1, selected
             Session Id: 0x145
             Next hop: 10.50.0.6 via ge-0/0/4.14 weight 0xf000
             Session Id: 0x143
             State: <Active Int>
```

```

Local AS: 5673.873
Age: 2:32:36      Metric: 400
Validation State: unverified
Area: 0.0.0.0
Task: OSPF
Announcement bits (4): 0-KRT 5-LDP 7-Resolve tree 2 8-
Resolve_IGP_FRR task
AS path: I

[edit]
lab@R1# run show route 10.50.250.7

inet.0: 547 destinations, 549 routes (547 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.7/32      *[OSPF/10] 02:32:41, metric 400
                   > to 10.50.0.10 via ge-0/0/4.16
                   to 10.50.0.6 via ge-0/0/4.14

[edit]
lab@R1# run show route 172.31.0.0/20

inet.0: 547 destinations, 549 routes (547 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.31.0.0/20      *[OSPF/150] 02:33:00, metric 600, tag 0
                   > to 10.50.0.10 via ge-0/0/4.16
                   to 10.50.0.6 via ge-0/0/4.14

inet.3: 19 destinations, 25 routes (4 active, 0 holddown, 19 hidden)

[edit]
lab@R1# run show route 172.31.0.0/20 extensive

inet.0: 547 destinations, 549 routes (547 active, 0 holddown, 1 hidden)
172.31.0.0/20 (1 entry, 1 announced)
State: <FlashAll>

TSI:
KRT in-kernel 172.31.0.0/20 -> {10.50.0.10, 10.50.0.6}
*OSPF Preference: 150
Next hop type: Router, Next hop index: 1048597
Address: 0xa896410
Next-hop reference count: 39
Next hop: 10.50.0.10 via ge-0/0/4.16 weight 0x1, selected
Session Id: 0x145
Next hop: 10.50.0.6 via ge-0/0/4.14 weight 0xf000
Session Id: 0x143
State: <Active Int Ext>
Local AS: 5673.873
Age: 2:33:03      Metric: 600
Validation State: unverified
Tag: 0
Task: OSPF
Announcement bits (3): 0-KRT 5-LDP 7-Resolve tree 2
AS path: I

inet.3: 19 destinations, 25 routes (4 active, 0 holddown, 19 hidden)

```

h. Verifying the IPv6 configuration

Below outputs show the separation of the traffic as required by the task for some prefixes.

```

[edit]
lab@R1# run show route table inet6.0

```



```
inet6.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
2001:abcd::1/128    *[Direct/0] 03:19:15
                    > via lo0.0
2001:abcd::2/128    *[OSPF3/10] 02:53:34, metric 1
                    > to fe80::206:a00:c0e:204 via ge-0/0/4.12
2001:abcd::3/128    *[OSPF3/10] 00:06:59, metric 2
                    > to fe80::206:a00:e0e:404 via ge-0/0/4.14
2001:abcd::4/128    *[OSPF3/10] 02:53:31, metric 1
                    > to fe80::206:a00:e0e:404 via ge-0/0/4.14
2001:abcd::5/128    *[OSPF3/10] 02:53:26, metric 2
                    > to fe80::206:a00:e0e:404 via ge-0/0/4.14
                    > to fe80::206:a00:100e:604 via ge-0/0/4.16
2001:abcd::6/128    *[OSPF3/10] 02:53:24, metric 1
                    > to fe80::206:a00:100e:604 via ge-0/0/4.16
2001:abcd::7/128    *[OSPF3/10] 00:06:59, metric 2
                    > to fe80::206:a00:e0e:404 via ge-0/0/4.14
2007:ab12::/36      *[OSPF3/150] 00:49:13, metric 1, tag 0
                    > to fe80::206:a00:e0e:404 via ge-0/0/4.14
```

```
[edit]
lab@R6# run show route table inet6.0 2007:ab12::/36
```

```
inet6.0: 111 destinations, 111 routes (111 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
2007:ab12::/36      *[OSPF3/150] 02:25:57, metric 100, tag 0
                    > to fe80::206:a00:2e0e:404 via ge-0/0/4.46
```

- Keep the same traffic separation between R4 and R5 in direction to Area 50 from the backbone area.

In the opposite direction, both ABRs will advertise the routes from Area 0 to Area 50 by default. To accomplish traffic separation, you can make the less preferred ABR to advertise only default route to Area 50.

a. R4 – IPv4 addressing

For the IPv4 traffic router R4 should be the less preferred ABR, hence it will only send a default route with higher metric. You have to configure the default route as NSSA route.

```
[edit]
lab@R4# set protocols ospf area 0.0.0.50 nssa default-lsa default-metric 500
```

```
[edit]
lab@R4# set protocols ospf area 0.0.0.50 nssa default-lsa type-7
```

```
[edit]
lab@R4# set protocols ospf area 0.0.0.50 nssa no-summaries
```

b. R5 – IPv6 addressing

For IPv6 traffic, router R5 should be the less preferred ABR, hence it will only send a default route with a higher metric. You have to configure the default route as NSSA route.

```
[edit]
lab@R5# set protocols ospf3 area 0.0.0.50 nssa default-lsa default-metric 500
```

```
[edit]
lab@R5# set protocols ospf3 area 0.0.0.50 nssa default-lsa type-7
```

```
[edit]
lab@R5# set protocols ospf3 area 0.0.0.50 nssa no-summaries
```

c. Verify the IPv4 configuration

```
[edit]
lab@R7# run show route 10.50.250.0/24

inet.0: 50 destinations, 50 routes (50 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.0/28      *[Aggregate/130] 03:40:27
                   > to 10.50.0.89 via ge-0/0/4.37
10.50.250.1/32     *[OSPF/10] 00:52:59, metric 400
                   > to 10.50.0.89 via ge-0/0/4.37
10.50.250.2/32     *[OSPF/10] 00:52:59, metric 400
                   > to 10.50.0.89 via ge-0/0/4.37
10.50.250.3/32     *[OSPF/10] 00:52:59, metric 100
                   > to 10.50.0.89 via ge-0/0/4.37
10.50.250.4/32     *[OSPF/10] 00:52:59, metric 300
                   > to 10.50.0.89 via ge-0/0/4.37
10.50.250.5/32     *[OSPF/10] 00:52:59, metric 200
                   > to 10.50.0.89 via ge-0/0/4.37
10.50.250.6/32     *[OSPF/10] 00:52:59, metric 300
                   > to 10.50.0.89 via ge-0/0/4.37
```

The default IPv4 route is received by router R7 advertised from router R4.

```
[edit]
lab@R7# run show route 0/0 exact

inet.0: 58 destinations, 58 routes (57 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 02:16:51, metric 600, tag 0
                   > to 10.50.0.93 via ge-0/0/4.47
```

d. Verify the IPv6 configuration

```
[edit]
lab@R7# run show route table inet6.0

inet6.0: 33 destinations, 33 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:abcd::/64     *[Aggregate/130] 00:56:13
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::1/128   *[OSPF3/10] 00:56:13, metric 2
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::2/128   *[OSPF3/10] 00:56:13, metric 2
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::3/128   *[OSPF3/10] 00:56:13, metric 1
                   > to fe80::206:a00:250e:304 via ge-0/0/4.37
2001:abcd::4/128   *[OSPF3/10] 00:56:13, metric 1
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::5/128   *[OSPF3/10] 00:56:13, metric 2
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
2001:abcd::6/128   *[OSPF3/10] 00:56:13, metric 2
                   > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
```

The default IPv6 route is received by router R7 from router R5.

```
[edit]
lab@R7# run show route table inet6.0 ::/0 exact

inet6.0: 55 destinations, 55 routes (55 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::/0              *[OSPF3/150] 00:13:16, metric 700, tag 0
```

```
> to fe80::206:a00:250e:304 via ge-0/0/4.37
```

- Due to performance issues router R8 is restricted to run a single OSPF process. However, it should still be able to exchange IPv4 and IPv6 with the rest of the network.

Because you are restricted to only use a single OSPF process on R8, the only option that supports IPv4 and IPv6 is OSPFv3. Although OSPFv3 can be used for routing only IPv4 traffic, you still need to enable the interfaces for IPv6 routing. That is required because OSPFv3 is using IPv6 link-local addresses to establish adjacencies.

a. R5

Starting with configuring the interface to R8 on router R5 for routing IPv6 traffic.

```
[edit]
lab@R5# set interfaces ge-0/0/4.58 family inet6
```

The IPv4 configuration in OSPFv3 protocol is configured under the realm ipv4-unicast stanza. Apart from that, the rest of configuration statements are the same like in OSPF.

```
[edit]
lab@R5# set protocols ospf3 realm ipv4-unicast area 0.0.0.8 interface ge-0/0/4.58
```

The interface to R8 is also added to the IPv6 configuration.

```
[edit]
lab@R5# set protocols ospf3 area 0.0.0.8 interface ge-0/0/4.58
```

```
[edit]
lab@R5# commit
commit complete
```

b. R6

Similar configuration is added to router R6 as well.

```
[edit]
lab@R6# set interfaces ge-0/0/4.68 family inet6
```

```
[edit]
lab@R6# set protocols ospf3 realm ipv4-unicast area 0.0.0.8 interface ge-0/0/4.68
```

```
[edit]
lab@R6# set protocols ospf3 area 0.0.0.8 interface ge-0/0/4.68
```

```
[edit]
lab@R6# commit
commit complete
```

c. R8

Router R8

```
[edit]
lab@R8# set interfaces ge-0/0/4.58 family inet6
```

```
[edit]
lab@R8# set interfaces ge-0/0/4.68 family inet6
```

```
[edit]
lab@R8# set protocols ospf3 realm ipv4-unicast area 0.0.0.8 interface ge-0/0/4.58
```

```
[edit]
lab@R8# set protocols ospf3 realm ipv4-unicast area 0.0.0.8 interface ge-0/0/4.68

[edit]
lab@R8# set protocols ospf3 realm ipv4-unicast area 0.0.0.8 interface lo0.0

[edit]
lab@R8# set protocols ospf3 area 0.0.0.8 interface ge-0/0/4.58

[edit]
lab@R8# set protocols ospf3 area 0.0.0.8 interface ge-0/0/4.68

[edit]
lab@R8# set protocols ospf3 area 0.0.0.8 interface lo0.0

[edit]
lab@R8# commit
commit complete
```

d. Verify the configuration

Verify that router R8 has established adjacencies with both routers, R5 and R6.

```
[edit]
lab@R8# run show ospf3 neighbor
```

ID	Interface	State	Pri	Dead
10.50.250.5	ge-0/0/4.58	Full	128	34
Neighbor-address fe80::206:a00:3a0e:504				
10.50.250.6	ge-0/0/4.68	Full	128	33
Neighbor-address fe80::206:a00:440e:604				

```
[edit]
lab@R8# run show ospf3 interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
ge-0/0/4.58	DR	0.0.0.8	10.50.250.8	10.50.250.5	1
ge-0/0/4.68	DR	0.0.0.8	10.50.250.8	10.50.250.6	1
lo0.0	DR	0.0.0.8	10.50.250.8	0.0.0.0	0

- R8 must receive only the loopback address routes of the other routers in the network plus a summary route of DC-1. Due to limited size of the RIB, make sure that the IPv6 routes appear as summarized in the R8 routing table. Including the IPv6 loopback addresses.

a. R5 – IPv4 addressing

Because router R8 is using the OSPFv3 process only, you need to leak the OSPF IPv4 routes on routers R5 and R6 to the OSPFv3 process. You have to do this only for the Loopback IPv4 addresses and the summary route from DC-1.

```
[edit]
lab@R5# set policy-options policy-statement from.ospf.to.ospf3 term 1 from route-
filter 172.31.0.0/20 exact

[edit]
lab@R5# set policy-options policy-statement from.ospf.to.ospf3 term 1 from route-
filter 10.50.250.0/28 longer

[edit]
lab@R5# set policy-options policy-statement from.ospf.to.ospf3 term 1 then accept

[edit]
lab@R5# set policy-options policy-statement from.ospf.to.ospf3 term
explicit.reject then reject
```

```
[edit]
lab@R5# set protocols ospf3 realm ipv4-unicast export from.ospf.to.ospf3
```

b. R6 – IPv4 addressing

The same configuration is applied to R6 as well.

```
[edit]
lab@R6# set policy-options policy-statement from.ospf.to.ospf3 term 1 from route-
filter 172.31.0.0/20 exact
```

```
[edit]
lab@R6# set policy-options policy-statement from.ospf.to.ospf3 term 1 from route-
filter 10.50.250.0/28 longer
```

```
[edit]
lab@R6# set policy-options policy-statement from.ospf.to.ospf3 term 1 then accept
```

```
[edit]
lab@R6# set policy-options policy-statement from.ospf.to.ospf3 term
explicit.reject then reject
```

```
[edit]
lab@R6# set protocols ospf3 realm ipv4-unicast export from.ospf.to.ospf3
```

c. R5 – IPv6 configuration

For the IPv6 configuration, only a summary route must be advertised to router R8.

```
[edit]
lab@R5# set protocols ospf3 area 0.0.0.0 area-range 2001:abcd::/64
```

```
[edit]
lab@R5# commit
commit complete
```

d. R6 – IPv6 configuration

The same configuration is applied to R6 as well.

```
[edit]
lab@R6# set protocols ospf3 area 0.0.0.0 area-range 2001:abcd::/64
```

```
[edit]
lab@R6# commit
commit complete
```

e. Verify the configuration

IPv4 and IPv6 routes are received on router R8 as OSPFv3 routes.

```
[edit]
lab@R8# run show route

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.1.0/24      *[Direct/0] 04:11:54
                  > via fxp0.0
10.10.1.8/32     *[Local/0] 04:11:54
                  Local via fxp0.0
10.50.0.48/30    *[Direct/0] 01:47:41
```

```

> via ge-0/0/4.58
10.50.0.50/32 * [Local/0] 01:47:41
               Local via ge-0/0/4.58
10.50.0.52/30 * [Direct/0] 01:47:41
               > via ge-0/0/4.68
10.50.0.54/32 * [Local/0] 01:47:41
               Local via ge-0/0/4.68
10.50.250.1/32 * [OSPF3/150] 00:23:31, metric 1, tag 0
               > to 10.50.0.53 via ge-0/0/4.68
10.50.250.2/32 * [OSPF3/150] 00:00:58, metric 2, tag 0
               > to 10.50.0.49 via ge-0/0/4.58
               to 10.50.0.53 via ge-0/0/4.68
10.50.250.3/32 * [OSPF3/150] 00:00:58, metric 1, tag 0
               > to 10.50.0.49 via ge-0/0/4.58
10.50.250.4/32 * [OSPF3/150] 00:00:58, metric 1, tag 0
               > to 10.50.0.49 via ge-0/0/4.58
               to 10.50.0.53 via ge-0/0/4.68
10.50.250.5/32 * [OSPF3/150] 00:00:58, metric 0, tag 0
               > to 10.50.0.49 via ge-0/0/4.58
10.50.250.6/32 * [OSPF3/150] 00:23:31, metric 0, tag 0
               > to 10.50.0.53 via ge-0/0/4.68
10.50.250.7/32 * [OSPF3/150] 00:00:58, metric 2, tag 0
               > to 10.50.0.49 via ge-0/0/4.58
10.50.250.8/32 * [Direct/0] 01:47:41
               > via lo0.0
172.31.0.0/20 * [OSPF3/150] 00:00:58, metric 3, tag 0
               > to 10.50.0.49 via ge-0/0/4.58

[edit]
lab@R8# run show route table inet6.0

inet6.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:abcd::/64 * [OSPF3/10] 00:19:51, metric 3
               > to fe80::206:a00:3a0e:504 via ge-0/0/4.58
               to fe80::206:a00:440e:604 via ge-0/0/4.68
2001:abcd::8/128 * [Direct/0] 02:07:30
                 > via lo0.0
2007:ab12::/36 * [OSPF3/150] 02:07:04, metric 2, tag 0
               to fe80::206:a00:3a0e:504 via ge-0/0/4.58
               > to fe80::206:a00:440e:604 via ge-0/0/4.68
fe80::206:a00:3a0e:804/128
               * [Local/0] 02:07:19
               Local via ge-0/0/4.58
fe80::206:a00:440e:804/128
               * [Local/0] 02:07:19
               Local via ge-0/0/4.68
fe80::a00:dd0f:fcc0:801/128
               * [Direct/0] 02:07:30
               > via lo0.0
ff02::5/128 * [OSPF3/10] 02:07:30, metric 1
             MultiRecv

```

- Make sure that rest of the network can reach the R8 loopback IPv4 and IPv6 addresses.

a. R5

Similarly, routers R5 and R6 should explicitly advertise the IPv4 Loopback address to the OSPF process in Area 0.

```

[edit]
lab@R5# set policy-options policy-statement from.ospf3.to.ospf term 1 from
protocol ospf3

```

```
[edit]
lab@R5# set policy-options policy-statement from.ospf3.to.ospf term 1 from route-
filter 10.50.250.8/32 exact

[edit]
lab@R5# set policy-options policy-statement from.ospf3.to.ospf term 1 then accept

[edit]
lab@R5# set policy-options policy-statement from.ospf3.to.ospf term
explicit.reject then reject

[edit]
lab@R5# set protocols ospf export from.ospf3.to.ospf

[edit]
lab@R5# commit
commit complete
```

b. R6

The same configuration is applied to R6 as well.

```
[edit]
lab@R6# set policy-options policy-statement from.ospf3.to.ospf term 1 from
protocol ospf3

[edit]
lab@R6# set policy-options policy-statement from.ospf3.to.ospf term 1 from route-
filter 10.50.250.8/32 exact

[edit]
lab@R6# set policy-options policy-statement from.ospf3.to.ospf term 1 then accept

[edit]
lab@R6# set policy-options policy-statement from.ospf3.to.ospf term
explicit.reject then reject

[edit]
lab@R6# set protocols ospf export from.ospf3.to.ospf

[edit]
lab@R6# commit
commit complete
```

c. Verify the configuration

The IPv4 Loopback address of router R8 is learned as an external OSPF route in Area 0 and Area 50.

```
[edit]
lab@R1# run show route 10.50.250.8/32

inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.8/32      *[OSPF/150] 00:00:15, metric 2, tag 0
                   > to 10.50.0.10 via ge-0/0/4.16

[edit]
lab@R7# run show route 10.50.250.8/32

inet.0: 47 destinations, 47 routes (47 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.8/32      *[OSPF/150] 00:01:03, metric 1, tag 0
```

```
> to 10.50.0.89 via ge-0/0/4.37
```

The IPv6 Loopback address of router R8 is learned as intra-area address in OSPFv3.

```
[edit]
lab@R1# run show route 2001:abcd::8/128

inet6.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:abcd::8/128    *[OSPF3/10] 00:04:16, metric 2
                  > to fe80::206:a00:100e:604 via ge-0/0/4.16

[edit]
lab@R7# run show route 2001:abcd::8/128

inet6.0: 55 destinations, 55 routes (55 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:abcd::8/128    *[OSPF3/10] 00:33:38, metric 300
                  > to fe80::206:a00:2f0e:404 via ge-0/0/4.47
```

Solution - Task 2.4: Advanced OSPF configuration

- 6) Ensure all links with a bandwidth of 1Gbps have an OSPF cost of 100 and those with a bandwidth of 2Gbps have an OSPF cost of 50. Do not configure the cost manually on the interfaces.

a. All devices

By default, OSPF generates the cost of the link as a reference of 100mbps interface. This means all interfaces with speed higher than that 100Mbps will have a metric of 1. To accomplish the required by the Task cost values, a reference of 100gbps should be used instead.

```
[edit]
lab@R1# set protocols ospf reference-bandwidth 100g

[edit]
lab@R1# set protocols ospf3 reference-bandwidth 100g

[edit]
lab@R5# set protocols ospf3 realm ipv4-unicast reference-bandwidth 100g

[edit]
lab@R5# set protocols ospf3 reference-bandwidth 100g
```

b. Verify the configuration

The configuration achieved the desired cost set to the interfaces.

```
[edit]
lab@R3# run show ospf interface ae3.0 extensive
```

Interface	State	Area	DR ID	BDR ID	Nbrs
ae3.0	BDR	0.0.0.50	10.50.250.4	10.50.250.3	1

```

Type: LAN, Address: 10.50.0.81, Mask: 255.255.255.252, MTU: 1500, Cost: 50
DR addr: 10.50.0.82, BDR addr: 10.50.0.81, Priority: 128
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Stub NSSA
Auth type: None
Protection type: None
```


Topology default (ID 0) -> Cost: 50

```
[edit]
lab@R3# run show ospf interface ge-0/0/4.35 extensive
Interface      State   Area      DR ID      BDR ID      Nbrs
ge-0/0/4.35    DR      0.0.0.50   10.50.250.3 10.50.250.5 1
  Type: LAN, Address: 10.50.0.85, Mask: 255.255.255.252, MTU: 1500, Cost: 100
  DR addr: 10.50.0.85, BDR addr: 10.50.0.86, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Stub NSSA
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 100
```

- 7) Make sure that all non backbone OSPF routers wait at least 360 seconds after the OSPF process has started before they forward traffic.

OSPF has an overload concept similar to the IS-IS routing protocol. Overload means that a device is the least preferred transit node and other routers should avoid it. OSPF is actually achieving this by sending the local LSAs with an infinitive metric. Configured with a certain timeout allows OSPF to automatically clear the overload state after the time elapsed.

a. R3

```
[edit]
lab@R3# set protocols ospf overload timeout 360

[edit]
lab@R3# set protocols ospf3 overload timeout 360
```

b. R7

```
[edit]
lab@R7# set protocols ospf overload timeout 360

[edit]
lab@R7# set protocols ospf3 overload timeout 360
```

c. R8

```
[edit]
lab@R8# set protocols ospf3 realm ipv4-unicast overload timeout 360

[edit]
lab@R8# set protocols ospf3 overload timeout 360
```

- 8) Configure R1 to protect the link to R6 using the alternate path through the network. R1 must have that path into the forwarding table ready to be used in case the direct link to R6 fails.

The task requires configuring OSPF to calculate an alternate path through the network. By default, OSPF installs only the best path for a particular route. The task requires you to configure LFA (Loop-Free Alternate Routes) or also known as IP fast-reroute. It is mainly used to provide fast-reroute capability for LDP in MPLS enabled networks.

a. R1

Configuring LFA is relatively simple. You just have to enable link-protection on the interfaces that need to be protected. Similarly to RSVP fast-reroute, LFA can provide link-protection and node-link-protection mechanisms.

```
[edit]
lab@R1# set protocols ospf area 0.0.0.0 interface ge-0/0/4.16 link-protection
```

```
[edit]
lab@R1# set protocols ospf3 area 0.0.0.0 interface ge-0/0/4.16 link-protection
```

Load-balancing policy is required to allow the forwarding engine or the PFE to store more than one next-hop for a given route. The load-balancing policy is applied to the forwarding table.

```
[edit]
lab@R1# set policy-options policy-statement load-balance term 1 then load-balance
per-packet
```

```
[edit]
lab@R1# set policy-options policy-statement load-balance term 1 then accept
```

```
[edit]
lab@R1# set routing-options forwarding-table export load-balance
```

b. Verify the configuration

R1 lists the alternate path through router R4 towards router R6.

```
[edit]
lab@R1# run show ospf backup spf detail 10.50.250.6
```

Topology default results:

Area 0.0.0.0 results:

```
10.50.250.6
Self to Destination Metric: 100
Parent Node: 10.50.250.1
Primary next-hop: ge-0/0/4.16 via 10.50.0.10
Backup next-hop: ge-0/0/4.14 via 10.50.0.6
Backup Neighbor: 10.50.250.6 via: Direct
  Neighbor to Destination Metric: 0, Neighbor to Self Metric: 100
  Self to Neighbor Metric: 100, Backup preference: 0x0
  Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.50.250.4 via: Direct
  Neighbor to Destination Metric: 100, Neighbor to Self Metric: 100
  Self to Neighbor Metric: 100, Backup preference: 0x0
  Track Item: 10.50.250.6
  Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.50.250.2 via: Direct
  Neighbor to Destination Metric: 200, Neighbor to Self Metric: 100
  Self to Neighbor Metric: 100, Backup preference: 0x0
  Track Item: 10.50.250.1
  Track Item: 10.50.250.4
  Track Item: 10.50.250.6
  Not evaluated, Reason: Interface is already covered
```

```
[edit]
lab@R1# run show ospf3 backup spf detail 10.50.250.6
```

Area 0.0.0.0 results:

```
10.50.250.6
Self to Destination Metric: 1
Parent Node: 10.50.250.1
Primary next-hop: ge-0/0/4.16 via fe80::206:a00:100e:604
Backup next-hop: ge-0/0/4.14 via fe80::206:a00:e0e:404
```

```

Backup Neighbor: 10.50.250.6 via: Direct
  Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1, Backup preference: 0x0
  Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.50.250.4 via: Direct
  Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1, Backup preference: 0x0
  Track Item: 10.50.250.6
  Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.50.250.2 via: Direct
  Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1, Backup preference: 0x0
  Track Item: 10.50.250.1
  Track Item: 10.50.250.4
  Track Item: 10.50.250.6
  Not evaluated, Reason: Interface is already covered

```

[edit]

lab@R1# **run show route forwarding-table destination 10.50.250.6**

Routing table: default.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
10.50.250.6/32	user	0		ulst	1048577	4	
			10.50.0.10	ucst	610	4	ge-0/0/4.16
			10.50.0.6	ucst	607	7	ge-0/0/4.14

Routing table: __master.anon__.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	515	1	

Routing table: __juniper_services__.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	537	2	

Routing table: __pfe_private__.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	550	2	

Part 3: Troubleshooting and Configuring BGP

This part is focused on troubleshooting and configuring 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.

Solution - Task 3.1: Internal BGP design and configuration

- 1) Configure the necessary connectivity to the dedicated route-reflector device. It already has partial configuration, moreover you have to ensure that appropriate IGP configuration is in place for providing connectivity to and from the route-reflector.

Start with configuring the connectivity to the route-reflector from routers R4 and R5.

a. R4

```
[edit]
lab@R4# set interfaces ge-0/0/4 unit 49 description "=== connection to route-
reflector ==="

[edit]
lab@R4# set interfaces ge-0/0/4 unit 49 vlan-id 49

[edit]
lab@R4# set interfaces ge-0/0/4 unit 49 family inet address 10.50.0.97/30

[edit]
lab@R4# set protocols ospf area 0.0.0.0 interface ge-0/0/4.49 interface-type p2p

[edit]
lab@R4# commit
commit complete
```

b. R5

```
[edit]
lab@R5# set interfaces ge-0/0/4 unit 59 description "=== connection to route-
reflector ==="

[edit]
lab@R5# set interfaces ge-0/0/4 unit 59 vlan-id 59

[edit]
lab@R5# set interfaces ge-0/0/4 unit 59 family inet address 10.50.0.101/30

[edit]
lab@R5# set protocols ospf area 0.0.0.0 interface ge-0/0/4.59 interface-type p2p

[edit]
lab@R5# commit
commit complete
```

c. Route-reflector

The route-reflector Loopback IP address must be advertised and reachable by all other routers. This is required because iBGP sessions are usually established between the Loopback addresses.

```
[edit]
lab@route-reflector# set protocols ospf area 0 interface lo0.0 passive
```

d. Verify the configuration

Verify that the OSPF adjacencies with the router-reflector are established.

```
[edit]
lab@R4# run show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.50.0.5	ge-0/0/4.14	Full	10.50.250.1	128	39
10.50.0.13	ge-0/0/4.24	Full	10.50.250.2	128	37
10.50.0.34	ge-0/0/4.45	Full	10.50.250.5	128	34
10.50.0.38	ge-0/0/4.46	Full	10.50.250.6	128	39
10.50.0.98	ge-0/0/4.49	Full	10.50.250.10	128	37
10.50.0.81	ae3.0	Full	10.50.250.3	128	37
10.50.0.94	ge-0/0/4.47	Full	10.50.250.7	128	37

The route-reflector has reachability to all Loopback addresses of the other routers in the network.

```
[edit]
lab@route-reflector# run show route 10.50.250.0/28
```

```
inet.0: 540 destinations, 744 routes (540 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

10.50.250.1/32	*[OSPF/10] 00:26:18, metric 201 > to 10.50.0.97 via ge-0/0/1.49 to 10.50.0.101 via ge-0/0/1.59
10.50.250.2/32	*[OSPF/10] 00:26:18, metric 201 > to 10.50.0.97 via ge-0/0/1.49 to 10.50.0.101 via ge-0/0/1.59
10.50.250.3/32	*[OSPF/10] 00:26:02, metric 101 > to 10.50.0.101 via ge-0/0/1.59
10.50.250.4/32	*[OSPF/10] 03:31:54, metric 1 > to 10.50.0.97 via ge-0/0/1.49
10.50.250.5/32	*[OSPF/10] 00:26:22, metric 1 > to 10.50.0.101 via ge-0/0/1.59
10.50.250.6/32	*[OSPF/10] 00:26:18, metric 101 > to 10.50.0.97 via ge-0/0/1.49 to 10.50.0.101 via ge-0/0/1.59
10.50.250.7/32	*[OSPF/10] 00:26:02, metric 201 > to 10.50.0.101 via ge-0/0/1.59
10.50.250.8/32	*[OSPF/150] 00:26:22, metric 100, tag 0 > to 10.50.0.101 via ge-0/0/1.59
10.50.250.10/32	*[Direct/0] 03:32:51 > via lo0.0

- 2) Design an iBGP network with two clusters. Both of the clusters must reside on the route-reflector named respectively – cluster-1 and cluster-2. All Area 0 routers should share one the cluster and the rest of the routers the other. iBGP sessions must tolerate a single link failure. Configure only a single IPv4 iBGP session on each router to the route reflector.

Next, you have to configure the iBGP connectivity. The sessions are built over the Loopback IP addresses to prevent failer due to possible physical links issues.

a. All devices

```
[edit]
lab@R1# set routing-options autonomous-system 5673.873
```

By default, all BGP sessions in JUNOS are external. Hence, configuring an internal BGP session requires manual changes.

```
[edit]
lab@R1# set protocols bgp group internal type internal
```

The Loopback IP address is used for source of the iBGP packets. Using the Loopback address will prevent the sessions to disconnect in the event of an interface failure.

```
[edit]
lab@R1# set protocols bgp group internal local-address 10.50.250.1
```

```
[edit]
lab@R1# set protocols bgp group internal neighbor 10.50.250.10
```

b. Route-reflector

The configuration on the router-reflector router should be divided in two clusters. You have to use the exact names provided in the task description for the configuration of the two clusters. Last, you need to add the correct neighbors to each cluster.

```
[edit]
lab@route-reflector# set routing-options autonomous-system 5673.873
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 type internal
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 local-address 10.50.250.10
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 cluster 1.1.1.1
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 10.50.250.1
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 10.50.250.2
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 10.50.250.4
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 10.50.250.5
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 neighbor 10.50.250.6
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 type internal
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 local-address 10.50.250.10
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 cluster 2.2.2.2
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 10.50.250.3
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 10.50.250.7
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 10.50.250.8
```

c. Verify the configuration

Verifying the BGP summary table reveals that iBGP sessions with all routers in the network are Established.

```
[edit]
lab@route-reflector# run show bgp summary
Groups: 4 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/Accepted/Damped...
10.50.250.1    371786601      10        9        0        0      2:57
0/0/0/0        0/0/0/0
10.50.250.2    371786601      10        8        0        0      2:44
0/0/0/0        0/0/0/0
10.50.250.3    371786601       8        9        0        0      2:57
0/0/0/0        0/0/0/0
10.50.250.4    371786601       7        8        0        0      2:53
0/0/0/0        0/0/0/0
10.50.250.5    371786601       7        8        0        0      2:49
0/0/0/0        0/0/0/0
10.50.250.6    371786601       8        7        0        0      2:45
0/0/0/0        0/0/0/0
10.50.250.7    371786601       3        4        0        0       29
0/0/0/0        0/0/0/0
10.50.250.8    371786601       3        2        0        0       5
0/0/0/0        0/0/0/0
```

Verify one of the routers, R3 in the example below, to confirm that the session is established from the other end too.

```
[edit]
lab@R3# run show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                1          1          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10    371786601      23       20        0        0      8:11
1/1/1/0        0/0/0/0
```

3) All iBGP sessions must use the BFD protocol for faster failure detection.

The task requires configuring the BFD protocol monitoring the BGP sessions. Since iBGP sessions can cross several hops, it is common to use the BFD protocol for faster failure detection. You can choose some reasonable values for liveness detection, if not given by the task.

a. All devices

```
[edit]
lab@R1# set protocols bgp group internal neighbor 10.50.250.10 bfd-liveness-
detection minimum-interval 500
```

```
[edit]
lab@R1# set protocols bgp group internal neighbor 10.50.250.10 bfd-liveness-
detection multiplier 3
```

b. Route-reflector

On the route-reflector you have to add the BFD configuration under both clusters.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 bfd-liveness-detection
minimum-interval 500
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 bfd-liveness-detection
multiplier 3
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 bfd-liveness-detection
minimum-interval 500
```

```
[edit]
lab@route-reflector# set protocols bgp group cluster-2 bfd-liveness-detection
multiplier 3
```

Verifying the BFD sessions table on the route-reflector router reveals that all sessions are UP.

```
[edit]
lab@route-reflector# run show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.50.250.1	Up		1.500	0.500	3
10.50.250.2	Up		1.500	0.500	3
10.50.250.3	Up		1.500	0.500	3
10.50.250.4	Up		1.500	0.500	3
10.50.250.5	Up		1.500	0.500	3
10.50.250.6	Up		1.500	0.500	3
10.50.250.7	Up		1.500	0.500	3
10.50.250.8	Up		1.500	0.500	3

8 sessions, 8 clients

Cumulative transmit rate 16.0 pps, cumulative receive rate 16.0 pps

- 4) Advertise a summary IPv4 route 10.50.250.0/28 representing the loopback addresses of all routers to iBGP. You have to make sure to do that on R4 and R5 only. Ensure that all routers have the summary present in the routing table learnt from BGP.

Next, you have to advertise the summary route representing the Loopback IPv4 addresses of all routers. The task restricts you to do that only on routers R4 and R5.

a. R4

First, you have to configure an aggregate route as requested by the task. Then you have to use that route in routing policy exported to the route-reflector router, where it will be reflected to the other routers.

```
[edit]
lab@R4# set routing-options aggregate route 10.50.250.0/28
```

```
[edit]
lab@R4# set policy-options policy-statement aggregate.v4 term 1 from route-filter
10.50.250.0/28 exact
```

```
[edit]
lab@R4# set policy-options policy-statement aggregate.v4 term 1 then accept
```

```
[edit]
lab@R4# set protocols bgp group internal export aggregate.v4
```


b. R5

The same configuration is applied to router R5 as well.

```
[edit]
lab@R5# set routing-options aggregate route 10.50.250.0/28

[edit]
lab@R5# set policy-options policy-statement aggregate.v4 term 1 from route-filter
10.50.250.0/28 exact

[edit]
lab@R5# set policy-options policy-statement aggregate.v4 term 1 then accept

[edit]
lab@R5# set protocols bgp group internal export aggregate.v4
```

c. Route-reflector

You may wonder why you have to amend the next-hop of the summary route when the route-reflector send it to routers R3 and R7. The reason is the traffic separation of IPv4 and IPv6 traffic in the IGP section. Router R4 is configured to advertise only a default route to Area 50. In case R5 fails, R3 and R7 would only have the default route from R4, which will cause their iBGP sessions to fail. That is because on R3 and R7 the next-hop of the route to the route reflector Loopback IP address will resolve to the summary route receive by the route-reflector iBGP sessions. To prevent that to happen you can change the next-hop of the summary route to a route resolved locally in the Area 50. This way R3 and R7 will be able to resolve the next-hop by a more specific route.

```
[edit]
lab@route-reflector# set policy-options policy-statement amend.aggregate term 1
from route-filter 10.50.250.0/28 exact

[edit]
lab@route-reflector# set policy-options policy-statement amend.aggregate term 1
then next-hop 10.50.0.93

[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 10.50.250.3
export amend.aggregate

[edit]
lab@route-reflector# set protocols bgp group cluster-2 neighbor 10.50.250.7
export amend.aggregate
```

d. Verify the configuration

The route-reflector receives the route from both R4 and R5.

```
[edit]
lab@route-reflector# run show route 10.50.250.0/28

inet.0: 26 destinations, 27 routes (26 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.0/28      *[BGP/170] 00:00:35, localpref 100, from 10.50.250.4
                   AS path: I, validation-state: unverified
                   > to 10.50.0.97 via ge-0/0/1.49
                   [BGP/170] 00:00:14, localpref 100, from 10.50.250.5
                   AS path: I, validation-state: unverified
                   > to 10.50.0.101 via ge-0/0/1.59
```

Below output shows what will happen if the next-hop of the summary route send to R3 and R7 is not change and the failure in R5 routing happens.

```
[edit]
lab@R3# run show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State    Pending
inet.0
                1          1          0          0          0          0
Peer           AS         InPkt    OutPkt    OutQ     Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10    371786601          23          20          0          0          8:11
1/1/1/0        0/0/0/0
```

```
[edit]
lab@R5# deactivate protocols ospf
```

```
[edit]
lab@R5# commit
commit complete
```

The session to the route-reflector goes to Active state and cannot be established again.

```
[edit]
lab@R3# 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...
10.50.250.10    371786601          24          21          0          1          15
Active
```

Solution - Task 3.2: External BGP peers and configuration

- 5) Configure eBGP sessions with the Upstream (U1, U2 and U3) peers using the diagram IP addresses. U3 requires a separate eBGP session for advertising IPv6 prefixes. Make sure you receive IPv4 and IPv6 routes from upstream U3.

Next, you have to successfully configure eBGP sessions to the external peers connected to the network.

a. R3

Starting with the configuration of the interfaces to U1 and U2. Again, do not forget to configure the interfaces to accept tagged traffic where needed.

```
[edit]
lab@R3# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R3# set interfaces ge-0/0/5 unit 2001 description "=== connection to U1 ==="

[edit]
lab@R3# set interfaces ge-0/0/5 unit 2001 vlan-id 2001

[edit]
lab@R3# set interfaces ge-0/0/5 unit 2001 family inet address 15.16.17.130/30

[edit]
lab@R3# set interfaces ge-0/0/5 unit 2002 description "=== connection to U2 ==="

[edit]
lab@R3# set interfaces ge-0/0/5 unit 2002 vlan-id 2002

[edit]
lab@R3# set interfaces ge-0/0/5 unit 2002 family inet address 15.16.18.82/30

[edit]
lab@R3# set protocols bgp group U1 peer-as 87.109

[edit]
lab@R3# set protocols bgp group U1 neighbor 15.16.17.129

[edit]
lab@R3# set protocols bgp group U2 peer-as 87.109

[edit]
lab@R3# set protocols bgp group U2 neighbor 15.16.18.81
```

b. R5

Configuring the interface on router R5 to peer U2.

```
[edit]
lab@R5# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R5# set interfaces ge-0/0/5 unit 2003 description "=== connection to U2 ==="

[edit]
lab@R5# set interfaces ge-0/0/5 unit 2003 vlan-id 2003

[edit]
lab@R5# set interfaces ge-0/0/5 unit 2003 family inet address 15.16.18.86/30

[edit]
```

```
lab@R5# set protocols bgp group U2 peer-as 87.109

[edit]
lab@R5# set protocols bgp group U2 neighbor 15.16.18.85
```

c. R6

The task requires that R6 should have two separate sessions to U3, one for IPv4 and other for IPv6. Below is shown how you can accomplish this on router R6.

```
[edit]
lab@R6# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R6# set interfaces ge-0/0/5 unit 2004 description "=== connection to U3 ==="

[edit]
lab@R6# set interfaces ge-0/0/5 unit 2004 vlan-id 2004

[edit]
lab@R6# set interfaces ge-0/0/5 unit 2004 family inet address 195.34.23.18/30

[edit]
lab@R6# set interfaces ge-0/0/5 unit 2004 family inet6 address
2001:ba:34:23::2/64

[edit]
lab@R6# set protocols bgp group U3-v4 peer-as 9687

[edit]
lab@R6# set protocols bgp group U3-v4 neighbor 195.34.23.17

[edit]
lab@R6# set protocols bgp group U3-v6 peer-as 9687

[edit]
lab@R6# set protocols bgp group U3-v6 neighbor 2001:ba:34:23::1
```

d. Verify the configuration

Verifying the BGP table on R3, R5 and R6 routers shows that eBGP sessions are successfully established and prefixes are received.

```
[edit]
lab@R3# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10 371786601 12 510 0 2 3:28
1/1/1/0 0/0/0/0
15.16.17.129 5701741 396 4 0 0 10
337/460/460/0 0/0/0/0
15.16.18.81 5701741 394 4 0 0 6
123/460/460/0 0/0/0/0

[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
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
```

```

10.50.250.10      371786601      10      403      0      2      2:59
0/1/1/0          0/0/0/0
15.16.18.85      5701741      395      3      0      0      2
460/460/460/0    0/0/0/0

[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
                461        461          0          0          0          0
inet6.0
                88         88          0          0          0          0
Peer           AS         InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10    371786601      13      404        0          1      5:06
1/1/1/0         0/0/0/0
195.34.23.17    9687          391        4          0          0      22
460/460/460/0   0/0/0/0
2001:ba:34:23::1 9687          81         3          0          0      7
Establ
inet6.0: 88/88/88/0

```

- 6) Configure eBGP session with the Partner (P1) peer using the diagram IP addresses. Make sure that you receive IPv4 and IPv6 prefixes from P1 over a single IPv4 eBGP session.

Unlike the two separate eBGP sessions with U3, with P1 you have to configure a single eBGP session carrying both IPv4 and IPv6 traffic.

a. R7

```

[edit]
lab@R7# set interfaces ge-0/0/5 unit 2012 description "=== connection to P1 ==="

[edit]
lab@R7# set interfaces ge-0/0/5 unit 2012 vlan-id 2012

[edit]
lab@R7# set interfaces ge-0/0/5 unit 2012 family inet address 136.41.0.2/30

```

According to the diagram the IPv6 addresses used for communication are IPv4-compatible IPv6 addresses.

```

[edit]
lab@R7# set interfaces ge-0/0/5 unit 2012 family inet6 address ::136.41.0.2/96

```

The command below allows the BGP process to accept prefixes with a next-hop different then the address of the BGP session.

```

[edit]
lab@R7# set protocols bgp group P1 accept-remote-nexthop

```

Although the command above is enough to allow router R7 to accept the route as active in the routing table, JUNOS by default assigns an IPv4-mapped address. You have to change that to the format used on the diagram to successfully transmit traffic to and from P1.

```

[edit]

```

```

lab@R7# set policy-options policy-statement P1.in term accept.ipv6 from family
inet6

[edit]
lab@R7# set policy-options policy-statement P1.in term accept.ipv6 from protocol
bgp

[edit]
lab@R7# set policy-options policy-statement P1.in term accept.ipv6 then next-hop
::136.41.0.1

[edit]
lab@R7# set policy-options policy-statement P1.in term accept.ipv6 then accept

[edit]
lab@R7# set protocols bgp group P1 import P1.in

```

b. Verify the configuration

Router R7 successfully establishes eBGP session to P1 and receives IPv4 and IPv6 routes.

```

[edit]
lab@R7# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                18         17         0         0         0         0
inet6.0
                16         0         0         0         0         0
Peer           AS          InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10   371786601    252      275      0        1    1:56:44
0/1/1/0       0/0/0/0
136.41.0.1     51707916    12        6        0        2    1:27
Establ
  inet.0: 17/17/17/0
  inet6.0: 16/16/16/0

```

Below is shown how you can troubleshoot BGP prefix issues.

```

[edit]
lab@R7# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                18         17         0         0         0         0
inet6.0
                16         0         0         0         0         0
Peer           AS          InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10   371786601    252      275      0        1    1:56:44
0/1/1/0       0/0/0/0
136.41.0.1     51707916    12        6        0        2    1:27
Establ
  inet.0: 17/17/17/0
  inet6.0: 0/16/0/0

[edit]
lab@R7# run show route table inet6.0 protocol bgp hidden extensive

inet6.0: 52 destinations, 52 routes (36 active, 0 holddown, 16 hidden)
fd17:f0f4:f691:d::/80 (1 entry, 0 announced)
  BGP
    Next hop type: Router, Next hop index: 0
    Address: 0xa6df508
    Next-hop reference count: 16
    Source: 136.41.0.1

```

```

Next hop: ::ffff:136.41.0.1 via ge-0/0/5.2012, selected
Session Id: 0x0
State: <Hidden Ext>
Local AS: 371786601 Peer AS: 51707916
Age: 34
Validation State: unverified
Task: BGP_51707916.136.41.0.1
AS path: 51707916 I
Router ID: 136.41.0.1
Hidden reason: protocol nexthop is not on the interface

```

```

[edit]
lab@R7# set protocols bgp traceoptions flag update

```

```

[edit]
lab@R7# set protocols bgp traceoptions file BGP_log

```

```

[edit]
lab@R7# commit
commit complete

```

```

Dec 2 15:01:09.492786 BGP RECV message type 2 (Update) length 237
Dec 2 15:01:09.492999 BGP_UNUSABLE_NEXTHOP: bgp_nexthop_sanity: peer 136.41.0.1
(External AS 51707916) next hop ::ffff:136.41.0.1 unexpectedly remote, ignoring
routes in this update
Dec 2 15:01:09.493020 bgp_rcv_nlri: fd17:f0f4:f691:d::/80
Dec 2 15:01:09.493046 bgp_rcv_nlri: fd17:f0f4:f691:d:1::/80
Dec 2 15:01:09.493058 bgp_rcv_nlri: fd17:f0f4:f691:d:2::/80
Dec 2 15:01:09.493073 bgp_rcv_nlri: fd17:f0f4:f691:d:3::/80
Dec 2 15:01:09.493085 bgp_rcv_nlri: fd17:f0f4:f691:d:4::/80
Dec 2 15:01:09.493095 bgp_rcv_nlri: fd17:f0f4:f691:d:5::/80
Dec 2 15:01:09.493106 bgp_rcv_nlri: fd17:f0f4:f691:d:6::/80
Dec 2 15:01:09.493115 bgp_rcv_nlri: fd17:f0f4:f691:d:7::/80
Dec 2 15:01:09.493125 bgp_rcv_nlri: fd17:f0f4:f691:d:8::/80
Dec 2 15:01:09.493139 bgp_rcv_nlri: fd17:f0f4:f691:d:9::/80
Dec 2 15:01:09.493151 bgp_rcv_nlri: fd17:f0f4:f691:d:a::/80
Dec 2 15:01:09.493161 bgp_rcv_nlri: fd17:f0f4:f691:d:b::/80
Dec 2 15:01:09.493172 bgp_rcv_nlri: fd17:f0f4:f691:d:c::/80
Dec 2 15:01:09.493181 bgp_rcv_nlri: fd17:f0f4:f691:d:d::/80
Dec 2 15:01:09.493191 bgp_rcv_nlri: fd17:f0f4:f691:d:e::/80
Dec 2 15:01:09.493216 bgp_rcv_nlri: fd17:f0f4:f691:d:f::/80
Dec 2 15:01:09.493233
Dec 2 15:01:09.493233 BGP RECV 136.41.0.1+58429 -> 136.41.0.2+179
Dec 2 15:01:09.493240 BGP RECV message type 2 (Update) length 30

```

After the `accept-remote-nexthop` command is applied and the routing policy changing the next-hop, prefixes are not hidden anymore.

```

[edit]
lab@R7# run show route table inet6.0 protocol bgp extensive

inet6.0: 52 destinations, 52 routes (52 active, 0 holddown, 0 hidden)
fd17:f0f4:f691:d::/80 (1 entry, 1 announced)
TSI:
KRT in-kernel fd17:f0f4:f691:d::/80 -> {::136.41.0.1}
    *BGP      Preference: 170/-101
              Next hop type: Router, Next hop index: 627
              Address: 0xa6db758
              Next-hop reference count: 32
              Source: 136.41.0.1
              Next hop: ::136.41.0.1 via ge-0/0/5.2012, selected
              Session Id: 0x148
              State: <Active Ext>
              Local AS: 371786601 Peer AS: 51707916
              Age: 19
              Validation State: unverified
              Task: BGP_51707916.136.41.0.1

```

```

Announcement bits (2): 0-KRT 5-Resolve tree 3
AS path: 51707916 I
Accepted
Localpref: 100
Router ID: 136.41.0.1

```

- 7) Configure an eBGP session with the Customer (C1 and C2) peers using the diagram IP addresses. C1 supports two separate eBGP sessions, one for IPv4 and another for IPv6, however the IPv4 and IPv6 addresses used for connectivity are not known. You have to find a way to discover them.

Again, you have to configure two separate sessions between router R4 to peer C1. As stated in the task description and shown on diagram, the peer IP addresses of C1 are not known. You have to find a way to discover them.

Configuring random addresses for each of the sessions and enabling traceoptions reveals the actual addresses used by peer C1.

a. R4

```

[edit]
lab@R4# set protocols bgp group C1-v4 peer-as 65457

[edit]
lab@R4# set protocols bgp group C1-v4 neighbor 170.0.15.1

[edit]
lab@R4# set protocols bgp group C1-v6 peer-as 65457

[edit]
lab@R4# set protocols bgp group C1-v6 neighbor 2001:b170:0:15::1

[edit]
lab@R4# commit
commit complete

[edit]
lab@R4# set protocols bgp traceoptions file BGP_log

[edit]
lab@R4# set protocols bgp traceoptions flag open

[edit]
lab@R4# commit
commit complete

[edit]
lab@R4# run show log BGP_log
Dec  2 14:53:07 trace_on: Tracing to "/var/log/BGP_log" started
Dec  2 15:00:14.055849 bgp_listen_accept: Connection attempt from unconfigured
neighbor: 170.0.15.10+62043
Dec  2 15:00:14.056205
Dec  2 15:00:14.056205 BGP SEND 170.0.15.4+179 -> 170.0.15.10+62043
Dec  2 15:00:14.056219 BGP SEND message type 1 (Open) length 37
Dec  2 15:00:14.056226 BGP SEND version 4 as 23456 holdtime 90 id 10.50.250.4
parmlen 8
Dec  2 15:00:14.056230 BGP SEND Refresh capability, code=128
Dec  2 15:00:14.056232 BGP SEND Refresh capability, code=2
Dec  2 15:00:14.056238
Dec  2 15:00:14.056238 BGP SEND 170.0.15.4+179 -> 170.0.15.10+62043
Dec  2 15:00:14.056245 BGP SEND message type 3 (Notification) length 21
Dec  2 15:00:14.056249 BGP SEND Notification code 6 (Cease) subcode 5 (Connection
Rejected)

```



```
Dec 2 15:00:14.056312 bgp_listen_accept:4984: NOTIFICATION sent to
170.0.15.10+62043 (proto): code 6 (Cease) subcode 5 (Connection Rejected),
Reason: Connection attempt from unconfigured neighbor: 170.0.15.10+62043
Dec 2 15:00:35.329175 bgp_listen_accept: Connection attempt from unconfigured
neighbor: 2001:b170:0:15::be+60905
Dec 2 15:00:35.329448
Dec 2 15:00:35.329448 BGP SEND 2001:b170:0:15::4+179 -> 2001:b170:0:15::be+60905
Dec 2 15:00:35.329462 BGP SEND message type 1 (Open) length 37
Dec 2 15:00:35.329477 BGP SEND version 4 as 23456 holdtime 90 id 10.50.250.4
parmlen 8
Dec 2 15:00:35.329480 BGP SEND Refresh capability, code=128
Dec 2 15:00:35.329483 BGP SEND Refresh capability, code=2
Dec 2 15:00:35.329490
Dec 2 15:00:35.329490 BGP SEND 2001:b170:0:15::4+179 -> 2001:b170:0:15::be+60905
Dec 2 15:00:35.329496 BGP SEND message type 3 (Notification) length 21
Dec 2 15:00:35.329503 BGP SEND Notification code 6 (Cease) subcode 5 (Connection
Rejected)
Dec 2 15:00:35.329520 bgp_listen_accept:4984: NOTIFICATION sent to
2001:b170:0:15::be+60905 (proto): code 6 (Cease) subcode 5 (Connection Rejected),
Reason: Connection attempt from unconfigured neighbor: 2001:b170:0:15::be+60905
```

Configuring the correct addresses successfully establish two eBGP sessions from router R4 to peer C1.

```
[edit]
lab@R4# set protocols bgp group C1-v4 peer-as 65457

[edit]
lab@R4# set protocols bgp group C1-v4 neighbor 170.0.15.10

[edit]
lab@R4# set protocols bgp group C1-v6 peer-as 65457

[edit]
lab@R4# set protocols bgp group C1-v6 neighbor 2001:b170:0:15::be
```

b. Verify the configuration

```
[edit]
lab@R4# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 19 19 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...
10.50.250.10 371786601 367 362 0 0 2:39:26
0/0/0/0
170.0.15.10 65457 8 7 0 0 1:50
19/19/19/0 0/0/0/0
2001:b170:0:15::be 65457 6 6 0 0 1:46
Establ
inet6.0: 16/16/16/0
```

- 8) The C2 customer is recently migrated from other provider and is currently configured to peer with AS7862. Until C2's configuration is corrected, make sure it successfully establishes an eBGP session with R2. To prevent C2 prefixes being discarded due to loop protection, make sure AS7862 is not added to the AS path.

Router R2 should establish a single IPv4 eBGP session with peer C2. The task mentions that the peer is configured by mistake to peer with different AS number from the local AS.

To configure router R2 to pretend it is part of AS7862 you have to use the `local-as` when configuring the eBGP neighbor.

a. R2

```
[edit]
lab@R2# set interfaces ge-0/0/5 vlan-tagging

[edit]
lab@R2# set interfaces ge-0/0/5 unit 2022 description "=== connection to C2 ==="

[edit]
lab@R2# set interfaces ge-0/0/5 unit 2022 vlan-id 2022

[edit]
lab@R2# set interfaces ge-0/0/5 unit 2022 family inet address 88.136.0.33/30

[edit]
lab@R2# set protocols bgp group C2 peer-as 65456
```

The `local-as` statement is used to make router R2 pretend it is part of AS7862. The `private` option is added here to hide that additional AS number from the `as-path` attribute when the C2 prefixes are advertised further out to the other eBGP peers.

```
[edit]
lab@R2# set protocols bgp group C2 local-as 7862

[edit]
lab@R2# set protocols bgp group C2 local-as private

[edit]
lab@R2# set protocols bgp group C2 neighbor 88.136.0.34
```

b. Verify the configuration

After committing the configuration, the session with neighbor C2 is established successfully and prefixes are received.

```
[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
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10 371786601 373 476 0 1 3:32:36
0/0/0/0 0/0/0/0
88.136.0.34 65456 4 3 0 0 4
18/18/18/0 0/0/0/0
```

Verify if the additional AS number is removed from the prefixes advertised out to external peers. The output below shows one of the C2 prefixes advertised to peer U1 from router R3. AS7862 is missing from the `as-path`.

```
[edit]
lab@R3# run show route advertising-protocol bgp 15.16.18.81 193.46.132.0/24

inet.0: 558 destinations, 1221 routes (554 active, 0 holddown, 8 hidden)
  Prefix Nexthop MED Lclpref AS path
* 193.46.132.0/24 Self 65456 I
```

To verify that the `private` options is also needed to meet the task requirements, see below the same prefix as shown in the example above but without the option used.

```
[edit]
lab@R3# run show route advertising-protocol bgp 15.16.18.81 193.46.132.0/24

inet.0: 558 destinations, 1221 routes (554 active, 0 holddown, 8 hidden)
  Prefix            Nexthop          MED      Lclpref   AS path
* 193.46.132.0/24    Self              7862     65456 I
```

9) Ensure all eBGP neighbors state changes are logged to syslog.

The task requires you to configure all routers that have eBGP sessions to generate a syslog message whenever there is a change in the state of the BGP peer.

a. All devices with external BGP peers

```
[edit]
lab@R2# set protocols bgp group C2 log-updown

[edit]
lab@R3# set protocols bgp group U1 log-updown

[edit]
lab@R3# set protocols bgp group U2 log-updown
```

Solution - Task 3.3: BGP policy configuration

10) All IPv4 prefixes received from Upstream, Partner and Customer peers must be active and reachable on all routers in your AS unless they are filtered by policy in some of the subsequent tasks.

a. R3

```
[edit]
lab@R3# set policy-options policy-statement nhs then next-hop self
```

```
[edit]
lab@R3# set protocols bgp group internal export nhs
```

b. R5

```
[edit]
lab@R5# set policy-options policy-statement nhs then next-hop self
```

```
[edit]
lab@R5# set protocols bgp group internal export nhs
```

c. R6

```
[edit]
lab@R6# set policy-options policy-statement nhs then next-hop self
```

```
[edit]
lab@R6# set protocols bgp group internal export nhs
```

d. R4

```
[edit]
lab@R4# set policy-options policy-statement nhs then next-hop self
```

```
[edit]
lab@R4# set protocols bgp group internal export nhs
```

e. R2

```
[edit]
lab@R2# set policy-options policy-statement nhs then next-hop self
```

```
[edit]
lab@R2# set protocols bgp group internal export nhs
```

f. Verify the configuration

Verifying on the route-reflector that routes are received and active from the border routers in the network.

```
[edit]
lab@route-reflector# run show bgp summary
Groups: 5 Peers: 8 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 1179 514 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.1 371786601 823 1686 0 0 4:56:42
0/0/0/0 0/0/0/0
```

10.50.250.2	371786601	824	1486	0	0	4:56:29
18/18/18/0	0/0/0/0					
10.50.250.3	371786601	1503	837	0	0	4:45:18
256/460/460/0	0/0/0/0					
10.50.250.4	371786601	633	1680	0	0	4:56:38
20/20/20/0	0/0/0/0					
10.50.250.5	371786601	1138	1469	0	2	4:44:22
204/205/205/0	0/0/0/0					
10.50.250.6	371786601	1426	1660	0	1	4:46:49
0/460/460/0	0/0/0/0					
10.50.250.7	371786601	671	843	0	1	4:46:50
16/16/16/0	0/0/0/0					
10.50.250.8	371786601	616	841	0	1	4:46:49
0/0/0/0	0/0/0/0					

11) Advertise a summary route representing all your internal loopback IPv4 address to all external eBGP peers. Make sure that you avoid amending policies to accomplish this where possible.

The next task is asking to advertise a local summary route to the eBGP peers. However, you are not allowed to amend the routing-option stanza, in other words you should use the route already present in the routing table. BGP is advertising by default any BGP learned route further to its peers. The only reason this might not happen is if the route is inactive. This is the case on some of the routers because they have the same route learned by OSPF protocol, which has a lower route preference compared to BGP. This requires the use of the advertise-inactive command used below to force BGP to advertise routes despite its inactive.

a. R3

```
[edit]
lab@R3# set protocols bgp group U1 advertise-inactive
```

```
[edit]
lab@R3# set protocols bgp group U2 advertise-inactive
```

b. R5

```
[edit]
lab@R5# set protocols bgp advertise-inactive
```

c. R4

```
[edit]
lab@R4# set policy-options policy-statement C1.out term local.aggregate from
route-filter 10.50.250.0/28 exact
```

```
[edit]
lab@R4# set policy-options policy-statement C1.out term local.aggregate then
accept
```

d. R2

```
[edit]
lab@R2# set policy-options policy-statement C2.out term local.aggregate from
route-filter 10.50.250.0/28 exact
```

```
[edit]
lab@R2# set policy-options policy-statement C2.out term local.aggregate then
accept
```

e. Verify the configuration

The outputs below show the local summary route advertised to all eBGP peers. The “*” in front of the line shows that the particular route is active in the local routing table.

```
[edit]
lab@R3# run show route advertising-protocol bgp 15.16.17.129 192.168.68.0/24

inet.0: 514 destinations, 976 routes (514 active, 0 holddown, 0 hidden)
  Prefix            Nexthop          MED      Lclpref    AS path
* 192.168.68.0/24      Self                      65456 I

[edit]
lab@R3# run show route advertising-protocol bgp 15.16.18.81 192.168.68.0/24

inet.0: 514 destinations, 976 routes (514 active, 0 holddown, 0 hidden)
  Prefix            Nexthop          MED      Lclpref    AS path
* 192.168.68.0/24      Self                      65456 I
```

On router R5 the “*” sign is missing in front of the local summary route, this means the particular route is not active. But despite this, it is successfully advertised to U2 due to the use of advertise-inactive.

```
[edit]
lab@R5# run show route advertising-protocol bgp 15.16.18.85 10.50.250.0/28

inet.0: 551 destinations, 817 routes (551 active, 0 holddown, 0 hidden)
  Prefix            Nexthop          MED      Lclpref    AS path
  10.50.250.0/28      Self                      I

[edit]
lab@R5# run show route 10.50.250.0/28

inet.0: 551 destinations, 817 routes (551 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.0/28      *[Aggregate/130] 04:06:54
                    Reject
                    [BGP/170] 04:02:00, localpref 100, from 10.50.250.10
                    AS path: I, validation-state: unverified
                    > to 10.50.0.33 via ge-0/0/4.45
```

```
[edit]
lab@R2# run show route advertising-protocol bgp 88.136.0.34 10.50.250.0/28

inet.0: 549 destinations, 550 routes (547 active, 0 holddown, 2 hidden)
  Prefix            Nexthop          MED      Lclpref    AS path
* 10.50.250.0/28      Self                      5673.873 I

[edit]
lab@R4# run show route advertising-protocol bgp 170.0.15.10 10.50.250.0/28

inet.0: 579 destinations, 580 routes (576 active, 0 holddown, 3 hidden)
  Prefix            Nexthop          MED      Lclpref    AS path
* 10.50.250.0/28      Self                      I
```

- 12) All three Upstream peers are advertising the same prefixes. Ensure that all internal routers are selecting the best path to the Upstream peer based on the MED attribute. Apart of R3 which has to prefer the link to U1 when sending traffic to AS 87.109.

The MED or metric BGP attribute by default is only compared between routes that are received by the same direct AS. To force the route-reflector and the border routers to use

it for all routes you have to amend the BGP path algorithm using the `always-compare-med` knob.

a. Route-reflector

```
[edit]
lab@route-reflector# set protocols bgp path-selection always-compare-med

[edit]
lab@route-reflector# commit
commit complete
```

b. R3

R3 should prefer routes from peer U1 over routes received by U2. You should be careful with the approach here. If you decide to use higher local-preference that will make the routes preferred, but this will also influence the decision on other routers, the route-reflector for instance, which would compromise the first requirement. The correct way for solving this task is to lower the default routing preference only for the routes received from peer U1. The change will only have local significance, hence achieving what is asked in the task.

```
[edit]
lab@R3# set protocols bgp path-selection always-compare-med

[edit]
lab@R3# set policy-options policy-statement U1.in term lower.preference then
preference 160
```

c. R5

```
[edit]
lab@R5# set protocols bgp path-selection always-compare-med
```

d. R6

```
[edit]
lab@R6# set protocols bgp path-selection always-compare-med
```

e. Verify the configuration

```
[edit]
lab@route-reflector# run show bgp summary
Groups: 5 Peers: 8 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0
1179 514 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.1 371786601 884 1810 0 0 5:18:42
0/0/0/0 0/0/0/0
10.50.250.2 371786601 885 1534 0 0 5:18:29
18/18/18/0 0/0/0/0
10.50.250.3 371786601 1547 962 0 0 5:07:18
256/460/460/0 0/0/0/0
10.50.250.4 371786601 679 1805 0 0 5:18:38
20/20/20/0 0/0/0/0
10.50.250.5 371786601 1179 1594 0 2 5:06:22
123/205/205/0 0/0/0/0
10.50.250.6 371786601 1474 1712 0 1 5:08:49
81/460/460/0 0/0/0/0
10.50.250.7 371786601 721 967 0 1 5:08:50
16/16/16/0 0/0/0/0
```

```
10.50.250.8      371786601      663      965      0      1      5:08:49
0/0/0/0          0/0/0/0
```

The change forces the route-reflector to prefer the routes with the lowest MED among the same routes received from peers U1, U2 and U3.

```
[edit]
lab@route-reflector# run show route

inet.0: 539 destinations, 1204 routes (539 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/4      *[BGP/170] 00:27:19, MED 50, localpref 100, from 10.50.250.3
                AS path: 5701741 1620 61671 I, validation-state: unverified
                > to 10.50.0.101 via ge-0/0/1.59
                [BGP/170] 00:38:52, MED 150, localpref 100, from 10.50.250.6
                AS path: 9687 1620 61671 I, validation-state: unverified
                > to 10.50.0.97 via ge-0/0/1.49
                > to 10.50.0.101 via ge-0/0/1.59
0.0.0.0/5      *[BGP/170] 00:27:19, MED 50, localpref 100, from 10.50.250.3
                AS path: 5701741 1620 61671 27075 I, validation-state:
unverified
                > to 10.50.0.101 via ge-0/0/1.59
                [BGP/170] 00:38:52, MED 150, localpref 100, from 10.50.250.6
                AS path: 9687 1620 61671 27075 I, validation-state:
unverified
                > to 10.50.0.97 via ge-0/0/1.49
                > to 10.50.0.101 via ge-0/0/1.59
1.64.0.0/10    *[BGP/170] 00:27:19, MED 50, localpref 100, from 10.50.250.3
                AS path: 5701741 1620 61671 I, validation-state: unverified
                > to 10.50.0.101 via ge-0/0/1.59
                [BGP/170] 00:38:52, MED 150, localpref 100, from 10.50.250.6
                AS path: 9687 1620 61671 I, validation-state: unverified
                > to 10.50.0.97 via ge-0/0/1.49
                > to 10.50.0.101 via ge-0/0/1.59
```

The change of the routing preference on router R3 makes the routes received by U1 to be the most preferred although the same always-compared-med options was used there too.

```
[edit]
lab@R3# run show route

inet.0: 528 destinations, 1195 routes (528 active, 0 holddown, 8 hidden)
+ = Active Route, - = Last Active, * = Both

1.64.0.0/10    *[BGP/160] 00:01:21, MED 50, localpref 100
                AS path: 5701741 1620 61671 I, validation-state: unverified
                > to 15.16.17.129 via ge-0/0/5.2001
                [BGP/170] 00:42:47, MED 150, localpref 100
                AS path: 5701741 1620 61671 I, validation-state: unverified
                > to 15.16.18.81 via ge-0/0/5.2002
1.84.160.0/20  *[BGP/160] 00:01:21, MED 150, localpref 100
                AS path: 5701741 1620 33112 I, validation-state: unverified
                > to 15.16.17.129 via ge-0/0/5.2001
                [BGP/170] 00:42:47, MED 60, localpref 100
                AS path: 5701741 1620 33112 I, validation-state: unverified
                > to 15.16.18.81 via ge-0/0/5.2002
                [BGP/170] 00:42:58, MED 60, localpref 100, from 10.50.250.10
                AS path: 5701741 1620 33112 I, validation-state: unverified
                > to 10.50.0.86 via ge-0/0/4.35
1.96.0.0/11    *[BGP/160] 00:01:21, MED 50, localpref 100
                AS path: 5701741 1620 33112 63164 40776 51777 I,
validation-state: unverified
                > to 15.16.17.129 via ge-0/0/5.2001
                [BGP/170] 00:42:47, MED 150, localpref 100
```



```

AS path: 5701741 1620 33112 63164 40776 51777 I,
validation-state: unverified
> to 15.16.18.81 via ge-0/0/5.2002

```

- 13) Accept only prefixes originated from P1 autonomous system on R7. Traffic from P1 must not use your AS for transit to Upstreams or Customers networks. You are allowed to do changes only to P1 policies to accomplish this.

a. R7

To ensure that only routes advertised by peer P1 are accepted by R7 you have to use as-path regular expression to match on those routes. Apart from that you also have to prevent those routes not to be advertised to other external peers. Assigning the well-known community no-export, will achieve that. Routes assigned with this community are not advertised further to eBGP peers. In the opposite direction you have to restrict the advertisement of the Upstream and Customer routes to peer P1.

```

[edit]
lab@R7# set policy-options policy-statement P1.in term P1.origin from protocol
bgp

[edit]
lab@R7# set policy-options policy-statement P1.in term P1.origin from as-path
P1.origin

[edit]
lab@R7# set policy-options policy-statement P1.in term P1.origin then community
add no.export

[edit]
lab@R7# set policy-options policy-statement P1.in term P1.origin then accept

[edit]
lab@R7# set policy-options policy-statement P1.in term reject.rest from protocol
bgp

[edit]
lab@R7# set policy-options policy-statement P1.in term reject.rest then reject

[edit]
lab@R7# set policy-options policy-statement P1.out term term reject.cutomers from
protocol bgp

[edit]
lab@R7# set policy-options policy-statement P1.out term term reject.cutomers from
community [ C1.routes C2.routes ]

[edit]
lab@R7# set policy-options policy-statement P1.out term term reject.cutomers then
reject

[edit]
lab@R7# set policy-options policy-statement P1.out term term reject.upstream from
protocol bgp

[edit]
lab@R7# set policy-options policy-statement P1.out term term reject.upstream from
family inet

[edit]
lab@R7# set policy-options policy-statement P1.out term term reject.cutomers from
as-path [ UlandU2 U3 ]

[edit]

```

```
lab@R7# set policy-options policy-statement P1.out term term reject.cutomers then reject
```

```
[edit]
```

```
lab@R7# set policy-options community no.export members no-export
```

```
[edit]
```

```
lab@R7# set policy-options community C1.routes members 65457:2021
```

```
[edit]
```

```
lab@R7# set policy-options community C2.routes members 65456:2022
```

To match 4byte AS numbers in regular expressions you need to escape the dot as shown below.

```
[edit]
```

```
lab@R7# set policy-options as-path P1.origin "789\.12"
```

```
[edit]
```

```
lab@R7# set policy-options as-path U1andU2 "87\.109"
```

```
[edit]
```

```
lab@R7# set policy-options as-path U3 "9687"
```

```
[edit]
```

```
lab@R7# set protocols bgp group P1 import P1.in
```

```
[edit]
```

```
lab@R7# set protocols bgp group P1 export P1.out
```

b. Verify the configuration

The output below shows that only routes advertised by peer P1 are received. However, the rest of the routes are filtered and hidden.

```
[edit]
```

```
lab@R7# run show route receive-protocol bgp 136.41.0.1
```

```
inet.0: 310 destinations, 311 routes (309 active, 0 holddown, 1 hidden)
  Prefix                Nexthop              MED      Lclpref    AS path
* 19.0.32.0/24          136.41.0.1              51707916 I
* 19.0.33.0/24          136.41.0.1              51707916 I
* 19.0.34.0/24          136.41.0.1              51707916 I
* 19.0.35.0/24          136.41.0.1              51707916 I
* 19.0.36.0/24          136.41.0.1              51707916 I
* 19.0.37.0/24          136.41.0.1              51707916 I
* 19.0.38.0/24          136.41.0.1              51707916 I
* 19.0.39.0/24          136.41.0.1              51707916 I
* 19.0.40.0/24          136.41.0.1              51707916 I
* 19.0.41.0/24          136.41.0.1              51707916 I
* 19.0.42.0/24          136.41.0.1              51707916 I
* 19.0.43.0/24          136.41.0.1              51707916 I
* 19.0.44.0/24          136.41.0.1              51707916 I
* 19.0.45.0/24          136.41.0.1              51707916 I
* 19.0.46.0/24          136.41.0.1              51707916 I
* 19.0.47.0/24          136.41.0.1              51707916 I

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

inet6.0: 52 destinations, 52 routes (52 active, 0 holddown, 0 hidden)
  Prefix                Nexthop              MED      Lclpref    AS path
* fd17:f0f4:f691:d::/80  ::ffff:136.41.0.1      51707916 I
  fd17:f0f4:f691:d:1::/80  ::ffff:136.41.0.1      51707916 I
* fd17:f0f4:f691:d:2::/80  ::ffff:136.41.0.1      51707916 I
```

```

*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:3::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:4::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:5::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:6::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:7::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:8::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:9::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:a::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:b::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:c::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:d::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:e::/80
*          ::ffff:136.41.0.1          51707916 I
  fd17:f0f4:f691:d:f::/80
*          ::ffff:136.41.0.1          51707916 I

```

[edit]

```
lab@R7# run show route receive-protocol bgp 136.41.0.1 hidden
```

```

inet.0: 310 destinations, 311 routes (309 active, 0 holddown, 1 hidden)
  Prefix          Nexthop          MED      Lclpref    AS path
  19.0.64.0/24      136.41.0.1          MED      Lclpref    51707916
5832736 873 I

```

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

```
inet6.0: 52 destinations, 52 routes (52 active, 0 holddown, 0 hidden)
```

In the export direction, only the local summary route is advertised from router R7 to peer P1.

[edit]

```
lab@R7# run show route advertising-protocol bgp 136.41.0.1
```

```

inet.0: 571 destinations, 573 routes (570 active, 0 holddown, 1 hidden)
  Prefix          Nexthop          MED      Lclpref    AS path
  10.50.250.0/28    Self              MED      Lclpref    I

```

14) IPv4 Customers prefixes should be advertised out of your AS to the Upstream peers only. The peering agreement with U3 does not allow you to send any routes with communities to AS 9687.

You have to advertise the customer routes to the Upstream peers only. You have to delete all community attached to the routes. Since BGP by default will advertise the Customer routes to the Upstream peers, you have to apply routing policy router R6 for peer U3 to delete any communities.

a. R6

[edit]

```
lab@R6# set policy-options policy-statement U3.out term delete.communities from
protocol bgp
```

```
[edit]
lab@R6# set policy-options policy-statement U3.out term delete.communities then
community delete match.all
```

To match any possible community you have to use the regular expression used below.

```
[edit]
lab@R6# set policy-options community match.all members *:*
```

```
[edit]
lab@R6# set protocols bgp group U3-v4 export U3.out
```

b. Verify the configuration

As seen at the output below some of the customer routes have communities assigned.

```
[edit]
lab@R6# run show route 193.46.132.0/24 extensive

inet.0: 555 destinations, 943 routes (551 active, 0 holddown, 5 hidden)
193.46.132.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 193.46.132.0/24 -> {indirect(1048574)}
Page 0 idx 1, (group U3-v4 type External) Type 1 val 0xb19a6ac (adv_entry)
  Advertised metrics:
    Nexthop: Self
    AS path: 5673.873 5673.873 5673.873 5673.873 [5673.873] I
    Communities:
Path 193.46.132.0 from 10.50.250.10 Vector len 4. Val: 1
  *BGP      Preference: 170/-101
            Next hop type: Indirect, Next hop index: 0
            Address: 0xa6ed32c
            Next-hop reference count: 32
            Source: 10.50.250.10
            Next hop type: Router, Next hop index: 689
            Next hop: 10.50.0.17 via ge-0/0/4.26, selected
            Label-switched-path r6-to-r2
            Label element ptr: 0xb0ad080
            Label parent element ptr: 0x0
            Label element references: 8
            Label element child references: 4
            Label element lsp id: 5
            Session Id: 0x162
            Protocol next hop: 10.50.250.2
            Indirect next hop: 0xa8f7d60 1048574 INH Session ID: 0x167
            State: <Active Int Ext>
            Local AS: 5673.873 Peer AS: 5673.873
            Age: 55:33      Metric2: 100
            Validation State: unverified
            Task: BGP_5673.873.10.50.250.10
            Announcement bits (3): 0-KRT 9-BGP_RT_Background 10-Resolve tree
2
            AS path: 65456 I (Originator)
            Cluster list: 1.1.1.1
            Originator ID: 10.50.250.2
            Communities: 65456:2022
            Accepted
            Localpref: 100
            Router ID: 10.50.250.10
            Indirect next hops: 1
                Protocol next hop: 10.50.250.2 Metric: 100
                Indirect next hop: 0xa8f7d60 1048574 INH Session ID:
0x167
```

Nevertheless, after the routing policy to delete the communities is applied to the peer U3, the routes are advertised without a community.

```
[edit]
lab@R6# run show route advertising-protocol bgp 195.34.23.17 193.46.132.0/24
extensive

inet.0: 555 destinations, 943 routes (551 active, 0 holddown, 5 hidden)
* 193.46.132.0/24 (1 entry, 1 announced)
  BGP group U3-v4 type External
    Nexthop: Self
    AS path: 5673.873 5673.873 5673.873 5673.873 [5673.873] I
```

- 15) The ingress traffic from AS 87.109 should prefer the link from U2 to R3. Ensure traffic going to C1 prefixes prefer the path through R5. U3 peer must be the least preferred exit point from your AS.

For the next task you need to configure as-path prepend to influence the entry point of traffic to your own AS.

a. R3

```
[edit]
lab@R3# set policy-options policy-statement U1.out term as-path.prepend then as-
path-prepend "5673.873 5673.873"

[edit]
lab@R3# set policy-options policy-statement U2.out term C1.routes.prepend from
protocol bgp

[edit]
lab@R3# set policy-options policy-statement U2.out term C1.routes.prepend from
community C1.routes

[edit]
lab@R3# set policy-options policy-statement U2.out term C1.routes.prepend then
as-path-prepend "5673.873 5673.873"

[edit]
lab@R3# set policy-options policy-statement U2.out term C1.routes.prepend then
accept
```

The prefixes advertised by C1 are matched by community. C1 and C2 prefixes are assigned by communities in Task 17) later on. This is a good example why it is a very good approach to read all tasks before you start to tackle the challenges one by one.

```
[edit]
lab@R3# set policy-options community C1.routes members 65457:2021

[edit]
lab@R3# set protocols bgp group U1 export U1.out

[edit]
lab@R3# set protocols bgp group U2 export U2.out
```

b. R5

```
[edit]
lab@R5# set policy-options policy-statement U2.out term C1.exclude.prepend from
protocol bgp

[edit]
```

```
lab@R5# set policy-options policy-statement U2.out term C1.exclude.prepend from
community C1.routes
```

```
[edit]
```

```
lab@R5# set policy-options policy-statement U2.out term C1.exclude.prepend then
accept
```

```
[edit]
```

```
lab@R5# set policy-options policy-statement U2.out term local.aggregate from
route-filter 10.50.250.0/28 exact
```

```
[edit]
```

```
lab@R5# set policy-options policy-statement U2.out term local.aggregate then as-
path-prepend "5673.873 5673.873"
```

```
[edit]
```

```
lab@R5# set policy-options policy-statement U2.out term local.aggregate then
accept
```

```
[edit]
```

```
lab@R5# set policy-options policy-statement U2.out term as-path.prepend then as-
path-prepend "5673.873 5673.873"
```

```
[edit]
```

```
lab@R5# set protocols bgp group U2 export U2.out
```

c. R6

```
[edit]
```

```
lab@R6# set policy-options policy-statement U3.out term as-path.prepend then as-
path-prepend "5673.873 5673.873 5673.873 5673.873"
```

d. Verify the configuration

The outputs below show that the correct as-path prepend policy is applied to the prefixes advertised to the Upstream peers. Router R6 is the least preferred entry point to your network. The traffic load is split between routers R3 and R5 according to the requirements.

```
[edit]
```

```
lab@R3# run show route advertising-protocol bgp 15.16.17.129
```

```
inet.0: 557 destinations, 1220 routes (553 active, 0 holddown, 8 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 10.50.250.0/28        Self                    MED      Lclpref    5673.873
5673.873 [5673.873] I
* 183.45.32.0/24        Self                    5673.873
5673.873 [5673.873] I
* 183.45.33.0/24        Self                    5673.873
5673.873 [5673.873] I
* 183.45.34.0/24        Self                    5673.873
5673.873 [5673.873] I
* 183.45.35.0/24        Self                    5673.873
5673.873 [5673.873] I
* 183.45.36.0/24        Self                    5673.873
5673.873 [5673.873] I
* 183.45.37.0/24        Self                    5673.873
5673.873 [5673.873] I
* 183.45.38.0/24        Self                    5673.873
5673.873 [5673.873] I
* 183.45.39.0/24        Self                    5673.873
5673.873 [5673.873] I
* 183.45.40.0/24        Self                    5673.873
5673.873 [5673.873] I
* 183.45.41.0/24        Self                    5673.873
5673.873 [5673.873] I
```

```

* 183.45.42.0/24          Self          5673.873
5673.873 [5673.873] I
* 183.45.43.0/24          Self          5673.873
5673.873 [5673.873] I
* 183.45.44.0/24          Self          5673.873
5673.873 [5673.873] I
* 183.45.45.0/24          Self          5673.873
5673.873 [5673.873] I

```

[edit]

```
lab@R3# run show route advertising-protocol bgp 15.16.18.81
```

```

inet.0: 557 destinations, 1220 routes (553 active, 0 holddown, 8 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 10.50.250.0/28        Self            I
* 183.45.32.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.33.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.34.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.35.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.36.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.37.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.38.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.39.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.40.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.41.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.42.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.43.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.44.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.45.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.46.0/24        Self            5673.873
5673.873 [5673.873] I
* 183.45.47.0/24        Self            5673.873
5673.873 [5673.873] I

```

```

* 193.46.132.0/24        Self            I
* 193.46.133.0/24        Self            I
* 193.46.134.0/24        Self            I
* 193.46.135.0/24        Self            I
* 193.46.136.0/24        Self            I
* 193.46.137.0/24        Self            I
* 193.46.138.0/24        Self            I
* 193.46.139.0/24        Self            I
* 193.46.140.0/24        Self            I
* 193.46.141.0/24        Self            I
* 193.46.142.0/24        Self            I
* 193.46.143.0/24        Self            I
* 193.46.144.0/24        Self            I
* 193.46.145.0/24        Self            I
* 193.46.146.0/24        Self            I
* 193.46.147.0/24        Self            I

```

[edit]

```
lab@R5# run show route advertising-protocol bgp 15.16.18.85
```

```
inet.0: 561 destinations, 896 routes (557 active, 0 holddown, 4 hidden)
```

Prefix	Nexthop	MED	Lclpref	AS path
* 157.19.92.128/27	Self			5673.873
5673.873 [5673.873] 9687 1620 33003 I				
* 159.98.63.0/25	Self			5673.873
5673.873 [5673.873] 9687 1620 33112 I				
* 168.79.190.64/26	Self			5673.873
5673.873 [5673.873] 9687 1620 33003 27075 45586 59077 22636 35682 I				
* 168.117.97.32/27	Self			5673.873
5673.873 [5673.873] 9687 1620 61671 36621 54297 35546 56619 21624 I				
* 169.82.37.32/28	Self			5673.873
5673.873 [5673.873] 9687 1620 33112 5697 17954 57862 33909 54284 I				
* 174.19.62.0/25	Self			5673.873
5673.873 [5673.873] 9687 1620 33003 22508 18720 11041 I				
* 176.126.44.0/26	Self			5673.873
5673.873 [5673.873] 9687 1620 61671 41164 28272 56893 33020 I				
* 180.45.182.160/28	Self			5673.873
5673.873 [5673.873] 9687 1620 61671 30404 56153 54005 29201 30202 63861 I				
* 183.45.32.0/24	Self			I
* 183.45.33.0/24	Self			I
* 183.45.34.0/24	Self			I
* 183.45.35.0/24	Self			I
* 183.45.36.0/24	Self			I
* 183.45.37.0/24	Self			I
* 183.45.38.0/24	Self			I
* 183.45.39.0/24	Self			I
* 183.45.40.0/24	Self			I
* 183.45.41.0/24	Self			I
* 183.45.42.0/24	Self			I
* 183.45.43.0/24	Self			I
* 183.45.44.0/24	Self			I
* 183.45.45.0/24	Self			I
* 183.45.46.0/24	Self			I
* 183.45.47.0/24	Self			I
* 187.43.189.64/26	Self			5673.873
5673.873 [5673.873] 9687 1620 61671 62186 25728 I				
* 188.11.87.32/27	Self			5673.873
5673.873 [5673.873] 9687 1620 33112 36621 27260 32724 30096 I				
* 190.132.176.0/25	Self			5673.873
5673.873 [5673.873] 9687 1620 33112 I				
* 193.46.132.0/24	Self			5673.873
5673.873 [5673.873] I				
* 193.46.133.0/24	Self			5673.873
5673.873 [5673.873] I				
* 193.46.134.0/24	Self			5673.873
5673.873 [5673.873] I				
* 193.46.135.0/24	Self			5673.873
5673.873 [5673.873] I				
* 193.46.136.0/24	Self			5673.873
5673.873 [5673.873] I				

[edit]

lab@R6# **run show route advertising-protocol bgp 195.34.23.17**

inet.0: 553 destinations, 930 routes (549 active, 0 holddown, 4 hidden)

Prefix	Nexthop	MED	Lclpref	AS path
* 184.0.0.0/5	Self			5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33003 27075 221 I				
* 184.64.152.0/21	Self			5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33112 63164 37414 53415 32006 15944 I				
* 185.0.0.0/10	Self			5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33112 62186 13238 4101 60484 56440 I				
* 185.187.114.0/24	Self			5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33112 I				


```

* 186.3.32.0/21          Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33112 41164 47078 15085 I
* 187.0.0.0/8           Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33003 I
* 187.44.101.0/24       Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33112 49129 I
* 188.0.0.0/7           Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33112 I
* 188.32.0.0/11         Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33003 63164 56345 48994 56097
22632 53000 9401 I
* 189.18.28.0/22        Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33112 63164 11993 I
* 189.192.0.0/12        Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 61671 36621 50753 24739 26907
37744 I
* 190.47.128.0/17        Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33003 22508 47842 6191 40599
24177 47198 I
* 190.56.32.0/19        Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 61671 30404 11114 53598 10591
4850 54617 I
* 190.192.0.0/12        Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33003 52411 13839 3913 I
* 191.8.0.0/16          Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33112 27075 1973 42864 26191 I
* 191.146.0.0/17        Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 33112 28515 8897 I
* 191.152.0.0/14        Self          5673.873
5673.873 5673.873 5673.873 [5673.873] 87.109 1620 61671 62186 62064 52989 I

```

16) Only advertise all IPv4 Upstream prefixes to C1 and C2 except routes with a mask higher than /28. Make sure the MED value is preserved and sent to customers (C1 and C2).

The task requires advertising the Upstream routes with a mask lower than /28 to peers C1 and C2. Besides that, you have to preserve the metric value of the routes.

a. R4

BGP will advertise all route to peers C1 and C2 without explicit policy, you have to filter the more specific routes only.

```

[edit]
lab@R4# set policy-options policy-statement C1.out term filter.specific from
protocol bgp

```

Rejecting the routes with mask higher than /28 the prefix-list below.

```

[edit]
lab@R4# set policy-options policy-statement C1.out term filter.specific from
route-filter 0.0.0.0/0 prefix-length-range /29-/32

```

```

[edit]
lab@R4# set policy-options policy-statement C1.out term filter.specific then
reject

```

```

[edit]
lab@R4# set policy-options policy-statement C1.out term local.aggregate from
route-filter 10.50.250.0/28 exact

```

```

[edit]
lab@R4# set policy-options policy-statement C1.out term local.aggregate then
accept

```

By default, BGP does not send a metric value to external peers unless explicitly set. However, you have to preserve the same value already assigned to the Upstream routes. To do that you can add a metric value of 0 to the routes advertised to the C1 and C2.

```
[edit]
lab@R4# set policy-options policy-statement C1.out term add.MED then metric add 0
```

```
[edit]
lab@R4# set protocols bgp group C1-v4 export C1.out
```

b. R2

```
[edit]
lab@R2# set policy-options policy-statement C2.out term filter.specific from
protocol bgp
```

```
[edit]
lab@R2# set policy-options policy-statement C2.out term filter.specific from
route-filter 0.0.0.0/0 prefix-length-range /29-/32
```

```
[edit]
lab@R2# set policy-options policy-statement C2.out term filter.specific then
reject
```

```
[edit]
lab@R2# set policy-options policy-statement C2.out term local.aggregate from
route-filter 10.50.250.0/28 exact
```

```
[edit]
lab@R2# set policy-options policy-statement C2.out term local.aggregate then
accept
```

```
[edit]
lab@R2# set policy-options policy-statement C2.out term add.MED then metric add 0
```

```
[edit]
lab@R2# set protocols bgp group C2 export C2.out
```

c. Verify the configuration

The output below proves that the MED value is preserved to the routes advertised to peers C1 and C2.

```
[edit]
lab@R2# run show route advertising-protocol bgp 88.136.0.34
```

```
inet.0: 318 destinations, 318 routes (318 active, 0 holddown, 0 hidden)
  Prefix                Nexthop      MED      Lclpref    AS path
* 0.0.0.0/4              Self                50                371786601
5701741 1620 61671 I
* 0.0.0.0/5              Self                50                371786601
5701741 1620 61671 27075 I
* 1.64.0.0/10            Self                50                371786601
5701741 1620 61671 I
* 1.96.0.0/11            Self                50                371786601
5701741 1620 33112 63164 40776 51777 I
* 1.176.0.0/12           Self                50                371786601
5701741 1620 33112 49129 16320 52954 I
* 2.0.0.0/7              Self                50                371786601
5701741 1620 61671 60177 3091 5721 1770 35283 6098 41407 I
* 3.189.0.0/16           Self                50                371786601
5701741 1620 33112 41164 7136 6370 35389 39823 16965 I
* 4.16.128.0/17          Self                50                371786601
5701741 1620 33003 49129 50466 22491 57331 43943 48281 I
```

```
* 7.0.0.0/8 Self 50 371786601
5701741 1620 33112 52411 37841 5705 22477 16750 I
* 8.0.0.0/5 Self 50 371786601
5701741 1620 33003 60177 56572 I
```

If you do not add the MED value of 0 the routes are advertised without it.

```
[edit]
lab@R2# run show route advertising-protocol bgp 88.136.0.34

inet.0: 549 destinations, 550 routes (547 active, 0 holddown, 2 hidden)
  Prefix            Nexthop          MED      Lclpref    AS path
* 1.64.0.0/10      Self                    5673.873  87.109
1620 61671 I
* 1.84.160.0/20    Self                    5673.873  87.109
1620 33112 I
* 1.96.0.0/11      Self                    5673.873  87.109
1620 33112 63164 40776 51777 I
* 1.161.192.0/21   Self                    5673.873  87.109
1620 33112 30404 32138 45045 I
* 1.176.0.0/12     Self                    5673.873  87.109
1620 33112 49129 16320 52954 I
* 2.0.0.0/7        Self                    5673.873  87.109
1620 61671 60177 3091 5721 1770 35283 6098 41407 I
* 3.189.0.0/16     Self                    5673.873  87.109
1620 33112 41164 7136 6370 35389 39823 16965 I
* 3.214.25.0/24    Self                    5673.873  87.109
1620 61671 62186 11282 7796 I
* 4.16.128.0/17    Self                    5673.873  87.109
1620 33003 49129 50466 22491 57331 43943 48281 I
* 4.84.128.0/18    Self                    5673.873  87.109
1620 33112 63164 28164 56111 39848 I
```

- 17) Because of the peer agreement with the Customers (C1 and C2) all routes advertised by their AS must be tagged with community 65457:2021 for C1 and 65456:2022 for C2. You must use a policy with a reject action to accomplish that. Use these communities to disable the communication between C1 and C2.

You have to eliminate the possibility of peers C1 and C2 to communicate with each other. To do that you have to use the assigned communities. You have to make sure that only the routes with the assigned communities are accepted. The task restricts you to assign the communities by a routing policy using a reject action.

a. R4

```
[edit]
lab@R4# set policy-options policy-statement C1.out term filter.C2.routes from
protocol bgp

[edit]
lab@R4# set policy-options policy-statement C1.out term filter.C2.routes from
community C2.routes

[edit]
lab@R4# set policy-options policy-statement C1.out term filter.C2.routes then
reject

[edit]
lab@R4# set policy-options policy-statement C1.in term only.C1.routes from
protocol bgp

[edit]
lab@R4# set policy-options policy-statement C1.in term only.C1.routes from
community except.C1
```

```
[edit]
lab@R4# set policy-options policy-statement C1.in term only.C1.routes then reject
```

To accomplish this requirement you have to make the community value to match the invert result of the regular expression.

```
[edit]
lab@R4# set policy-options community except.C1 invert-match
```

```
[edit]
lab@R4# set policy-options community except.C1 members 65457:2021
```

```
[edit]
lab@R4# set protocols bgp group C1-v4 import C1.in
```

b. R2

```
[edit]
lab@R2# set policy-options policy-statement C2.out term filter.C1.routes from
protocol bgp
```

```
[edit]
lab@R2# set policy-options policy-statement C2.out term filter.C1.routes from
community C1.routes
```

```
[edit]
lab@R2# set policy-options policy-statement C2.out term filter.C1.routes then
reject
```

```
[edit]
lab@R2# set policy-options policy-statement C2.in term only.C2.routes from
protocol bgp
```

```
[edit]
lab@R2# set policy-options policy-statement C2.in term only.C2.routes from
community except.C2
```

```
[edit]
lab@R2# set policy-options policy-statement C2.in term only.C2.routes then reject
```

```
[edit]
lab@R2# set policy-options community except.C2 invert-match
```

```
[edit]
lab@R2# set policy-options community except.C2 members 65456:2022
```

```
[edit]
lab@R2# set protocols bgp group C2 import C2.in
```

c. Verify the configuration

Only routes with the correct community are matched and therefore accepted.

```
[edit]
lab@R2# run show route receive-protocol bgp 88.136.0.34 extensive
```

```
inet.0: 318 destinations, 318 routes (316 active, 0 holddown, 2 hidden)
* 193.46.132.0/24 (1 entry, 1 announced)
  Accepted
  Nexthop: 88.136.0.34
  AS path: 65456 I
  Communities: 65456:2022

* 193.46.133.0/24 (1 entry, 1 announced)
  Accepted
  Nexthop: 88.136.0.34
```

```

AS path: 65456 I
Communities: 65456:2022

* 193.46.134.0/24 (1 entry, 1 announced)
  Accepted
  Nexthop: 88.136.0.34
  AS path: 65456 I
  Communities: 65456:2022

* 193.46.135.0/24 (1 entry, 1 announced)
  Accepted
  Nexthop: 88.136.0.34
  AS path: 65456 I
  Communities: 65456:2022

```

No C2 routes are advertised from router R2 to the peer C1.

```

[edit]
lab@R2# run show route advertising-protocol bgp 88.136.0.34 183.45.32.0/24

```

Only routes with the correct community are matched and therefore accepted.

```

[edit]
lab@R4# run show route receive-protocol bgp 170.0.15.10 extensive

inet.0: 568 destinations, 569 routes (565 active, 0 holddown, 3 hidden)
* 183.45.32.0/24 (1 entry, 1 announced)
  Accepted
  Nexthop: 170.0.15.10
  AS path: 65457 I
  Communities: 65457:2021

* 183.45.33.0/24 (1 entry, 1 announced)
  Accepted
  Nexthop: 170.0.15.10
  AS path: 65457 I
  Communities: 65457:2021

* 183.45.34.0/24 (1 entry, 1 announced)
  Accepted
  Nexthop: 170.0.15.10
  AS path: 65457 I
  Communities: 65457:2021

* 183.45.35.0/24 (1 entry, 1 announced)
  Accepted
  Nexthop: 170.0.15.10
  AS path: 65457 I
  Communities: 65457:2021

```

None of the C1 routes are advertised from router R2 to the peer C2.

```

[edit]
lab@R4# run show route advertising-protocol bgp 170.0.15.10 193.46.132.0/24

```

18) Do not accept IPv4 routes that are part of RFC1918 or any default routes from any of the external peers.

To protect the network from DDOS attacks you are asked to filter the private address space and any default routes.

a. All eBGP neighbors

```

[edit]

```

```

lab@R3# set policy-options policy-statement U1.in term reject.rfc1918 from route-
filter 10.0.0.0/8 orlonger reject

[edit]
lab@R3# set policy-options policy-statement U1.in term reject.rfc1918 from route-
filter 172.16.0.0/12 orlonger reject

[edit]
lab@R3# set policy-options policy-statement U1.in term reject.rfc1918 from route-
filter 192.168.0.0/16 orlonger reject

[edit]
lab@R3# set policy-options policy-statement U1.in term reject.any.default from
route-filter 0.0.0.0/0 through 0.0.0.0/32

[edit]
lab@R3# set policy-options policy-statement U1.in term reject.any.default then
reject

```

b. Verify the configuration

After the routing policy is applied a few routes are filtered on the border routers.

```

[edit]
lab@R3# run show route hidden

inet.0: 528 destinations, 1195 routes (528 active, 0 holddown, 8 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/4          [BGP ] 00:01:48, MED 150, localpref 100
                   AS path: 5701741 1620 61671 I, validation-state: unverified
> to 15.16.18.81 via ge-0/0/5.2002
0.0.0.0/5          [BGP ] 00:01:48, MED 50, localpref 100
                   AS path: 5701741 1620 61671 I, validation-state: unverified
> to 15.16.17.129 via ge-0/0/5.2001
0.0.0.0/5          [BGP ] 00:01:48, MED 150, localpref 100
                   AS path: 5701741 1620 61671 27075 I, validation-state:
unverified
> to 15.16.18.81 via ge-0/0/5.2002
10.0.0.0/10        [BGP ] 00:01:48, MED 50, localpref 100
                   AS path: 5701741 1620 61671 27075 I, validation-state:
unverified
> to 15.16.17.129 via ge-0/0/5.2001
10.0.0.0/10        [BGP ] 00:01:48, MED 150, localpref 100
                   AS path: 5701741 1620 33112 60177 54491 19593 62822 59959
22793 22889 I, validation-state: unverified
> to 15.16.18.81 via ge-0/0/5.2002
10.140.184.64/26  [BGP ] 00:01:48, MED 50, localpref 100
                   AS path: 5701741 1620 33112 60177 54491 19593 62822 59959
22793 22889 I, validation-state: unverified
> to 15.16.17.129 via ge-0/0/5.2001
10.140.184.64/26  [BGP ] 00:01:48, MED 150, localpref 100
                   AS path: 5701741 1620 61671 I, validation-state: unverified
> to 15.16.18.81 via ge-0/0/5.2002
10.140.184.64/26  [BGP ] 00:01:48, MED 150, localpref 100
                   AS path: 5701741 1620 61671 I, validation-state: unverified
> to 15.16.17.129 via ge-0/0/5.2001

```

19) Ensure that 4byte AS numbers are displayed in as-dot notation format on all routers.

JUNOS supports representing the 4byte AS numbers in two different formats, the default numeric format and AS-dot format. The task requires to change the format to AS-dot representation.

- a. All devices including the route-reflector

```
[edit]
lab@R1# set routing-options autonomous-system asdot-notation
```

- b. Verify the configuration

After the command above is applied to all routers, the representation of the 4bytes AS numbers in all outputs changed to AS-dot format.

```
[edit]
lab@R3# run show route advertising-protocol bgp 15.16.17.129

inet.0: 557 destinations, 1220 routes (553 active, 0 holddown, 8 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 10.50.250.0/28         Self              0
5673.873 [5673.873] I
* 183.45.32.0/24         Self              0
5673.873 [5673.873] I
* 183.45.33.0/24         Self              0
5673.873 [5673.873] I
* 183.45.34.0/24         Self              0
5673.873 [5673.873] I
* 183.45.35.0/24         Self              0
5673.873 [5673.873] I
* 183.45.36.0/24         Self              0
5673.873 [5673.873] I
* 183.45.37.0/24         Self              0
5673.873 [5673.873] I
```

Solution - Task 3.4: BGP general requirements

20) No hidden (except filtered by applied policies) or unresolved routes should be present in AS 5673.873.

21) No static routes are allowed if not explicit permitted by a particular task.

22) Ensure optimal routing in the network.

It is recommended at the end of the section to do some checks to ensure all tasks are completed and solved correctly.

- a. Route-reflector

If you carefully verify the number of the routes received by the route-reflector from the border routers, you will notice that they are in the range of 460 unique routes.

```
[edit]
lab@route-reflector# run show bgp summary
Groups: 5 Peers: 8 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State    Pending
inet.0
                1179        514          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.1     371786601      884      1810        0        0      5:18:42
0/0/0/0         0/0/0/0
10.50.250.2     371786601      885      1534        0        0      5:18:29
18/18/18/0      0/0/0/0
```

```

10.50.250.3      371786601      1547      962      0      0      5:07:18
256/460/460/0    0/0/0/0
10.50.250.4      371786601      679      1805      0      0      5:18:38
20/20/20/0      0/0/0/0
10.50.250.5      371786601      1179     1594      0      2      5:06:22
123/205/205/0    0/0/0/0
10.50.250.6      371786601      1474     1712      0      1      5:08:49
81/460/460/0     0/0/0/0
10.50.250.7      371786601      721      967      0      1      5:08:50
16/16/16/0       0/0/0/0
10.50.250.8      371786601      663      965      0      1      5:08:49
0/0/0/0          0/0/0/0

```

However, if you check the same output on routers R2 and R7 you will see that the routes received from the route-reflector are almost half.

```

[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
                287        285          0          0          0          0
inet6.0
                16          0          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last  Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10    5673.873      228        6        0        1          11
Establ
inet.0: 269/269/269/0
88.136.0.34     65456         35       774        0        0         15:02
Establ
inet.0: 16/18/16/0

```

```

[edit]
lab@R7# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                259        257          0          0          0          0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last  Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10    5673.873      283        9        0        1         1:50
Establ
inet.0: 241/242/242/0
136.41.0.1      789.12         36       39        0        0         15:13
Establ
inet.0: 16/17/16/0
inet6.0: 16/16/16/0

```

Obviously something is wrong, because it is expected all the routes to be reflected by the route-reflector. Verify the configuration on the route-reflector. It turns out that the initial configuration was applied using apply-groups to disable the client reflection for routers R2 and R7.

```

groups {
  initial {
    protocols {
      bgp {
        group cluster-1 {
          type internal;
          neighbor 10.50.250.2 {
            no-client-reflect;
          }
        }
        group cluster-2 {

```


Deleting the configuration groups should fix route reflection to routers R2 and R7.

```
[edit]
lab@R7# run show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State      Pending
inet.0
                512         510           0           0           0
inet6.0
                17          16           0           0           0
Peer           AS         InPkt      OutPkt      OutQ      Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10   5673.873         503         11         0         2         2:48
Establ 0/0/0
  inet.0: 494/495/495/0
136.41.0.1     789.12          57         63         0         0         25:12
Establ
  inet.0: 16/17/16/0
  inet6.0: 16/16/16/0
```

49)

Part 4: MPLS configuration

The following part focuses 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.

Solution - Task 4.1: MPLS and RSVP configuration

- 1) Enable the RSVP protocol on all routers in OSPF except R8 and the route-reflector.
Ensure that additional information like link coloring and bandwidth availability is exchanged.
 - a. All devices apart from R8 and the route-reflector

Enabling RSVP protocol is straightforward, as configuring the interfaces required by the Task under the `protocols rsvp` stanza.

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.12

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.14

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.16
```

OSPF protocol does not exchange traffic engineering information by default. You have to enable the functionality explicitly.

```
[edit]
lab@R1# set protocols ospf traffic-engineering
```

- b. Verify the configuration

The RSVP interface table displays the interface enabled for RSVP.

```
[edit]
lab@R1# run show rsvp interface
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW
ge-0/0/4.12	Up	2	100%	700Mbps	700Mbps	0bps
ge-0/0/4.14	Up	1	100%	700Mbps	700Mbps	0bps
ge-0/0/4.16	Up	1	100%	300Mbps	300Mbps	0bps

The traffic engineering database is populated using OSPF. The information is used by constrained shortest path first algorithm (cspf) to build LSPs with additional path requirements.

```
[edit]
lab@R1# run show ted database
TED database: 0 ISIS nodes 5 INET nodes
```

ID	Type	Age(s)	LnkIn	LnkOut	Protocol
10.50.250.1	Rtr	723	3	3	OSPF(0.0.0.0)

To: 10.50.250.2, Local: 10.50.0.1, Remote: 10.50.0.2
Local interface index: 332, Remote interface index: 0

```

To: 10.50.250.4, Local: 10.50.0.5, Remote: 10.50.0.6
Local interface index: 333, Remote interface index: 0
To: 10.50.250.6, Local: 10.50.0.9, Remote: 10.50.0.10
Local interface index: 334, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.50.250.2                       Rtr   337    3    3 OSPF(0.0.0.0)
To: 10.50.250.1, Local: 10.50.0.2, Remote: 10.50.0.1
Local interface index: 332, Remote interface index: 0
To: 10.50.250.4, Local: 10.50.0.13, Remote: 10.50.0.14
Local interface index: 333, Remote interface index: 0
To: 10.50.250.6, Local: 10.50.0.17, Remote: 10.50.0.18
Local interface index: 334, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.50.250.4                       Rtr  9464    4    4 OSPF(0.0.0.0)
To: 10.50.250.2, Local: 10.50.0.14, Remote: 10.50.0.13
Local interface index: 333, Remote interface index: 0
To: 10.50.250.1, Local: 10.50.0.6, Remote: 10.50.0.5
Local interface index: 332, Remote interface index: 0
To: 10.50.250.6, Local: 10.50.0.37, Remote: 10.50.0.38
Local interface index: 335, Remote interface index: 0
To: 10.50.250.5, Local: 10.50.0.33, Remote: 10.50.0.34
Local interface index: 334, Remote interface index: 0

```

The RSVP neighbor table is displayed on router R1. However, the neighbor sessions will be established when there is a LSP signaled.

```

[edit]
lab@R1# run show rsvp neighbor
RSVP neighbor: 3 learned
Address      Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.50.0.2    0  1/0    1:30:28      9   506/506   544
10.50.0.6    5  1/0    1:30:18      9   473/473   15
10.50.0.10   0  1/0    1:30:24      9   599/599   503

```

2) Configure the appropriate colors of the links according the diagram. Also configure the available bandwidth for the LSPs as follows:

- GOLD = 700Mbps.
- SILVER = 500Mbps.
- BRONZE = 300Mbps.

a. All devices apart from R8 and the route-reflector

Administrative groups or link coloring are manually assigned to the interfaces and used to calculate the paths for LSP's in the network. The name of the groups and the values must be consistent on all devices from the MPLS domain.

First you have to define the groups with the values assigned to each of it.

```

[edit]
lab@R1# set protocols mpls admin-groups GOLD 1

[edit]
lab@R1# set protocols mpls admin-groups SILVER 2

[edit]
lab@R1# set protocols mpls admin-groups BRONZE 3

```

The line below will prevent the mpls protocol to be enabled on the fxp0.0 interface.

```

[edit]

```

```
lab@R1# set protocols mpls interface fxp0.0 disable
```

Family mpls was already enabled on all interface through apply-groups. However, to successfully enable labeled traffic to run over the interfaces, you have to add them under the protocols mpls stanza.

```
[edit]
lab@R1# set protocols mpls interface all
```

Assign the link colors to the interfaces as shown on the diagram.

```
[edit]
lab@R1# set protocols mpls interface ge-0/0/4.12 admin-group GOLD
```

```
[edit]
lab@R1# set protocols mpls interface ge-0/0/4.14 admin-group GOLD
```

```
[edit]
lab@R1# set protocols mpls interface ge-0/0/4.16 admin-group BRONZE
```

The bandwidth used by RSVP for creating LSPs are deliberately made lower to prevent congestion. The configuration lines to achieve it, are shown below.

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.12 bandwidth 700m
```

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.14 bandwidth 700m
```

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.16 bandwidth 300m
```

b. Verify the configuration

The assigned administrative group and the state of the MPLS enabled interfaces are checked with the commands below.

```
[edit]
lab@R1# run show mpls interface
Interface      State      Administrative groups (x: extended)
ge-0/0/4.12    Up         GOLD
ge-0/0/4.14    Up         GOLD
ge-0/0/4.16    Up         BRONZE
```

```
[edit]
lab@R2# run show mpls interface
Interface      State      Administrative groups (x: extended)
ge-0/0/4.12    Up         GOLD
ge-0/0/4.24    Up         BRONZE
ge-0/0/4.26    Up         SILVER
```

```
[edit]
lab@R3# run show mpls interface
Interface      State      Administrative groups (x: extended)
ge-0/0/4.35    Up         <none>
ge-0/0/4.37    Up         <none>
ge-0/0/5.2001  Up         <none>
ge-0/0/5.2002  Up         <none>
ae3.0          Up         <none>
```

```
[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						
ge-0/0/4.12	Up	2	100%	700Mbps	700Mbps	0bps
0bps						
ge-0/0/4.14	Up	1	100%	700Mbps	700Mbps	0bps
0bps						
ge-0/0/4.16	Up	1	100%	300Mbps	300Mbps	0bps
0bps						

3) Configure a LSP using the criteria below:

Next few steps requires from you to configure a number LSPs following strict criteria.

- Configure LSPs named *r1-to-r5* and *r5-to-r1* connecting R1 with R5 without using CSPF. However, you need to ensure they cross the link between R1 and R6. Protect those LSPs from failure of the link between R1 and R6. You have to ensure that the protecting LSPs have at least 100Mbps of available bandwidth.

a. R1

First step is configuring the destination of the LSP. Configure a short but meaningful name for the LSP helps with troubleshooting.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r5 to 10.50.250.5
```

The CSPF protocol is enabled by default for RSVP LSPs, therefore it must be disabled explicitly.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r5 no-cspf
```

`node-link-protection` will request protection along the path. This is required to fulfil the requirement of protecting the link between R1 and R6.

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r5 node-link-protection
```

```
[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r5 primary to-R6
```

```
[edit]
lab@R1# set protocols mpls path to-R6 10.50.0.10 strict
```

Configure a bypass LSP protecting the path from R1 to R6.

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.16 link-protection bypass to-r6 to 10.50.250.6
```

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.16 link-protection bypass to-r6 bandwidth 100m
```

b. R5

The LSPs are unidirectional, therefore LSPs should also be configured in the opposite direction. R5 should have similar configuration except for the RSVP bypass part, because it does not have direct connection to the segment that has to be protected.

```
[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r1 to 10.50.250.1

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r1 no-cspf

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r1 link-protection

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r1 primary to-R6
```

The path must force the LSP to go over the link between R6 and R1.

```
[edit]
lab@R5# set protocols mpls path to-R6 10.50.250.6 loose

[edit]
lab@R5# set protocols mpls path to-R6 10.50.0.9 strict
```

c. R6

To protect the LSP from router R5 to R1, you have to configure a bypass LSP originated on router R6.

```
[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.16 link-protection bypass to-r1 to
10.50.250.1

[edit]
lab@R6# set protocols rsvp interface ge-0/0/4.16 link-protection bypass to-r1
bandwidth 100m
```

d. Verify the configuration

The output below show the result of the configuration and the successful completion of the task.

```
[edit]
lab@R1# run show mpls lsp extensive
Ingress LSP: 3 sessions

10.50.250.5
  From: 10.50.250.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: to-R6 (primary)
  Link protection desired
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary to-R6 State: Up
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
  10.50.250.6(flag=0x20) 10.50.0.10(Label=299782) 10.50.250.5(flag=0x20)
10.50.0.45(Label=3)
  10 Dec 3 11:56:52.094 Link-protection Up
  9 Dec 3 11:56:43.586 Record Route: 10.50.250.6(flag=0x20)
10.50.0.10(Label=299782) 10.50.250.5(flag=0x20) 10.50.0.45(Label=3)
  8 Dec 3 11:56:43.586 Up
  7 Dec 3 11:56:43.078 Originate Call
  6 Dec 3 11:56:43.077 Clear Call
  5 Dec 3 09:31:17.292 Selected as active path
```

```

4 Dec 3 09:31:17.292 Record Route: 10.50.250.4(flag=0x20)
10.50.0.6(Label=299777) 10.50.250.5(flag=0x20) 10.50.0.34(Label=3)
3 Dec 3 09:31:17.291 Up
2 Dec 3 09:30:31.804 No Route toward dest[7 times]
1 Dec 3 09:28:07.782 Originate Call
Created: Sat Dec 3 09:28:07 2016

[edit]
lab@R5# run show mpls lsp extensive
Ingress LSP: 2 sessions

10.50.250.1
  From: 10.50.250.5, State: Up, ActiveRoute: 0, LSPname: r5-to-r1
  ActivePath: to-R6 (primary)
  Link protection desired
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary to-R6 State: Up
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
  20=Node-ID):
    10.50.250.6(flag=0x21) 10.50.0.46(flag=1 Label=299783)
  10.50.250.1(flag=0x20) 10.50.0.9(Label=3)
    10 Dec 3 11:25:05.280 Record Route: 10.50.250.6(flag=0x21) 10.50.0.46(flag=1
  Label=299783) 10.50.250.1(flag=0x20) 10.50.0.9(Label=3)
    9 Dec 3 11:25:02.122 Record Route: 10.50.250.6(flag=0x20)
  10.50.0.46(Label=299783) 10.50.250.1(flag=0x20) 10.50.0.9(Label=3)
    8 Dec 3 11:25:02.122 Up
    7 Dec 3 11:25:01.679 Originate Call
    6 Dec 3 11:25:01.677 Clear Call
    5 Dec 3 09:11:50.213 Selected as active path
    4 Dec 3 09:11:50.212 Record Route: 10.50.250.4(flag=0x20)
  10.50.0.33(Label=299776) 10.50.250.1(flag=0x20) 10.50.0.5(Label=3)
    3 Dec 3 09:11:50.212 Up
    2 Dec 3 09:11:04.991 No Route toward dest[4 times]
    1 Dec 3 09:10:55.986 Originate Call
Created: Sat Dec 3 09:10:55 2016

```

The bypass LSP is established to protect the link between R1 and R6. The LSP path is signaled through router R2.

```

[edit]
lab@R1# run show mpls lsp bypass extensive
Ingress LSP: 1 sessions

10.50.250.6
  From: 10.50.250.1, LSPstate: Up, ActiveRoute: 0
  LSPname: to-r6
  LSPTYPE: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299779
  Resv style: 1 SE, Label in: -, Label out: 299779
  Time left: -, Since: Sat Dec 3 11:56:46 2016
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 28595 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  ActiveResv 1, PreemptionCnt 0, Update threshold 0%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 100Mbps
    MaxAvailableBW 100Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  PATH rcvfrom: localclient

```

```

Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 10.50.0.2 (ge-0/0/4.12) 1 pkts
RESV rcvfrom: 10.50.0.2 (ge-0/0/4.12) 1 pkts, Entropy label: Yes
Explct route: 10.50.0.2 10.50.0.18
Record route: <self> 10.50.0.2 10.50.0.18
  4 Dec  3 11:56:46 Record Route:  10.50.0.2 10.50.0.18
  3 Dec  3 11:56:46 Up
  2 Dec  3 11:56:46 CSPF: computation result accepted
  1 Dec  3 11:56:46 Originate Call
Total 1 displayed, Up 1, Down 0

```

In the opposite direction from router R5 to R1, an other LSP is established again through router R2.

```

[edit]
lab@R6# run show mpls lsp bypass extensive
Ingress LSP: 1 sessions

```

```

10.50.250.1
  From: 10.50.250.6, LSPstate: Up, ActiveRoute: 0
  LSPname: to-r1
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299780
  Resv style: 1 SE, Label in: -, Label out: 299780
  Time left: -, Since: Sat Dec  3 10:54:55 2016
  Tspec: rate 100Mbps size 100Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 55570 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  ActiveResv 1, PreemptionCnt 0, Update threshold 0%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 100Mbps
    MaxAvailableBW 100Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.50.0.17 (ge-0/0/4.26) 1 pkts
  RESV rcvfrom: 10.50.0.17 (ge-0/0/4.26) 1 pkts, Entropy label: Yes
  Explct route: 10.50.0.17 10.50.0.1
  Record route: <self> 10.50.0.17 10.50.0.1
    4 Dec  3 10:54:55 Record Route:  10.50.0.17 10.50.0.1
    3 Dec  3 10:54:55 Up
    2 Dec  3 10:54:55 CSPF: computation result accepted
    1 Dec  3 10:54:53 Originate Call
Total 1 displayed, Up 1, Down 0

```

```

[edit]
lab@R6# run show rsvp interface ge-0/0/4.16 extensive
ge-0/0/4.16 Index 338, State Ena/Up
  NoAuthentication, Aggregate, Reliable, LinkProtection
  HelloInterval 9(second)
  Address 10.50.0.10
  ActiveResv 3, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 300Mbps
  ct0: StaticBW 300Mbps, AvailableBW 300Mbps
    MaxAvailableBW 300Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: On, Bypass: 1, LSP: 1, Protected LSP: 1, Unprotected LSP: 0
    1 Dec 13 12:56:47 New bypass to-r1
  Bypass: to-r1, State: Up, Type: LP, LSP: 1, Backup: 0
    4 Dec 13 12:56:48 Record Route:  10.50.0.17 10.50.0.1
    3 Dec 13 12:56:48 Up

```



```

2 Dec 13 12:56:48 CSPF: computation result accepted
1 Dec 13 12:56:47 Originate Call

[edit]
lab@R1# run show rsvp interface ge-0/0/4.16 extensive
ge-0/0/4.16 Index 338, State Ena/Up
  NoAuthentication, Aggregate, Reliable, LinkProtection
  HelloInterval 9(second)
  Address 10.50.0.9
  ActiveResv 2, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 300Mbps
  ct0: StaticBW 300Mbps, AvailableBW 300Mbps
    MaxAvailableBW 300Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: On, Bypass: 1, LSP: 1, Protected LSP: 0, Unprotected LSP: 1
    2 Dec 13 16:39:51 No suitable bypass with sufficient bandwidth Bypass-
>10.50.0.10->10.50.0.45[11 times]
    1 Dec 13 14:37:01 New bypass to-r6
  Bypass: to-r6, State: Up, Type: LP, LSP: 0, Backup: 0
    4 Dec 13 14:37:02 Record Route: 10.50.0.2 10.50.0.18
    3 Dec 13 14:37:02 Up
    2 Dec 13 14:37:02 CSPF: computation result accepted
    1 Dec 13 14:37:01 Originate Call

```

- Configure LSPs named *r2-to-r5* and *r5-to-r2* providing connectivity between R2 and R5. Make sure to restrict those LSPs to use links with maximum available bandwidth. Protect the LSPs with one-to-one backup method not shared by other LSPs. The protecting LSPs should not be restricted by any means.

a. R2

Next pair of LSPs should be configured between routers R2 and R5. You have to establish the LSPs over the links with the most bandwidth and protect them using fast-reroute. The fast-reroute protection method must be used instead of node protection, because of the requirement of one-to-one protection not shared by other LSPs. Node or Link protection is not considered because it provides facility backup for all the LSPs that need protection.

```

[edit]
lab@R2# set protocols mpls label-switched-path r2-to-r5 to 10.50.250.5

```

Specify the links colored with the highest bandwidth among all possible paths between R2 and R5.

```

[edit]
lab@R2# set protocols mpls label-switched-path r2-to-r5 admin-group include-any GOLD

```

```

[edit]
lab@R2# set protocols mpls label-switched-path r2-to-r5 admin-group include-any SILVER

```

The protection LSPs are configured to not inherit any constraints imposed by the primary LSP.

```

[edit]
lab@R2# set protocols mpls label-switched-path r2-to-r5 fast-reroute no-include-any

```

b. R5

Similar configuration is applied to router R5 for the opposite direction LSP.

```
[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r2 to 10.50.250.2

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r2 admin-group include-any
GOLD

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r2 admin-group include-any
SILVER

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r2 fast-reroute no-include-
any
```

c. Verify the configuration

The output below displays the result of the configuration applied to routers R2 and R5. The LSPs are established over the path with highest available bandwidth. In addition the LSPs are protected by detour LSPs.

```
[edit]
lab@R2# run show mpls lsp extensive
Ingress LSP: 3 sessions

10.50.250.5
  From: 10.50.250.2, State: Up, ActiveRoute: 0, LSPname: r2-to-r5
  ActivePath: (primary)
  FastReroute desired
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Include Any: SILVER GOLD
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 200)
    10.50.0.18 S 10.50.0.45 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
    20=Node-ID):
    10.50.0.18(flag=1) 10.50.0.45
  19 Dec 3 10:46:12.427 Make-before-break: Cleaned up old instance: Hold dead
  expiry
  18 Dec 3 10:44:47.402 Record Route: 10.50.0.18(flag=1) 10.50.0.45
  17 Dec 3 10:44:47.298 Fast-reroute Detour Up
  16 Dec 3 10:44:45.340 Make-before-break: Switched to new instance
  15 Dec 3 10:44:45.340 Fast-reroute Detour Down
  14 Dec 3 10:44:44.333 Record Route: 10.50.0.18 10.50.0.45
  13 Dec 3 10:44:44.333 Up
  12 Dec 3 10:44:43.899 Originate make-before-break call
  11 Dec 3 10:44:43.899 CSPF: computation result accepted 10.50.0.18
10.50.0.45
  10 Dec 3 10:44:41.603 CSPF failed: no route toward 10.50.250.5
  9 Dec 3 10:44:41.603 CSPF: link down/deleted:
10.50.0.46(10.50.250.6:335) (10.50.250.6)->10.50.0.45(10.50.250.5:0) (10.50.250.5)
  8 Dec 3 08:59:24.967 Record Route: 10.50.0.18(flag=1) 10.50.0.45
  7 Dec 3 08:59:24.876 Fast-reroute Detour Up
  6 Dec 3 08:59:21.895 Selected as active path
  5 Dec 3 08:59:21.894 Record Route: 10.50.0.18 10.50.0.45
  4 Dec 3 08:59:21.894 Up
  3 Dec 3 08:59:21.463 Originate Call
  2 Dec 3 08:59:21.463 CSPF: computation result accepted 10.50.0.18
10.50.0.45
  1 Dec 3 08:58:51.745 CSPF failed: no route toward 10.50.250.5[4 times]
```

Created: Sat Dec 3 08:56:54 2016

[edit]

lab@R2# **run show RSVP session extensive**

Ingress RSVP: 3 sessions

10.50.250.5

From: 10.50.250.2, LSPstate: Up, ActiveRoute: 0

LSPname: r2-to-r5, LSPpath: Primary

LSPTYPE: Static Configured

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: 299780

Resv style: 1 FF, Label in: -, Label out: 299780

Time left: -, Since: Sat Dec 3 10:44:43 2016

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Port number: sender 2 receiver 64055 protocol 0

FastReroute desired

PATH rcvfrom: localclient

Adspec: sent MTU 1500

Path MTU: received 1500

PATH sentto: 10.50.0.18 (ge-0/0/4.26) 4 pkts

RESV rcvfrom: 10.50.0.18 (ge-0/0/4.26) 2 pkts, Entropy label: Yes

Explot route: 10.50.0.18 10.50.0.45

Record route: <self> 10.50.0.18 10.50.0.45

Detour is Up

Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Detour adspec: sent MTU 1500

Path MTU: received 1500

Detour PATH sentto: 10.50.0.14 (ge-0/0/4.24) 3 pkts

Detour RESV rcvfrom: 10.50.0.14 (ge-0/0/4.24) 2 pkts, Entropy label: Yes

Detour Explot route: 10.50.0.14 10.50.0.34

Detour Record route: <self> 10.50.0.14 10.50.0.34

Detour Label out: 299782

[edit]

lab@R5# **run show MPLS lsp extensive**

Ingress LSP: 2 sessions

10.50.250.2

From: 10.50.250.5, State: Up, ActiveRoute: 0, LSPname: r5-to-r2

ActivePath: (primary)

FastReroute desired

LSPTYPE: Static Configured, Penultimate hop popping

LoadBalance: Random

Encoding type: Packet, Switching type: Packet, GPID: IPv4

*Primary State: Up

Priorities: 7 0

SmartOptimizeTimer: 180

Include Any: SILVER GOLD

Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 200)

10.50.0.46 S 10.50.0.17 S

Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):

10.50.0.46(flag=1) 10.50.0.17

19 Dec 3 11:09:31.498 Make-before-break: Cleaned up old instance: Hold dead expiry

18 Dec 3 11:08:07.888 Record Route: 10.50.0.46(flag=1) 10.50.0.17

17 Dec 3 11:08:07.309 Fast-reroute Detour Up

16 Dec 3 11:08:05.328 Make-before-break: Switched to new instance

15 Dec 3 11:08:05.328 Fast-reroute Detour Down

14 Dec 3 11:08:04.321 Record Route: 10.50.0.46 10.50.0.17

13 Dec 3 11:08:04.321 Up

12 Dec 3 11:08:03.860 Originate make-before-break call

11 Dec 3 11:08:03.860 CSPF: computation result accepted 10.50.0.46

10.50.0.17

10 Dec 3 11:07:37.571 CSPF failed: no route toward 10.50.250.2

9 Dec 3 11:07:37.571 CSPF: link down/deleted:

10.50.0.18(10.50.250.6:333) (10.50.250.6)->10.50.0.17(10.50.250.2:0) (10.50.250.2)

```

8 Dec 3 09:11:58.951 Record Route: 10.50.0.46(flag=1) 10.50.0.17
7 Dec 3 09:11:58.379 Fast-reroute Detour Up
6 Dec 3 09:11:55.395 Selected as active path
5 Dec 3 09:11:55.394 Record Route: 10.50.0.46 10.50.0.17
4 Dec 3 09:11:55.394 Up
3 Dec 3 09:11:55.030 Originate Call
2 Dec 3 09:11:55.030 CSPF: computation result accepted 10.50.0.46
10.50.0.17
1 Dec 3 09:11:25.503 CSPF failed: no route toward 10.50.250.2[2 times]
Created: Sat Dec 3 09:10:55 2016
Total 2 displayed, Up 2, Down 0

[edit]
lab@R5# run show rsvp session extensive
Ingress RSVP: 2 sessions

10.50.250.2
From: 10.50.250.5, LSPstate: Up, ActiveRoute: 0
LSPname: r5-to-r2, LSPpath: Primary
LSptype: Static Configured
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 299781
Resv style: 1 FF, Label in: -, Label out: 299781
Time left: -, Since: Sat Dec 3 11:08:03 2016
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 4085 protocol 0
FastReroute desired
PATH rcvfrom: localclient
Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 10.50.0.46 (ge-0/0/4.56) 5 pkts
RESV rcvfrom: 10.50.0.46 (ge-0/0/4.56) 2 pkts, Entropy label: Yes
Explot route: 10.50.0.46 10.50.0.17
Record route: <self> 10.50.0.46 10.50.0.17
Detour is Up
Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Detour adspec: sent MTU 1500
Path MTU: received 1500
Detour PATH sentto: 10.50.0.33 (ge-0/0/4.45) 3 pkts
Detour RESV rcvfrom: 10.50.0.33 (ge-0/0/4.45) 1 pkts, Entropy label: Yes
Detour Explot route: 10.50.0.33 10.50.0.13
Detour Record route: <self> 10.50.0.33 10.50.0.13
Detour Label out: 299784
Total 2 displayed, Up 2, Down 0

```

- Configure an LSP between R4 and R6 named *r4-to-r6* and *r6-to-r4* having a secondary path different from the primary. Make sure that the primary path is established only over links marked with color SILVER, whereas the secondary path must be using links marked with GOLD and BRONZE. The LSPs should have at least 250Mbps of available bandwidth. If the primary path is preempted for any reason, ensure that the secondary path is established beforehand. However, to preserve resources on the routers, the secondary path must not be signaled if not needed.

a. R4

The LSPs between routers R4 and R6, on per direction, must follow the given requirements.

```

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r6 to 10.50.250.6

```

The `adaptive` options accomplishes the requirement of keeping the old path established until the new rerouted path is calculated and ready to take the traffic. In addition to that it also avoids the double counting of the reserved bandwidth.

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r6 adaptive
```

Requesting bandwidth reservation along the path of the LSP.

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r6 bandwidth 250m
```

The primary path is configured to use only links colored with SILVER administrative-group.

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r6 primary to-r6-primary
admin-group include-all SILVER
```

The secondary path is configured to be signaled over any links having administrative-groups GOLD or BRONZE. That path will not be established until the primary path is UP, since the `standby` option is not used.

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r6 secondary to-r6-secondary
admin-group include-any GOLD
```

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r6 secondary to-r6-secondary
admin-group include-any BRONZE
```

Empty primary and secondary paths are configured.

```
[edit]
lab@R4# set protocols mpls path to-r6-secondary
```

```
[edit]
lab@R4# set protocols mpls path to-r6-primary
```

b. R6

On R6 an LSP is created towards R4

```
[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r4 to 10.50.250.4
```

```
[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r4 bandwidth 250m
```

```
[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r4 adaptive
```

```
[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r4 primary to-r4-primary
admin-group include-all SILVER
```

```
[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r4 secondary to-r4-secondary
admin-group include-all GOLD
```

```
[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r4 secondary to-r4-secondary
admin-group include-all BRONZE
```

```
[edit]
```

```
lab@R6# set protocols mpls path to-r4-primary

[edit]
lab@R6# set protocols mpls path to-r4-secondary
```

c. Verify the configuration

The output below confirms a path is signaled between routers R4 and R6.

```
[edit]
lab@R4# run show mpls lsp extensive
Ingress LSP: 2 sessions

10.50.250.6
  From: 10.50.250.4, State: Up, ActiveRoute: 0, LSPname: r4-to-r6
  ActivePath: to-r6-primary (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary to-r6-primary State: Up
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Include All: SILVER
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 100)
  10.50.0.38 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
  20=Node-ID):
    10.50.0.38
    6 Dec 3 09:08:20.152 Selected as active path
    5 Dec 3 09:08:20.152 Record Route: 10.50.0.38
    4 Dec 3 09:08:20.151 Up
    3 Dec 3 09:08:20.034 Originate Call
    2 Dec 3 09:08:20.034 CSPF: computation result accepted 10.50.0.38
    1 Dec 3 09:07:51.012 CSPF failed: no route toward 10.50.250.6[3 times]
  Secondary to-r6-secondary State: Dn
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Include Any: BRONZE GOLD
  No computed ERO.
    6 Dec 3 09:08:50.030 Clear Call
    5 Dec 3 09:08:20.442 Record Route: 10.50.0.5 10.50.0.10
    4 Dec 3 09:08:20.442 Up
    3 Dec 3 09:08:20.035 Originate Call
    2 Dec 3 09:08:20.035 CSPF: computation result accepted 10.50.0.5 10.50.0.10
    1 Dec 3 09:07:51.012 CSPF failed: no route toward 10.50.250.6[2 times]
  Created: Sat Dec 3 09:06:51 2016

[edit]
lab@R6# run show mpls lsp extensive
Ingress LSP: 4 sessions

10.50.250.4
  From: 10.50.250.6, State: Up, ActiveRoute: 0, LSPname: r6-to-r4
  ActivePath: to-r4-primary (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary to-r4-primary State: Up
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Include All: SILVER
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 100)
  10.50.0.37 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
  20=Node-ID):
    10.50.0.37
    6 Dec 3 08:56:39.543 Selected as active path
```

```

5 Dec 3 08:56:39.542 Record Route: 10.50.0.37
4 Dec 3 08:56:39.542 Up
3 Dec 3 08:56:39.340 Originate Call
2 Dec 3 08:56:39.340 CSPF: computation result accepted 10.50.0.37
1 Dec 3 08:56:09.414 CSPF failed: no route toward 10.50.250.4
Secondary-to-r4-secondary State: Dn
Priorities: 7 0
SmartOptimizeTimer: 180
Include All: BRONZE GOLD
No computed ERO.
1 Dec 3 08:56:39.340 CSPF failed: no route toward 10.50.250.4
Created: Sat Dec 3 08:56:09 2016

```

The RSVP sessions are displayed as Shared Explicit (SE), caused by the adaptive option configured. By default the style is fixed filter (FF).

```

[edit]
lab@R4# run show rsvp session
Ingress RSVP: 2 sessions
To          From          State   Rt  Style Labelin Labelout LSPname
10.50.250.6 10.50.250.4 Up      0  1  SE      -        3 r4-to-r6
10.50.250.7 10.50.250.4 Up      0  1  FF      -        3 r4-to-r7
Total 2 displayed, Up 2, Down 0

[edit]
lab@R6# run show rsvp session
Ingress RSVP: 5 sessions
To          From          State   Rt  Style Labelin Labelout LSPname
10.50.250.1 10.50.250.6 Up      0  1  FF      -        3 r6-to-r1
10.50.250.1 10.50.250.6 Up      0  1  SE      -        299786 to-r1
10.50.250.2 10.50.250.6 Up      0  1  FF      -        3 r6-to-r2
10.50.250.4 10.50.250.6 Up      0  1  SE      -        3 r6-to-r4
10.50.250.7 10.50.250.6 Up      0  1  FF      -        299793 r6-to-r7
Total 5 displayed, Up 5, Down 0

```

- Configure LSPs from R7 to R4 and R6 with the primary path that follows the shortest route through the network. Protect the LSPs with a secondary path that completely avoids the primary path. Make sure that the secondary path is signaled and ready for use. Configure LSPs from R6 and R4 to R7 without restrictions.

a. R7

Unlike the LSPs from the previous Task, here you have to configure protection using standby LSPs signaled. Since, R7 not part of Area 0 CSPF protocol cannot be used and should be disabled.

```

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r4 to 10.50.250.4

[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-r4 primary to-r4-primary

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r4 secondary to-r4-secondary
standby

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r6 to 10.50.250.6

```

```
[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r6 no-cspf

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r6 primary to-r4-primary

[edit]
lab@R7# set protocols mpls label-switched-path r7-to-r6 secondary to-r4-secondary
standby
```

The primary path is forced to pass through router R4, which meets the requirement for the shortest path to router R6.

```
[edit]
lab@R7# set protocols mpls path to-r4-primary 10.50.0.93 strict

[edit]
lab@R7# set protocols mpls path to-r4-secondary
```

b. R4

The LSP in the opposite direction is configured to follow the best IGP path without constraints.

```
[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r7 to 10.50.250.4

[edit]
lab@R4# set protocols mpls label-switched-path r4-to-r7 no-cspf
```

c. R6

```
[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r7 to 10.50.250.7

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r7 no-cspf
```

d. Verify the configuration

The outputs below show the LSPs are established using the shortest path from router R7 to routers R4 and R6. The secondary paths are established over diverse paths and are ready to be used.

```
[edit]
lab@R7# run show mpls lsp extensive
Ingress LSP: 2 sessions

10.50.250.4
  From: 10.50.250.7, State: Up, ActiveRoute: 0, LSPname: r7-to-r4
  ActivePath: to-r4-primary (primary)
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary to-r4-primary State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
      10.50.0.93
      4 Dec 3 09:06:12.667 Selected as active path
      3 Dec 3 09:06:12.666 Record Route: 10.50.0.93
      2 Dec 3 09:06:12.666 Up
      1 Dec 3 09:06:12.652 Originate Call
  Standby to-r4-secondary State: Up
```



```

Priorities: 7 0
SmartOptimizeTimer: 180
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
10.50.0.89 10.50.0.86 10.50.0.33
3 Dec 3 09:06:42.741 Record Route: 10.50.0.89 10.50.0.86 10.50.0.33
2 Dec 3 09:06:42.741 Up
1 Dec 3 09:06:42.510 Originate Call
Created: Sat Dec 3 09:06:13 2016

```

```

10.50.250.6
From: 10.50.250.7, State: Up, ActiveRoute: 0, LSPname: r7-to-r6
ActivePath: to-r4-primary (primary)
LSPTYPE: Static Configured, Penultimate hop popping
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary to-r4-primary State: Up
Priorities: 7 0
SmartOptimizeTimer: 180
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
10.50.0.93 10.50.0.38
4 Dec 3 09:06:12.890 Selected as active path
3 Dec 3 09:06:12.889 Record Route: 10.50.0.93 10.50.0.38
2 Dec 3 09:06:12.889 Up
1 Dec 3 09:06:12.653 Originate Call
Standby to-r4-secondary State: Up
Priorities: 7 0
SmartOptimizeTimer: 180
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
10.50.0.89 10.50.0.86 10.50.0.46
3 Dec 3 09:06:42.249 Record Route: 10.50.0.89 10.50.0.86 10.50.0.46
2 Dec 3 09:06:42.249 Up
1 Dec 3 09:06:42.013 Originate Call
Created: Sat Dec 3 09:06:13 2016
Total 2 displayed, Up 2, Down 0

```

```

[edit]
lab@R6# run show mpls lsp
Ingress LSP: 4 sessions
To          From          State Rt P    ActivePath    LSPname
10.50.250.7 10.50.250.6    Up      0 *          r6-to-r7
Total 4 displayed, Up 4, Down 0

```

```

[edit]
lab@R4# run show mpls lsp
Ingress LSP: 2 sessions
To          From          State Rt P    ActivePath    LSPname
10.50.250.7 10.50.250.4    Up      0 *          r4-to-r7
Total 2 displayed, Up 2, Down 0

```

```

[edit]
lab@R7# run show route 10.50.250.4

inet.0: 572 destinations, 573 routes (571 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.4/32      *[OSPF/10] 01:14:48, metric 300
> to 10.50.0.89 via ge-0/0/4.37

```

```

[edit]
lab@R7# run show route 10.50.250.6

inet.0: 572 destinations, 573 routes (571 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```

```
10.50.250.6/32      *[OSPF/10] 01:14:52, metric 300
> to 10.50.0.89 via ge-0/0/4.37
```

- 4) Make sure that all devices use soft-state reduction for RSVP messages to increase the scalability and reliable message delivery.

The `aggregate` option enables the reduction of RSVP refresh messages exchanged. At the same time `reliable` will ensure that the exchange of messages is done in reliable way. This will instruct the RSVP protocol to include the Message ID object and ensures they are transmit in an ordered manner.

a. All devices part of Area 0

```
[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.12 aggregate

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.12 reliable

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.14 aggregate

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.14 reliable

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.16 aggregate

[edit]
lab@R1# set protocols rsvp interface ge-0/0/4.16 reliable
```

b. Verify the configuration

```
[edit]
lab@R1# run show rsvp interface extensive
RSVP interface: 3 active

ge-0/0/4.12 Index 332, State Ena/Up
  NoAuthentication, Aggregate, Reliable, NoLinkProtection
  HelloInterval 9(second)
  Address 10.50.0.1
  ActiveResv 3, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 700Mbps
  ct0: StaticBW 700Mbps, AvailableBW 600Mbps
  MaxAvailableBW 700Mbps = (bc0*subscription)
  ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7]
0bps
  Protection: Off

ge-0/0/4.14 Index 333, State Ena/Up
  NoAuthentication, Aggregate, Reliable, NoLinkProtection
  HelloInterval 9(second)
  Address 10.50.0.5
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 700Mbps
  ct0: StaticBW 700Mbps, AvailableBW 700Mbps
  MaxAvailableBW 700Mbps = (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 334, State Ena/Up
  NoAuthentication, Aggregate, Reliable, LinkProtection
```

```

HelloInterval 9(second)
Address 10.50.0.9
ActiveResv 2, PreemptionCnt 0, Update threshold 10%
Subscription 100%,
bc0 = ct0, StaticBW 300Mbps
ct0: StaticBW 300Mbps, AvailableBW 300Mbps
  MaxAvailableBW 300Mbps = (bc0*subscription)
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
Protection: On, Bypass: 1, LSP: 1, Protected LSP: 1, Unprotected LSP: 0
  1 Dec  3 11:56:46 New bypass to-r6
Bypass: to-r6, State: Up, Type: LP, LSP: 1, Backup: 0
  4 Dec  3 11:56:46 Record Route:  10.50.0.2 10.50.0.18
  3 Dec  3 11:56:46 Up
  2 Dec  3 11:56:46 CSPF: computation result accepted
  1 Dec  3 11:56:46 Originate Call

```

Solution - Task 4.2: MPLS and LDP configuration

- 5) Assume that R8 supports only LDP protocol. Configure R5 and R6 to provide MPLS transport gateway functionality to the rest of network for R8. CEs attached to R8 should be able to exchange labeled packets with the CEs behind R1 and R2.

The task mentions that router R8 does not support RSVP. In order to enable services to be deployed between router R8 and routers R1 and R2, you have to configure R5 and R6 to transport LDP traffic in RSVP signaled LSP. You have to configure LDP targeted sessions between routers R1, R2, R5 and R6.

a. R1

In JUNOS configuring LDP over RSVP can be done by enabling `ldp-tunneling` for the RSVP LSPs connecting the two routers configured for LDP. Targeted LDP sessions are established between the Loopback IP addresses.

```

[edit]
lab@R1# set protocols ldp interface lo0.0

[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r5 ldp-tunneling

[edit]
lab@R1# set protocols mpls label-switched-path r1-to-r6 ldp-tunneling

```

b. R2

```

[edit]
lab@R2# set protocols ldp interface lo0.0

[edit]
lab@R2# set protocols mpls label-switched-path r2-to-r5 ldp-tunneling

[edit]
lab@R2# set protocols mpls label-switched-path r2-to-r6 ldp-tunneling

```

c. R5

Routers R5, R6 with R8 establish direct LDP sessions over the logical interface between them.

```

[edit]
lab@R5# set protocols ldp interface ge-0/0/4.58

[edit]

```

```
lab@R5# set protocols ldp interface lo0.0

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r1 ldp-tunneling

[edit]
lab@R5# set protocols mpls label-switched-path r5-to-r2 ldp-tunneling
```

d. R6

```
[edit]
lab@R6# set protocols ldp interface ge-0/0/4.68

[edit]
lab@R6# set protocols ldp interface lo0.0

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r1 ldp-tunneling

[edit]
lab@R6# set protocols mpls label-switched-path r6-to-r2 ldp-tunneling
```

e. R8

```
[edit]
lab@R8# set protocols ldp interface ge-0/0/4.58

[edit]
lab@R8# set protocols ldp interface ge-0/0/4.68
```

Enabling LDP on the Loopback interface on R8 is needed to establish targeted sessions to routers R1 and R2.

```
[edit]
lab@R8# set protocols ldp interface lo0.0
```

f. Verify the configuration

The output below show the result of the commands applied.

```
[edit]
lab@R8# run show ldp interface
```

Interface	Label space ID	Nbr count	Next hello
lo0.0	10.50.250.8:0	3	0
ge-0/0/4.58	10.50.250.8:0	1	3
ge-0/0/4.68	10.50.250.8:0	1	2

```
[edit]
lab@R8# run show ldp neighbor
```

Address	Interface	Label space ID	Hold time
10.50.250.1	lo0.0	10.50.250.1:0	38
10.50.250.2	lo0.0	10.50.250.2:0	43
10.50.250.6	lo0.0	10.50.250.6:0	42
10.50.0.49	ge-0/0/4.58	10.50.250.5:0	12
10.50.0.53	ge-0/0/4.68	10.50.250.6:0	14

```
[edit]
lab@R8# run show ldp database
```

Input label database, 10.50.250.8:0--10.50.250.1:0
Labels received: 5

Label	Prefix
3	10.50.250.1/32
299778	10.50.250.5/32
299776	10.50.250.6/32
299783	10.50.250.8/32

```
262147      FEC129 NoCtrlWord ETHERNET 000a012c:0000012c 0a32fa01 0a32fa08
```

Output label database, 10.50.250.8:0--10.50.250.1:0

Labels advertised: 6

Label	Prefix
299779	10.50.250.1/32
299778	10.50.250.2/32
299777	10.50.250.5/32
299780	10.50.250.6/32
3	10.50.250.8/32

```
262145      FEC129 NoCtrlWord ETHERNET 000a012c:0000012c 0a32fa08 0a32fa01
```

Input label database, 10.50.250.8:0--10.50.250.2:0

Labels received: 5

Label	Prefix
3	10.50.250.2/32
299778	10.50.250.5/32
299777	10.50.250.6/32
299781	10.50.250.8/32
299776	L2CKT CtrlWord ETHERNET VC 3

Output label database, 10.50.250.8:0--10.50.250.2:0

Labels advertised: 6

Label	Prefix
299779	10.50.250.1/32
299778	10.50.250.2/32
299777	10.50.250.5/32
299780	10.50.250.6/32
3	10.50.250.8/32
299776	L2CKT CtrlWord ETHERNET VC 3

[edit]

```
lab@R8# run show route table inet.3
```

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.1/32	*[LDP/9] 00:04:04, metric 1 > to 10.50.0.53 via ge-0/0/4.68, Push 299777
10.50.250.2/32	*[LDP/9] 00:04:04, metric 1 > to 10.50.0.53 via ge-0/0/4.68, Push 299778
10.50.250.5/32	*[LDP/9] 00:05:32, metric 1 > to 10.50.0.49 via ge-0/0/4.58
10.50.250.6/32	*[LDP/9] 00:04:04, metric 1 > to 10.50.0.53 via ge-0/0/4.68

[edit]

```
lab@R1# run show ldp interface
```

Interface	Label space ID	Nbr count	Next hello
lo0.0	10.50.250.1:0	2	0

[edit]

```
lab@R1# run show ldp neighbor
```

Address	Interface	Label space ID	Hold time
10.50.250.5	lo0.0	10.50.250.5:0	31
10.50.250.6	lo0.0	10.50.250.6:0	42
10.50.250.8	lo0.0	10.50.250.8:0	44

[edit]

```
lab@R1# run show ldp database
```

Input label database, 10.50.250.1:0--10.50.250.5:0

Labels received: 4

Label	Prefix
299777	10.50.250.1/32
299776	10.50.250.2/32
3	10.50.250.5/32
299781	10.50.250.8/32

```

Output label database, 10.50.250.1:0--10.50.250.5:0
Labels advertised: 4
  Label      Prefix
    3        10.50.250.1/32
 299778      10.50.250.5/32
 299776      10.50.250.6/32
 299783      10.50.250.8/32

[edit]
lab@R1# run show route table inet.3

inet.3: 12 destinations, 18 routes (4 active, 0 holddown, 12 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.2/32      * [RSVP/7/1] 03:16:10, metric 100
> to 10.50.0.2 via ge-0/0/4.12, label-switched-path r1-to-r2
10.50.250.5/32      * [LDP/6] 00:48:08, metric 1
> to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r5
  to 10.50.0.2 via ge-0/0/4.12, label-switched-path r1-to-r5
  [RSVP/7/1] 00:48:08, metric 200
> to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r5
  to 10.50.0.2 via ge-0/0/4.12, label-switched-path to-r6
10.50.250.6/32      * [LDP/6] 03:13:56, metric 1
> to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r6
  [RSVP/7/1] 03:13:56, metric 100
> to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r6
10.50.250.8/32      * [LDP/6] 00:00:10, metric 1
> to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r6

```

6) Make sure that LDP neighbors are not sending traps during state transitions.

The last task is straightforward and simply requires disabling of SNMP traps when LDP state transitions from one state to another.

a. All LDP enabled devices

```

[edit]
lab@R1# set protocols ldp log-updown trap disable

[edit]
lab@R2# set protocols ldp log-updown trap disable

```

50)

Part 5: IPv6 tunneling and Multicast configuration

Solution - Task 5.1: IPv6 tunneling configuration

The next part is focused on enabling IPV6 and Multicast traffic in the network. It is highly recommended to read the whole part before starting to solve the tasks.

- 1) Configure IPv6 connectivity between P1, C1 and U3 peers. Ensure IPv6 is encapsulated in labeled packets within your AS. Make sure that IPv6 routes are exchanged between all three peers. You are not allowed to configure routing-instances of any type or enable native IPv6 transport in the core of your network.

The solution is to enable the network for 6PE, which means to transport the IPv6 traffic inside LSPs. To enable this functionality, you have to configure the BGP protocol to exchange transport MPLS labels. Besides that, you also have to enable the tunneling of the IPv6 in LSPs.

- a. All devices including the route-reflector

```
[edit]
lab@R1# set protocols bgp group internal family inet unicast
```

You have to configure explicit-null label allocation when 6PE is configured. This is required due to the way IPv6 traffic is transported. JUNOS uses the explicit-null label (value 2) to make the egress PE aware that the encapsulated traffic in the labeled packet is IPv6 traffic.

```
[edit]
lab@R1# set protocols bgp group internal family inet6 labeled-unicast explicit-null
```

The two clusters on the router-reflector should reflect IPv6 labeled routes as well.

```
[edit]
lab@route-reflector# set protocols bgp group cluster-1 family inet unicast
```

```
[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-1 family inet6 labeled-unicast explicit-null
```

- b. R4, R6 and R7

Next, you have to configure the `ipv6-tunneling` knob under `protocols mpls`. This will actually copy the routes from the `inet.3` table into the `inet6.3` table and in the same time it will transform them in IPv4 mapped IPv6 addresses.

```
[edit]
lab@R4# set protocols mpls ipv6-tunneling
```

```
[edit]
lab@R6# set protocols mpls ipv6-tunneling
```

```
[edit]
lab@R7# set protocols mpls ipv6-tunneling
```

c. Verify the configuration

The result of the configuration applied for 6PE configuration is shown

R4 receives and resolves the IPv6 routes from external peers P1 and U3.

```
[edit]
lab@R4# run show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 513 510 0 0 0 0
inet6.0 120 120 0 0 0 0
bgp.l3vpn.0 0 0 0 0 0 0
bgp.l2vpn.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.50.250.10 5673.873 1471 422 0 0 3:09:19
Establ
inet.0: 494/494/494/0
inet6.0: 104/104/104/0
bgp.l3vpn.0: 0/0/0/0
bgp.l2vpn.0: 0/0/0/0
170.0.15.10 65457 385 1332 0 0 3:12:27
Establ
inet.0: 16/19/16/0
2001:b170:0:15::be 65457 383 593 0 0 3:12:20
Establ
inet6.0: 16/16/16/0
```

IPv6 routes are advertised to the external peer C1 by router R4.

```
[edit]
lab@R4# run show route advertising-protocol bgp 2001:b170:0:15::be

inet6.0: 162 destinations, 163 routes (162 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
* ::/48 Self 9687 1620 4563
982 324 I
* 2001:240::/32 Self 9687 1620 33112
60177 60637 30709 57763 I
* 2001:250:204::/48 Self 9687 1620 33112
5697 30284 25817 58254 24799 I
* 2001:254::/32 Self 9687 1620 61671
52411 64333 44013 43197 43321 6648 I
* 2001:2b8:17::/48 Self 9687 1620 61671
41164 28389 7809 45703 I
* 2001:2b8:d4::/48 Self 9687 1620 33112
27075 14574 55821 I
* 2001:2d8:1::/48 Self 9687 1620 61671
49129 48750 6221 24807 I
* 2001:2e0::/32 Self 9687 1620 33112
60177 23789 9452 28275 31473 I
* 2001:3e0::/32 Self 9687 1620 33112
60177 5037 32401 1078 9850 39463 42575 I
* 2001:410:102::/48 Self 9687 1620 61671
41164 18823 8392 I
* 2001:418:1::/48 Self 9687 1620 61671
30404 11114 53598 10591 4850 54617 I
```



```
* 2001:420:1100::/41      Self                      9687 1620 33003
63164 51521 34880 I
```

The output below is taken from R7 and shows a route advertised by router R4. The next-hop is resolved in table inet6.3 Two labels are pushed. The Label 2 on the bottom of the lable stack and on top a transport Label, resolved from the LSP.

```
[edit]
lab@R7# run show route table inet6.0 extensive 2001:fe:183:45::/80

inet6.0: 71 destinations, 71 routes (71 active, 0 holddown, 0 hidden)
2001:fe:183:45::/80 (1 entry, 1 announced)
TSI:
KRT in-kernel 2001:fe:183:45::/80 -> {indirect(1048577)}
Page 0 idx 1, (group P1 type External) Type 1 val 0xaf3d000 (adv_entry)
  Advertised metrics:
    Nexthop: Self
    AS path: [5673.873] 65457 I
    Communities: 65457:2021
Path 2001:fe:183:45:: from 10.50.250.10 Vector len 4. Val: 1
  *BGP      Preference: 170/-101
            Next hop type: Indirect, Next hop index: 0
            Address: 0xa6fb2f4
            Next-hop reference count: 32
            Source: 10.50.250.10
            Next hop type: Router, Next hop index: 1048584
            Next hop: 10.50.0.93 via ge-0/0/4.47 weight 0x1, selected
            Label-switched-path r7-to-r4
            Label operation: Push 2
            Label TTL action: prop-ttl
            Load balance label: Label 2: None;
            Label element ptr: 0xaf14200
            Label parent element ptr: 0xaf137e0
            Label element references: 1
            Label element child references: 0
            Label element lsp id: 0
            Session Id: 0x176
            Next hop: 10.50.0.89 via ge-0/0/4.37 weight 0x2001
            Label-switched-path r7-to-r4
            Label operation: Push 2, Push 299782(top)
            Label TTL action: prop-ttl, prop-ttl(top)
            Load balance label: Label 2: None; Label 299782: None;
            Label element ptr: 0xaf139c0
            Label parent element ptr: 0xaf13540
            Label element references: 1
            Label element child references: 0
            Label element lsp id: 0
            Session Id: 0x145
            Protocol next hop: ::ffff:10.50.250.4
            Label operation: Push 2
            Label TTL action: prop-ttl
            Load balance label: Label 2: None;
            Indirect next hop: 0xaca9600 1048577 INH Session ID: 0x17b
            State: <Active Int Ext>
            Local AS: 5673.873 Peer AS: 5673.873
            Age: 1:00:05      Metric2: 300
            Validation State: unverified
            Task: BGP_5673.873.10.50.250.10
            Announcement bits (3): 0-KRT 4-BGP_RT_Background 5-Resolve tree 3
            AS path: 65457 I (Originator)
            Cluster list: 1.1.1.1
            Originator ID: 10.50.250.4
            Communities: 65457:2021
            Accepted
            Route Label: 2
            Localpref: 100
            Router ID: 10.50.250.10
```

```

Indirect next hops: 1
  Protocol next hop: ::ffff:10.50.250.4 Metric: 300
  Label operation: Push 2
  Label TTL action: prop-ttl
  Load balance label: Label 2: None;
  Indirect next hop: 0xaca9600 1048577 INH Session ID:
0x17b
  Indirect path forwarding next hops: 2
    Next hop type: Router
    Next hop: 10.50.0.93 via ge-0/0/4.47 weight 0x1
    Session Id: 0x176
    Next hop: 10.50.0.89 via ge-0/0/4.37 weight
0x2001
      Session Id: 0x145
      ::ffff:10.50.250.4/128 Originating RIB: inet6.3
      Metric: 300 Node path count: 1
      Forwarding nexthops: 2
      Nexthop: 10.50.0.93 via ge-0/0/4.47

```

- 2) Avoid changing the resolution table or enabling MPLS on the interfaces of the route-reflector to resolve any hidden routes on the route-reflector. You are not allowed to have any routes except the Loopback addresses in inet.3 on the RR.

The Task restricts you from enabling MPLS on the route-reflector device. However, the route-reflector needs the routing table inet6.3 populated in order to reflect the routes to the clients. In the next part, the same configuration is needed for the inet.3 routing table. JUNOS provides several ways to populate the two routing tables, but the task directs you to one in particular. The solution is to configure rib-groups to copy the Loopback interface routes from the inet.0 routing table to the inet.3 and inet6.3 table. Use routing policy to filter specific routes with the rib-group.

a. Route-reflector

Rip-groups are common way in JUNOS to copy routes from one routing table to another.

```

[edit]
lab@route-reflector# set routing-options rib-groups leak.to.inet3 import-rib
inet.0

[edit]
lab@route-reflector# set routing-options rib-groups leak.to.inet3 import-rib
inet.3

[edit]
lab@route-reflector# set routing-options rib-groups leak.to.inet3 import-rib
inet6.3

[edit]
lab@route-reflector# set routing-options rib-groups leak.to.inet3 import-policy
leak.to.inet.3.loopback

[edit]
lab@route-reflector# set policy-options policy-statement leak.to.inet.3.loopback
term ipv6 from route-filter ::/0 prefix-length-range /128-/128

[edit]
lab@route-reflector# set policy-options policy-statement leak.to.inet.3.loopback
term ipv6 to rib inet6.3

[edit]
lab@route-reflector# set policy-options policy-statement leak.to.inet.3.loopback
term ipv6 then accept

```

```
[edit]
lab@route-reflector# set policy-options policy-statement leak.to.inet.3.loopback
term ipv4 from route-filter 0.0.0.0/0 prefix-length-range /32-/32

[edit]
lab@route-reflector# set policy-options policy-statement leak.to.inet.3.loopback
term ipv4 to rib inet.3

[edit]
lab@route-reflector# set policy-options policy-statement leak.to.inet.3.loopback
term ipv4 then accept

[edit]
lab@route-reflector# set policy-options policy-statement leak.to.inet.3.loopback
term reject.rest then reject
```

Since all routes that need to be copied are learned by OSPF, it is enough to apply the rib-group to OSPF.

```
[edit]
lab@route-reflector# set protocols ospf rib-group leak.to.inet3
```

b. Verify the configuration

The required routes are present in the inet.3 and inet6.3 table after the changes are committed.

```
[edit]
lab@route-reflector# run show route table inet.3

inet.3: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.1/32      *[OSPF/10] 02:16:51, metric 101
                   > to 10.50.0.97 via ge-0/0/1.49
10.50.250.2/32      *[OSPF/10] 02:16:51, metric 101
                   > to 10.50.0.97 via ge-0/0/1.49
10.50.250.3/32      *[OSPF/10] 02:16:51, metric 101
                   > to 10.50.0.101 via ge-0/0/1.59
10.50.250.4/32      *[OSPF/10] 02:16:51, metric 1
                   > to 10.50.0.97 via ge-0/0/1.49
10.50.250.5/32      *[OSPF/10] 02:16:51, metric 1
                   > to 10.50.0.101 via ge-0/0/1.59
10.50.250.6/32      *[OSPF/10] 02:16:51, metric 101
                   > to 10.50.0.97 via ge-0/0/1.49
                   > to 10.50.0.101 via ge-0/0/1.59
10.50.250.7/32      *[OSPF/10] 02:16:51, metric 201
                   > to 10.50.0.101 via ge-0/0/1.59
10.50.250.8/32      *[OSPF/150] 00:19:01, metric 100, tag 0
                   > to 10.50.0.101 via ge-0/0/1.59

[edit]
lab@route-reflector# run show route table inet6.3

inet6.3: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::ffff:10.50.250.1/128
                   *[OSPF/10] 02:18:08, metric 101
                   > to 10.50.0.97 via ge-0/0/1.49
::ffff:10.50.250.2/128
                   *[OSPF/10] 02:18:08, metric 101
                   > to 10.50.0.97 via ge-0/0/1.49
::ffff:10.50.250.3/128
```

```

        *[OSPF/10] 02:18:08, metric 101
        > to 10.50.0.101 via ge-0/0/1.59
::ffff:10.50.250.4/128
        *[OSPF/10] 02:18:08, metric 1
        > to 10.50.0.97 via ge-0/0/1.49
::ffff:10.50.250.5/128
        *[OSPF/10] 02:18:08, metric 1
        > to 10.50.0.101 via ge-0/0/1.59
::ffff:10.50.250.6/128
        *[OSPF/10] 02:18:08, metric 101
        > to 10.50.0.97 via ge-0/0/1.49
        to 10.50.0.101 via ge-0/0/1.59
::ffff:10.50.250.7/128
        *[OSPF/10] 02:18:08, metric 201
        > to 10.50.0.101 via ge-0/0/1.59
::ffff:10.50.250.8/128
        *[OSPF/150] 00:20:18, metric 100, tag 0
        > to 10.50.0.101 via ge-0/0/1.59
::ffff:192.168.50.50/128
        *[OSPF/10] 02:18:08, metric 101
        > to 10.50.0.97 via ge-0/0/1.49

```

[edit]

lab@route-reflector# **run show route summary**

Autonomous system number: 5673.873

Router ID: 10.50.250.10

```

inet.0: 540 destinations, 744 routes (540 active, 0 holddown, 0 hidden)
  Direct:      4 routes,      4 active
  Local:       3 routes,      3 active
  OSPF:        22 routes,     22 active
  BGP:         715 routes,    511 active

inet.3: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  OSPF:        9 routes,      9 active

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

bgp.l3vpn.0: 39 destinations, 39 routes (39 active, 0 holddown, 0 hidden)
  BGP:         39 routes,     39 active

inet6.0: 123 destinations, 123 routes (19 active, 0 holddown, 104 hidden)
  Direct:      2 routes,      2 active
  OSPF3:       1 routes,      1 active
  BGP:        120 routes,     16 active

inet6.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  OSPF:        1 routes,      1 active

bgp.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
  BGP:         3 routes,      3 active

```

Solution - Task 5.2: Multicast configuration

- 3) Configure all additional interfaces to connect the Src1, Rcv1 and Rcv2 to the network. Make sure those interfaces are reachable by all routers.

First, you have to configure the interfaces connecting to the Source and Receiver peers. After that, you should add them to the OSPF configuration to provide IP reachability through the network.

a. R5

```
[edit]
lab@R5# set interfaces ge-0/0/5 unit 1101 description "=== connection to Src1
=== "

[edit]
lab@R5# set interfaces ge-0/0/5 unit 1101 vlan-id 1101

[edit]
lab@R5# set interfaces ge-0/0/5 unit 1101 family inet address 192.168.40.5/24

[edit]
lab@R5# set protocols ospf area 0.0.0.0 interface ge-0/0/5.1101 passive
```

b. R7

```
[edit]
lab@R7# set interfaces ge-0/0/5 unit 1102 description "=== connection to Rsv1
=== "

[edit]
lab@R7# set interfaces ge-0/0/5 unit 1102 vlan-id 1102

[edit]
lab@R7# set interfaces ge-0/0/5 unit 1102 family inet address 192.168.41.7/24

[edit]
lab@R7# set protocols ospf area 0.0.0.50 interface ge-0/0/5.1102 passive
```

c. R4

```
[edit]
lab@R4# set interfaces ge-0/0/5 unit 1103 description "=== connection to Rsv2
=== "

[edit]
lab@R4# set interfaces ge-0/0/5 unit 1103 vlan-id 1103

[edit]
lab@R4# set interfaces ge-0/0/5 unit 1103 family inet address 192.168.42.4/24

[edit]
lab@R4# set protocols ospf area 0.0.0.0 interface ge-0/0/5.1103 passive
```

d. Verify the configuration

To verify the configuration you can check if the routing information is present in the routing table.

```
[edit]
lab@R7# run show route 192.168.40.100

inet.0: 572 destinations, 573 routes (571 active, 0 holddown, 1 hidden)
```

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

```
192.168.40.0/24    *[OSPF/10] 03:29:15, metric 300
> to 10.50.0.89 via ge-0/0/4.37
```

- 4) Configure PIM on all routers for OSPF Area 0 and Area 50. Make sure no PIM neighbor can be established on the management interface.

- a. All devices apart from R8

Next, you have to configure the PIM protocol connecting the Receivers with the Sources. PIM is actually used to build the multicast distribution tree. Do not forget to disable PIM on the fxp0.0 interface. PIM is enabled by default in sparse mode.

```
[edit]
lab@R1# set protocols pim interface all

[edit]
lab@R1# set protocols pim interface fxp0.0 disable
```

- b. Verify the configuration

After the PIM is enabled on all routers in Area 0 and Area 50 you should examine the PIM neighbor table. The output below shows the PIM neighbors on router R7.

```
[edit]
lab@R7# run show pim neighbors
B = Bidirectional Capable, G = Generation Identifier
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit

Instance: PIM.master
Interface      IP V Mode      Option      Uptime Neighbor addr
ge-0/0/4.37    4 2           HPLGT      03:25:43 10.50.0.89
ge-0/0/4.47    4 2           HPLGT      03:25:43 10.50.0.93
ge-0/0/5.1102  4 2           HPLGT      03:25:43 192.168.41.10
ge-0/0/5.200   4 2           HPLGT      03:25:43 10.50.200.4
```

- 5) Configure IP address 192.168.50.50 as RP for ASM multicast. However, to eliminate the possibility of RP failure dedicate R1 and R2 as RPs in the network. You are not allowed to use different protocols apart from PIM to accomplish this requirement.

The role of the Rendezvous Point (RP) is critical for Any Source Multicast operation. The RP is where the Receivers are connected with the Sources initially to join the Shared Tree. The task requires the configuration of redundant RPs. Two different methods are available in JUNOS to introduce redundant RPs. However, the task description restricts you to use only the PIM protocol. This limits your options to the PIM Anycast RP method for redundant RP configuration. This also suggest that PIM sparse mode has to be used. This is the default mode in JUNOS.

The anycast nature of the method suggest that you have single RP shared by two PIM routers. The same dedicated IP address is advertised from two point in the network. This way any non-PIM router will select the closest RP router. The RP routers also need to exchange information about the sources they have registered.

- a. R1

Because, a second IP is configured on the Loopback interface you have to set the unique IP address as primary to keep it as default address.

```
[edit]
lab@R1# set interfaces lo0 unit 0 family inet address 10.50.250.1/32 primary
```

The anycast IP address is configured on the Loopback0 interface.

```
[edit]
lab@R1# set interfaces lo0 unit 0 family inet address 192.168.50.50/32
```

The local router is configured as RP using the shared address of routers R1 and R2.

```
[edit]
lab@R1# set protocols pim rp local family inet address 192.168.50.50
```

Next, the PIM Anycast PM configuration is added. Using the unique Loopback address as a local anycast-pm address and the remote address of the other RP router. Because anycast-PIM supports IPv6 addresses, you have to configure the correct family.

```
[edit]
lab@R1# set protocols pim rp local family inet anycast-pim rp-set address
10.50.250.2
```

```
[edit]
lab@R1# set protocols pim rp local family inet anycast-pim local-address
10.50.250.1
```

b. R2

Router R2 configuration is similar to the PIM anycast RP configuration on router R1.

```
[edit]
lab@R2# set interfaces lo0 unit 0 family inet address 10.50.250.2/32 primary
```

```
[edit]
lab@R2# set interfaces lo0 unit 0 family inet address 192.168.50.50/32
```

```
[edit]
lab@R2# set protocols pim rp local family inet address 192.168.50.50
```

```
[edit]
lab@R2# set protocols pim rp local family inet anycast-pim rp-set address
10.50.250.1
```

```
[edit]
lab@R2# set protocols pim rp local family inet anycast-pim local-address
10.50.250.2
```

c. R4, R5 and R7

Static RP is configured on the rest of the routers pointing to the shared IP address.

```
[edit]
lab@R4# set protocols pim rp static address 192.168.50.50
```

d. Verify the configuration

```
[edit]
lab@R5# run show pim rps
Instance: PIM.master
```

```
address-family INET
RP address      Type      Mode      Holdtime  Timeout  Groups  Group prefixes
```

```
192.168.50.50 static sparse 0 None 0 224.0.0.0/4
```

```
address-family INET6
```

```
[edit]
```

```
lab@R1# run show pim rps extensive
```

```
Instance: PIM.master
```

```
address-family INET
```

```
RP: 192.168.50.50
```

```
Learned via: static configuration
```

```
Mode: Sparse
```

```
Time Active: 04:26:26
```

```
Holdtime: 150
```

```
Device Index: 0
```

```
Subunit: 0
```

```
Interface: None
```

```
Static RP Override: Off
```

```
Group Ranges:
```

```
224.0.0.0/4
```

```
Anycast PIM rpset:
```

```
10.50.250.2
```

```
Anycast PIM local address used: 10.50.250.1
```

```
[edit]
```

```
lab@R2# run show pim rps extensive
```

```
Instance: PIM.master
```

```
address-family INET
```

```
RP: 192.168.50.50
```

```
Learned via: static configuration
```

```
Mode: Sparse
```

```
Time Active: 03:33:10
```

```
Holdtime: 150
```

```
Device Index: 0
```

```
Subunit: 0
```

```
Interface: None
```

```
Static RP Override: Off
```

```
Group Ranges:
```

```
224.0.0.0/4
```

```
Anycast PIM rpset:
```

```
10.50.250.1
```

```
Anycast PIM local address used: 10.50.250.2
```

```
address-family INET6
```

- 6) Rcv1 is statically joined to group 239.255.255.2. Whereas Rcs2 is configured for source specific multicast, joined to group 239.255.255.3. Make sure multicast transmission from the source to the receivers is possible.

- a. R1, R2, R4, R5 and R7

The task description requires that traffic for the two groups is sourced by Src1. Group 239.255.255.3 is configured for source specific multicast. This requires extending the default SSM range to accept join message for groups that are within the 239/8 subnet. However, since the 239.255.255.2 group is part of the same range you have to enable the devices to accept these joins also.

```
[edit]
```

```
lab@R1# set routing-options multicast ssm-groups 239.0.0.0/8
```

```
[edit]
```



```
lab@R1# set routing-options multicast asm-override-ssm
```

b. Verify the configuration

The VR device is statically joined to groups, simulating Rcv1 and Rcv2. Once traffic is started with the ping command from the VR Src1 routing-instance, traffic reaches the Receivers.

The join messages (G,*) represent the Share Tree from the Rcv. to the RP for the ASM traffic. SSM traffic directly flows over the Source Tree since the Receiver knows the Source.

```
[edit]
lab@R4# run show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.255.255.2
Source: *
RP: 192.168.50.50
Flags: sparse,rptree,wildcard
Upstream interface: ge-0/0/4.24
Upstream neighbor: 10.50.0.13
Upstream state: Join to RP
Uptime: 00:10:31
Downstream neighbors:
  Interface: ge-0/0/4.45
    10.50.0.34 State: Join Flags: SRW Timeout: 163
    Uptime: 00:10:31 Time since last Join: 00:00:46
Number of downstream interfaces: 1
Number of downstream neighbors: 1

Group: 239.255.255.2
Source: 192.168.40.100
Flags: sparse
Upstream interface: ge-0/0/4.45
Upstream neighbor: 10.50.0.34
Upstream state: Prune to RP
Keepalive timeout:
Uptime: 00:10:31
Downstream neighbors:
  Interface: ge-0/0/4.45 (pruned)
    10.50.0.34 State: Prune Flags: SR Timeout: 163
    Uptime: 00:10:31 Time since last Prune: 00:00:46
Number of downstream interfaces: 1
Number of downstream neighbors: 1

Group: 239.255.255.3
Source: 192.168.40.100
Flags: sparse,spt
Upstream interface: ge-0/0/4.45
Upstream neighbor: 10.50.0.34
Upstream state: Join to Source
Keepalive timeout: 351
Uptime: 00:14:35
Downstream neighbors:
  Interface: ge-0/0/5.1103
    192.168.42.10 State: Join Flags: S Timeout: 177
    Uptime: 00:14:35 Time since last Join: 00:00:33
Number of downstream interfaces: 1
Number of downstream neighbors: 1

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

The command below shows the multicast forwarding table with traffic statistics for each group.

```
[edit]
lab@R4# run show multicast route extensive
Instance: master Family: INET

Group: 239.255.255.3
  Source: 192.168.40.100/32
  Upstream interface: ge-0/0/4.45
  Downstream interface list:
    ge-0/0/5.1103
  Number of outgoing interfaces: 1
  Session description: Organisational Local Scope
  Statistics: 0 kbps, 0 pps, 39 packets
  Next-hop ID: 1048580
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0
  Uptime: 00:15:01

Instance: master Family: INET6
```

Below are shown the outputs for the same set of commands on the router R5 directly attached to the Source.

```
[edit]
lab@R5# run show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.255.255.2
  Source: *
  RP: 192.168.50.50
  Flags: sparse,rptree,wildcard
  Upstream interface: ge-0/0/4.45
  Upstream neighbor: 10.50.0.33
  Upstream state: Join to RP
  Uptime: 00:12:19
  Downstream neighbors:
    Interface: ge-0/0/4.35
      10.50.0.85 State: Join Flags: SRW Timeout: 171
      Uptime: 00:12:19 Time since last Join: 00:00:38
  Number of downstream interfaces: 1
  Number of downstream neighbors: 1

Group: 239.255.255.2
  Source: 192.168.40.100
  Flags: sparse,spt
  Upstream interface: ge-0/0/5.1101
  Upstream neighbor: Direct
  Upstream state: Local Source, Prune to RP
  Keepalive timeout: 0
  Uptime: 00:12:59
  Downstream neighbors:
    Interface: ge-0/0/4.35
      10.50.0.85 State: Join Flags: S Timeout: 197
      Uptime: 00:00:12 Time since last Join: 00:00:12
  Number of downstream interfaces: 1
  Number of downstream neighbors: 1

Group: 239.255.255.3
  Source: 192.168.40.100
  Flags: sparse,spt
```

```

Upstream interface: ge-0/0/5.1101
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 00:16:29
Downstream neighbors:
  Interface: ge-0/0/4.45
    10.50.0.33 State: Join Flags: S Timeout: 205
    Uptime: 00:16:29 Time since last Join: 00:00:04
Number of downstream interfaces: 1
Number of downstream neighbors: 1

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

[edit]
lab@R5# run show multicast route extensive
Instance: master Family: INET

Group: 239.255.255.2
Source: 192.168.40.100/32
Upstream interface: ge-0/0/5.1101
Downstream interface list:
  ge-0/0/4.35
Number of outgoing interfaces: 1
Session description: Organisational Local Scope
Statistics: 0 kbps, 0 pps, 24 packets
Next-hop ID: 1048590
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0
Uptime: 00:00:36

Group: 239.255.255.3
Source: 192.168.40.100/32
Upstream interface: ge-0/0/5.1101
Downstream interface list:
  ge-0/0/4.45
Number of outgoing interfaces: 1
Session description: Organisational Local Scope
Statistics: 0 kbps, 0 pps, 5 packets
Next-hop ID: 1048575
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0
Uptime: 02:05:40

Instance: master Family: INET6

[edit]
root@vr-device# run ping 239.255.255.2 routing-instance Src1 bypass-routing ttl
20 interface ge-0/0/8.1101
PING 239.255.255.2 (239.255.255.2): 56 data bytes
^C
--- 239.255.255.2 ping statistics ---
24 packets transmitted, 0 packets received, 100% packet loss

[edit]
root@vr-device# run ping 239.255.255.3 routing-instance Src1 bypass-routing ttl
20 interface ge-0/0/8.1101
PING 239.255.255.3 (239.255.255.3): 56 data bytes
^C
--- 239.255.255.3 ping statistics ---
18 packets transmitted, 0 packets received, 100% packet loss

```

Part 6: MPLS VPNs configuration

Solution - Task 6.1: VPNA configuration

The last part is focused on configuring VPN services. This includes L3VPNs, L2VPNs and VPLS. It is highly recommended to read the whole part before starting to solve the tasks.

- 1) Configure a Layer 3 VPNA connecting the sites: VPNA-CE1, VPNA-CE2, VPNA-CE3 and VPNA-CE4. Sites CE1 and CE2 peer using BGP protocol. Whereas sites CE3 and CE4 are peering with the corresponding PEs using OSPF Area 0. Use your local AS number to define the route-distinguisher and route-target communities. Make sure you receive two prefixes from each site and share them between the sites.

- a. All devices including the route-reflector

First, you have to enable the MP-BGP to support vpnv4 routes.

```
[edit]
lab@R1# set protocols bgp group internal family inet-vpn unicast
```

- b. All devices with L3VPN eBGP neighbors

In the first section, a firewall filter protecting the RE was configured allowing the BGP packets to reach the RE for successful BGP communication. Automatic prefix-list are used to support new BGP neighbors without the need of amending the configuration. To meet the requirements, you have to create a new prefix-list for that supports routing-instances as well.

```
[edit]
lab@R5# set policy-options prefix-list bgp-neighbors-RI apply-path "routing-
instances <*> protocols bgp group <*> neighbor <*>"
```

```
[edit]
lab@R5# set firewall filter protect-re-ipv4 term accept-bgp from source-prefix-
list bgp-neighbors-RI
```

- 2) A backdoor connection is shared between CE1 and CE2. Since this connection is intended for backup link only, make sure that the L3VPN connection is preferred for sending traffic between CE1 and CE2. The backdoor link is configured with MED of 300.

- a. R5

A new interface is configured to connect with CE site.

```
[edit]
lab@R5# set interfaces ge-0/0/5 unit 101 description "=== connection to VPNA-CE1
==="
```

```
[edit]
lab@R5# set interfaces ge-0/0/5 unit 101 vlan-id 101
```

```
[edit]
lab@R5# set interfaces ge-0/0/5 unit 101 family inet address 192.168.15.6/30
```

L3VPNs are configured inside routing-instance (type `vrf`). The vrf instances require the route-distinguisher and route-target policies applied to the routes.

```
[edit]
lab@R5# set routing-instances VPNA instance-type vrf
```

Next, adding the CE interface to the VPNA routing-instance. At least one interface must be configured under routing-instance of type `vrf`.

```
[edit]
lab@R5# set routing-instances VPNA interface ge-0/0/5.101
```

The route-distinguisher should be based on the local AS number. Because it is 4byte AS number, you have to use the format below. A letter 'L' must be appended to the 4byte AS number.

```
[edit]
lab@R5# set routing-instances VPNA route-distinguisher 371786601L:100
```

The same rule is used for configuring the route-target communities. Here is used the `vrf-target` command to set the same communities for import and export to the VPNA.

```
[edit]
lab@R5# set routing-instances VPNA vrf-target target:371786601L:100
```

The `vrf-table-label` is used to program JUNOS using single vrf label for all routes part of the same VPN.

```
[edit]
lab@R5# set routing-instances VPNA vrf-table-label
```

Back door link exists between the sites VPNA-C1 and VPNA-C2 assigned with MED value of 300. The back door must be used only when any or all links to the MPLS network are not available. If you configure the eBGP sessions with both CE sites, the MED value is not send and hence not compared. The next step in the path selection algorithm prefers the internal routes to the same routes received from external peers. The solution for solving the requirement is to configure the independent domain feature using the AS number of the CE sites. This will allow BGP peering with the CE sites using iBGP sessions.

```
[edit]
lab@R5# set routing-instances VPNA routing-options autonomous-system 64999
```

```
[edit]
lab@R5# set routing-instances VPNA routing-options autonomous-system independent-domain
```

The use of independent domain allow you to configure internal BGP session with VPNA-CE1 and VPNA-CE2.

```
[edit]
lab@R5# set routing-instances VPNA protocols bgp group CE1 type internal
```

```
[edit]
lab@R5# set routing-instances VPNA protocols bgp group CE1 local-address 192.168.15.6
```

```
[edit]
lab@R5# set routing-instances VPNA protocols bgp group CE1 neighbor 192.168.15.5
```

b. R8

Similar configuration is applied to router R8 connecting VPNA-CE2 to the network.

```
[edit]
lab@R8# set interfaces ge-0/0/5 unit 102 description "=== connection to VPNA-CE2
=== "

[edit]
lab@R8# set interfaces ge-0/0/5 unit 102 vlan-id 102

[edit]
lab@R8# set interfaces ge-0/0/5 unit 102 family inet address 192.168.15.10/30

[edit]
lab@R8# set routing-instances VPNA instance-type vrf

[edit]
lab@R8# set routing-instances VPNA interface ge-0/0/5.102

[edit]
lab@R8# set routing-instances VPNA route-distinguisher 371786601L:100

[edit]
lab@R8# set routing-instances VPNA vrf-target target:371786601L:100

[edit]
lab@R8# set routing-instances VPNA vrf-table-label

[edit]
lab@R8# set routing-instances VPNA routing-options autonomous-system 64999

[edit]
lab@R8# set routing-instances VPNA routing-options autonomous-system independent-
domain

[edit]
lab@R8# set routing-instances VPNA protocols bgp group CE2 type internal

[edit]
lab@R8# set routing-instances VPNA protocols bgp group CE2 local-address
192.168.15.10

[edit]
lab@R8# set routing-instances VPNA protocols bgp group CE2 neighbor 192.168.15.9
```

c. R1

The other two sites, CE3 and CE4, are using OSPF to exchange routing information with the MPLS network. Apart from that, the rest of the configuration is identical to the sites with BGP routing protocol.

```
[edit]
lab@R1# set interfaces ge-0/0/5 unit 103 description "=== connection to VPNA-CE3
=== "

[edit]
lab@R1# set interfaces ge-0/0/5 unit 103 vlan-id 103

[edit]
lab@R1# set interfaces ge-0/0/5 unit 103 family inet address 192.168.15.14/30

[edit]
lab@R1# set routing-instances VPNA instance-type vrf

[edit]
```

```

lab@R1# set routing-instances VPNA interface ge-0/0/5.103

[edit]
lab@R1# set routing-instances VPNA route-distinguisher 371786601L:100

[edit]
lab@R1# set routing-instances VPNA vrf-export target:371786601L:100

[edit]
lab@R1# set routing-instances VPNA vrf-target target:371786601L:100

[edit]
lab@R1# set routing-instances VPNA vrf-table-label

```

Unlike the BGP sites, OSPF protocol requires an export policy to redistribute the routes from the MP-BGP.

```

[edit]
lab@R1# set routing-instances VPNA protocols ospf export VPNA-export.to.ospf

[edit]
lab@R1# set routing-instances VPNA protocols ospf area 0.0.0.0 interface ge-0/0/5.103

[edit]
lab@R1# set policy-options policy-statement VPNA-export.to.ospf from protocol bgp

[edit]
lab@R1# set policy-options policy-statement VPNA-export.to.ospf then accept

```

d. R2

Similar configuration is applied to R2 as well.

```

[edit]
lab@R2# set interfaces ge-0/0/5 unit 104 description "=== connection to VPNCA-CE4 ==="

[edit]
lab@R2# set interfaces ge-0/0/5 unit 104 vlan-id 104

[edit]
lab@R2# set interfaces ge-0/0/5 unit 104 family inet address 192.168.15.18/30

[edit]
lab@R2# set routing-instances VPNA instance-type vrf

[edit]
lab@R2# set routing-instances VPNA interface ge-0/0/5.104

[edit]
lab@R2# set routing-instances VPNA route-distinguisher 371786601L:100

[edit]
lab@R2# set routing-instances VPNA vrf-export target:371786601L:100

[edit]
lab@R2# set routing-instances VPNA vrf-target target:371786601L:100

[edit]
lab@R2# set routing-instances VPNA vrf-table-label

[edit]
lab@R2# set routing-instances VPNA protocols ospf export VPNA-export.to.ospf

[edit]

```

```
lab@R2# set routing-instances VPNA protocols ospf area 0.0.0.0 interface ge-
0/0/5.104

[edit]
lab@R2# set policy-options policy-statement VPNA-export.to.ospf from protocol bgp

[edit]
lab@R2# set policy-options policy-statement VPNA-export.to.ospf then accept
```

e. Verify the configuration

Router R5 has successfully received two prefixes from VPNA-CE1.

```
[edit]
lab@R5# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0
      842      504      0      0      0      0
inet6.0
      16       0      0      0      0      0
bgp.l3vpn.0
      21      21      0      0      0      0
bgp.l2vpn.0
       0       0      0      0      0      0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.15.5 64999 5 11 0 0 1:26
Establ
  VPNA.inet.0: 2/2/2/0

[edit]
lab@R5# run show route table VPNA.inet.0

VPNA.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.15.4/30 * [Direct/0] 03:04:10
                  > via ge-0/0/5.101
192.168.15.6/32 * [Local/0] 03:04:10
                  Local via ge-0/0/5.101
192.168.15.8/30 * [BGP/170] 00:02:51, localpref 100, from 10.50.250.10
                  AS path: I, validation-state: unverified
                  > to 10.50.0.50 via ge-0/0/4.58, Push 16
192.168.15.12/30 * [BGP/170] 00:02:51, localpref 100, from 10.50.250.10
                  AS path: 5673.873 I, validation-state: unverified
                  > to 10.50.0.46 via ge-0/0/4.56, label-switched-path r5-to-r1
192.168.15.16/30 * [BGP/170] 00:02:51, localpref 100, from 10.50.250.10
                  AS path: 5673.873 I, validation-state: unverified
                  > to 10.50.0.46 via ge-0/0/4.56, label-switched-path r5-to-r2
                  > to 10.50.0.33 via ge-0/0/4.45, label-switched-path r5-to-r2
192.168.16.0/24 * [BGP/170] 00:02:39, localpref 100
                  AS path: I, validation-state: unverified
                  > to 192.168.15.5 via ge-0/0/5.101
192.168.17.0/24 * [BGP/170] 00:02:39, localpref 100
                  AS path: I, validation-state: unverified
                  > to 192.168.15.5 via ge-0/0/5.101
192.168.18.0/24 * [BGP/170] 00:02:51, localpref 100, from 10.50.250.10
                  AS path: I, validation-state: unverified
                  > to 10.50.0.50 via ge-0/0/4.58, Push 16
192.168.19.0/24 * [BGP/170] 00:02:51, localpref 100, from 10.50.250.10
                  AS path: I, validation-state: unverified
                  > to 10.50.0.50 via ge-0/0/4.58, Push 16
192.168.20.0/24 * [BGP/170] 00:02:51, MED 1, localpref 100, from 10.50.250.10
                  AS path: 5673.873 I, validation-state: unverified
                  > to 10.50.0.46 via ge-0/0/4.56, label-switched-path r5-to-r1
192.168.21.0/24 * [BGP/170] 00:02:51, MED 1, localpref 100, from 10.50.250.10
                  AS path: 5673.873 I, validation-state: unverified
```



```

> to 10.50.0.46 via ge-0/0/4.56, label-switched-path r5-to-r1
192.168.22.0/24 * [BGP/170] 00:02:51, MED 1, localpref 100, from 10.50.250.10
AS path: 5673.873 I, validation-state: unverified
> to 10.50.0.46 via ge-0/0/4.56, label-switched-path r5-to-r2
to 10.50.0.33 via ge-0/0/4.45, label-switched-path r5-to-r2
192.168.23.0/24 * [BGP/170] 00:02:51, MED 1, localpref 100, from 10.50.250.10
AS path: 5673.873 I, validation-state: unverified
> to 10.50.0.46 via ge-0/0/4.56, label-switched-path r5-to-r2
to 10.50.0.33 via ge-0/0/4.45, label-switched-path r5-to-r2

```

Router R5 is advertising the routes from the other three sites to the VPNA-CE1. Two prefixes advertised by CE2 received from the MPLS core, are advertised as internal because of the independent-domain configuration.

```

[edit]
lab@R5# run show route advertising-protocol bgp 192.168.15.5

VPNA.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 192.168.15.8/30       Self              100         I
* 192.168.15.12/30      Self              100       5673.873 I
* 192.168.15.16/30      Self              100       5673.873 I
* 192.168.18.0/24       Self              100         I
* 192.168.19.0/24       Self              100         I
* 192.168.20.0/24       Self              1         100       5673.873 I
* 192.168.21.0/24       Self              1         100       5673.873 I
* 192.168.22.0/24       Self              1         100       5673.873 I
* 192.168.23.0/24       Self              1         100       5673.873 I

```

Similarly, router R8 receives two routes from VPNA-CE2 and advertises the routes for the other VPNA sites.

```

[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
                    511         511         0             0         0         0
inet6.0
                    16          0          0             0         0         0
bgp.l3vpn.0
                    30         30          0             0         0         0
bgp.l2vpn.0
                    2           2           0             0         0         0
Peer           AS      InPkt    OutPkt    OutQ    Flaps  Last  Up/Dwn
State|#Active/Received/Accepted/Damped...
192.168.15.9   64999      24       35        0         0          9:40
Establ
  VPNA.inet.0: 2/2/2/0

```

```

[edit]
lab@R8# run show route table VPNA.inet.0

VPNA.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.15.4/30 * [BGP/170] 00:01:54, localpref 100, from 10.50.250.10
AS path: I, validation-state: unverified
> to 10.50.0.49 via ge-0/0/4.58, Push 18
192.168.15.8/30 * [Direct/0] 07:47:26
> via ge-0/0/5.102
192.168.15.10/32 * [Local/0] 07:47:26
Local via ge-0/0/5.102
192.168.15.12/30 * [BGP/170] 00:09:50, localpref 100, from 10.50.250.10
AS path: 5673.873 I, validation-state: unverified
> to 10.50.0.53 via ge-0/0/4.68, Push 16, Push 299795(top)

```

```

192.168.15.16/30 * [BGP/170] 00:09:50, localpref 100, from 10.50.250.10
                  AS path: 5673.873 I, validation-state: unverified
                  > to 10.50.0.53 via ge-0/0/4.68, Push 16, Push 299778 (top)
192.168.16.0/24 * [BGP/170] 00:01:53, localpref 100, from 10.50.250.10
                  AS path: I, validation-state: unverified
                  > to 10.50.0.49 via ge-0/0/4.58, Push 18
192.168.17.0/24 * [BGP/170] 00:01:53, localpref 100, from 10.50.250.10
                  AS path: I, validation-state: unverified
                  > to 10.50.0.49 via ge-0/0/4.58, Push 18
192.168.18.0/24 * [BGP/170] 00:09:47, localpref 100
                  AS path: I, validation-state: unverified
                  > to 192.168.15.9 via ge-0/0/5.102
192.168.19.0/24 * [BGP/170] 00:09:47, localpref 100
                  AS path: I, validation-state: unverified
                  > to 192.168.15.9 via ge-0/0/5.102
192.168.20.0/24 * [BGP/170] 00:09:50, MED 1, localpref 100, from 10.50.250.10
                  AS path: 5673.873 I, validation-state: unverified
                  > to 10.50.0.53 via ge-0/0/4.68, Push 16, Push 299795 (top)
192.168.21.0/24 * [BGP/170] 00:09:50, MED 1, localpref 100, from 10.50.250.10
                  AS path: 5673.873 I, validation-state: unverified
                  > to 10.50.0.53 via ge-0/0/4.68, Push 16, Push 299795 (top)
192.168.22.0/24 * [BGP/170] 00:09:50, MED 1, localpref 100, from 10.50.250.10
                  AS path: 5673.873 I, validation-state: unverified
                  > to 10.50.0.53 via ge-0/0/4.68, Push 16, Push 299778 (top)
192.168.23.0/24 * [BGP/170] 00:09:50, MED 1, localpref 100, from 10.50.250.10
                  AS path: 5673.873 I, validation-state: unverified
                  > to 10.50.0.53 via ge-0/0/4.68, Push 16, Push 299778 (top)

```

[edit]

lab@R8# **run show route advertising-protocol bgp 192.168.15.9**

```

VPNA.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 192.168.15.4/30       Self              100      100        I
* 192.168.15.12/30      Self              100      100        5673.873 I
* 192.168.15.16/30      Self              100      100        5673.873 I
* 192.168.16.0/24       Self              100      100        I
* 192.168.17.0/24       Self              100      100        I
* 192.168.20.0/24       Self              1         100        5673.873 I
* 192.168.21.0/24       Self              1         100        5673.873 I
* 192.168.22.0/24       Self              1         100        5673.873 I
* 192.168.23.0/24       Self              1         100        5673.873 I

```

The routes received from sites CE1 and CE2 are advertised to the OSPF based sites also.

[edit]

lab@R1# **run show route table VPNA**

```

VPNA.inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.15.4/30 * [BGP/170] 00:03:18, localpref 100, from 10.50.250.10
                  AS path: I, validation-state: unverified
                  > to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r5
                  > to 10.50.0.2 via ge-0/0/4.12, label-switched-path r1-to-r5
192.168.15.8/30 * [BGP/170] 00:13:23, localpref 100, from 10.50.250.10
                  AS path: I, validation-state: unverified
                  > to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r6
192.168.15.12/30 * [Direct/0] 03:00:49
                  > via ge-0/0/5.103
192.168.15.14/32 * [Local/0] 03:00:49
                  Local via ge-0/0/5.103
192.168.15.16/30 * [BGP/170] 02:11:40, localpref 100, from 10.50.250.10
                  AS path: I, validation-state: unverified
                  > to 10.50.0.2 via ge-0/0/4.12, label-switched-path r1-to-r2
192.168.16.0/24 * [BGP/170] 00:03:16, localpref 100, from 10.50.250.10
                  AS path: I, validation-state: unverified
                  > to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r5

```

```

to 10.50.0.2 via ge-0/0/4.12, label-switched-path r1-to-r5
192.168.17.0/24 * [BGP/170] 00:03:16, localpref 100, from 10.50.250.10
AS path: I, validation-state: unverified
> to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r5
to 10.50.0.2 via ge-0/0/4.12, label-switched-path r1-to-r5
192.168.18.0/24 * [BGP/170] 00:13:19, localpref 100, from 10.50.250.10
AS path: I, validation-state: unverified
> to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r6
192.168.19.0/24 * [BGP/170] 00:13:19, localpref 100, from 10.50.250.10
AS path: I, validation-state: unverified
> to 10.50.0.10 via ge-0/0/4.16, label-switched-path r1-to-r6
192.168.20.0/24 * [OSPF/10] 03:00:49, metric 1
> to 192.168.15.13 via ge-0/0/5.103
192.168.21.0/24 * [OSPF/10] 03:00:49, metric 1
> to 192.168.15.13 via ge-0/0/5.103
192.168.22.0/24 * [BGP/170] 02:11:40, MED 1, localpref 100, from 10.50.250.10
AS path: I, validation-state: unverified
> to 10.50.0.2 via ge-0/0/4.12, label-switched-path r1-to-r2
192.168.23.0/24 * [BGP/170] 02:11:40, MED 1, localpref 100, from 10.50.250.10
AS path: I, validation-state: unverified
> to 10.50.0.2 via ge-0/0/4.12, label-switched-path r1-to-r2
224.0.0.5/32 * [OSPF/10] 10:00:20, metric 1

```

- 3) The routes received by sites CE3 and CE4 must be marked uniquely per site. All routes received from the MPLS backbone must be advertised to the OSPF sites as external. However, those routes must not be recognizable as coming from AS 5673.873 when examining the CE3 and CE4 routing tables.

a. R1

This task adds a few more requirements to the OSPF sites. The command lines below show the correct solution.

An explicit export policy should be configured to assign the additional extended community “site of origin” to mark the routes per site. Do not forget to add the route-target community for VPNA.

```

[edit]
lab@R1# set policy-options policy-statement VPNA-export.target term 1 from
protocol direct

[edit]
lab@R1# set policy-options policy-statement VPNA-export.target term 1 then
community add VPNA-target

[edit]
lab@R1# set policy-options policy-statement VPNA-export.target term 1 then
community add SoO-CE3

[edit]
lab@R1# set policy-options policy-statement VPNA-export.target term 1 then accept

[edit]
lab@R1# set policy-options policy-statement VPNA-export.target term 2 from
protocol ospf

[edit]
lab@R1# set policy-options policy-statement VPNA-export.target term 2 then
community add VPNA-target

[edit]
lab@R1# set policy-options policy-statement VPNA-export.target term 2 then
community add SoO-CE3

```

```
[edit]
lab@R1# set policy-options policy-statement VPNA-export.target term 2 then accept

[edit]
lab@R1# set policy-options community SoO-CE3 members origin:371786601L:3

[edit]
lab@R1# set policy-options community VPNA-target members target:371786601L:100

[edit]
lab@R1# set routing-instances VPNA vrf-export VPNA-export.target
```

Disabling the OSPF domain-id will redistribute all routes received from MB-BGP via the MPLS core network to the local OSPF process as external, regardless of the communities they carry.

```
[edit]
lab@R1# set routing-instances VPNA protocols ospf domain-id disable
```

The command below will reset the external route tag field from the OSPF routes.

```
[edit]
lab@R1# set routing-instances VPNA protocols ospf domain-vpn-tag 0
```

b. R2

The same configuration is applied to router R2.

```
[edit]
lab@R2# set policy-options policy-statement VPNA-export.target term 1 from
protocol direct

[edit]
lab@R2# set policy-options policy-statement VPNA-export.target term 1 then
community add VPNA-target

[edit]
lab@R2# set policy-options policy-statement VPNA-export.target term 1 then
community add SoO-CE4

[edit]
lab@R2# set policy-options policy-statement VPNA-export.target term 1 then accept

[edit]
lab@R2# set policy-options policy-statement VPNA-export.target term 2 from
protocol ospf

[edit]
lab@R2# set policy-options policy-statement VPNA-export.target term 2 then
community add VPNA-target

[edit]
lab@R2# set policy-options policy-statement VPNA-export.target term 2 then
community add SoO-CE4

[edit]
lab@R2# set policy-options policy-statement VPNA-export.target term 2 then accept

[edit]
lab@R2# set policy-options community SoO-CE4 members origin:371786601L:4

[edit]
lab@R2# set policy-options community VPNA-target members target:371786601L:100

[edit]
```

```
lab@R2# set routing-instances VPNA vrf-export VPNA-export.target
[edit]
lab@R2# set routing-instances VPNA protocols ospf domain-id disable
[edit]
lab@R2# set routing-instances VPNA protocols ospf domain-vpn-tag 0
```

c. Verify the configuration

Router R1 is advertising all routes from the rest of the VPNA sites as external to the local VPNA-CE3 site.

```
[edit]
lab@R1# run show ospf database instance VPNA
```

```
OSPF database, Area 0.0.0.0
Type      ID          Adv Rtr          Seq            Age  Opt  Cksum  Len
Router    *192.168.15.14    192.168.15.14    0x800000018    250  0x22  0x639d  36
Router    192.168.20.1      192.168.20.1      0x800000016    1473 0x22  0x18df  60
Network   192.168.15.13     192.168.20.1      0x800000004    1473 0x22  0x1120  32
OSPF AS SCOPE link state database
Type      ID          Adv Rtr          Seq            Age  Opt  Cksum  Len
Extern    *192.168.15.4     192.168.15.14    0x800000003    250  0x22  0xa10f  36
Extern    *192.168.15.8     192.168.15.14    0x800000007    250  0x22  0x7137  36
Extern    *192.168.15.16    192.168.15.14    0x800000006    250  0x22  0x237e  36
Extern    *192.168.16.0     192.168.15.14    0x800000003    250  0x22  0xd0df  36
Extern    *192.168.17.0     192.168.15.14    0x800000003    250  0x22  0xc5e9  36
Extern    *192.168.18.0     192.168.15.14    0x800000003    250  0x22  0xbaf3  36
Extern    *192.168.19.0     192.168.15.14    0x800000003    250  0x22  0xaaffd 36
Extern    *192.168.22.0     192.168.15.14    0x800000001    250  0x22  0x9c0f  36
Extern    *192.168.23.0     192.168.15.14    0x800000001    250  0x22  0x9119  36
```

```
[edit]
lab@R2# run show ospf database instance VPNA
```

```
OSPF database, Area 0.0.0.0
Type      ID          Adv Rtr          Seq            Age  Opt  Cksum  Len
Router    *192.168.15.18    192.168.15.18    0x800000013    287  0x22  0x975d  36
Router    192.168.22.1      192.168.22.1      0x800000012    2600 0x22  0xcc20  60
Network   *192.168.15.18    192.168.15.18    0x800000006    1144 0x22  0xb95e  32
OSPF AS SCOPE link state database
Type      ID          Adv Rtr          Seq            Age  Opt  Cksum  Len
Extern    *192.168.15.4     192.168.15.18    0x800000003    287  0x22  0x8923  36
Extern    *192.168.15.8     192.168.15.18    0x800000006    287  0x22  0x5b4a  36
Extern    *192.168.15.12    192.168.15.18    0x800000006    287  0x22  0x336e  36
Extern    *192.168.16.0     192.168.15.18    0x800000003    287  0x22  0xb8f3  36
Extern    *192.168.17.0     192.168.15.18    0x800000003    287  0x22  0xadfd  36
Extern    *192.168.18.0     192.168.15.18    0x800000003    287  0x22  0xa208  36
Extern    *192.168.19.0     192.168.15.18    0x800000003    287  0x22  0x9712  36
Extern    *192.168.20.0     192.168.15.18    0x800000001    287  0x22  0x9a0f  36
Extern    *192.168.21.0     192.168.15.18    0x800000001    287  0x22  0x8f19  36
```

Verify on the VR device if the routes are received with a 0 tag.

```
[edit]
root@vr-device# run show route table VPNA-CE3
```

```
VPNA-CE3.inet.0: 16 destinations, 18 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
192.168.15.4/30    *[OSPF/150] 00:00:03, metric 0, tag 0
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.15.8/30    *[OSPF/150] 00:00:03, metric 0, tag 0
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.15.12/30   *[Direct/0] 03:09:21
```

```

> via ge-0/0/8.103
192.168.15.13/32 * [Local/0] 03:09:21
                  Local via ge-0/0/8.103
192.168.15.16/30 * [OSPF/150] 00:00:03, metric 0, tag 0
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.16.0/24 * [OSPF/150] 00:00:03, metric 0, tag 0
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.17.0/24 * [OSPF/150] 00:00:03, metric 0, tag 0
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.18.0/24 * [OSPF/150] 00:00:03, metric 0, tag 0
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.19.0/24 * [OSPF/150] 00:00:03, metric 0, tag 0
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.20.0/24 * [Direct/0] 10:05:00
                  > via lo0.10
                  [OSPF/10] 10:05:00, metric 0
                  > via lo0.10
192.168.20.1/32 * [Local/0] 10:05:00
                  Local via lo0.10
192.168.21.0/24 * [Direct/0] 10:05:00
                  > via lo0.10
                  [OSPF/10] 10:05:00, metric 0
                  > via lo0.10
192.168.21.1/32 * [Local/0] 10:05:00
                  Local via lo0.10
192.168.22.0/24 * [OSPF/150] 00:00:03, metric 1, tag 0
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.23.0/24 * [OSPF/150] 00:00:03, metric 1, tag 0
                  > to 192.168.15.14 via ge-0/0/8.103
224.0.0.5/32 * [OSPF/10] 10:05:05, metric 1
                MultiRecv

```

For comparison the output below is taken again from VR device but without the command `domain-vpn-tag 0` applied.

```

[edit]
root@vr-device# run show route table VPNA-CE3

VPNA-CE3.inet.0: 16 destinations, 18 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.15.4/30 * [OSPF/150] 00:05:08, metric 0, tag 3489661801
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.15.8/30 * [OSPF/150] 00:05:08, metric 0, tag 3489661801
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.15.12/30 * [Direct/0] 03:07:48
                  > via ge-0/0/8.103
192.168.15.13/32 * [Local/0] 03:07:48
                  Local via ge-0/0/8.103
192.168.15.16/30 * [OSPF/150] 00:05:08, metric 0, tag 3489661801
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.16.0/24 * [OSPF/150] 00:05:08, metric 0, tag 3489661801
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.17.0/24 * [OSPF/150] 00:05:08, metric 0, tag 3489661801
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.18.0/24 * [OSPF/150] 00:05:08, metric 0, tag 3489661801
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.19.0/24 * [OSPF/150] 00:05:08, metric 0, tag 3489661801
                  > to 192.168.15.14 via ge-0/0/8.103
192.168.20.0/24 * [Direct/0] 10:03:27
                  > via lo0.10
                  [OSPF/10] 10:03:27, metric 0
                  > via lo0.10
192.168.20.1/32 * [Local/0] 10:03:27
                  Local via lo0.10
192.168.21.0/24 * [Direct/0] 10:03:27
                  > via lo0.10

```

```

[edit]
lab@R1# run show route table VPNA 192.168.22.0/24 extensive

VPNA.inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
192.168.22.0/24 (1 entry, 1 announced)
TSI:
OSPF area : 0.0.0.0, LSA ID : 192.168.22.0, LSA type : Summary
KRT in-kernel 192.168.22.0/24 -> {indirect(1048597)}
    *BGP      Preference: 170/-101
              Route Distinguisher: 5673.873:100
              Next hop type: Indirect, Next hop index: 0
              Address: 0xa6ff0f8
              Next-hop reference count: 9
              Source: 10.50.250.10
              Next hop type: Router, Next hop index: 694
              Next hop: 10.50.0.2 via ge-0/0/4.12, selected
              Label-switched-path rl-to-r2
              Label operation: Push 16
              Label TTL action: prop-ttl
              Load balance label: Label 16: None;
              Label element ptr: 0xb0afd20
              Label parent element ptr: 0xb0ac780
              Label element references: 1
              Label element child references: 0
              Label element lsp id: 0
              Session Id: 0x163
              Protocol next hop: 10.50.250.2
              Label operation: Push 16
              Label TTL action: prop-ttl
              Load balance label: Label 16: None;
              Indirect next hop: 0xa8fbd20 1048597 INH Session ID: 0x183
              State: <Secondary Active Int Ext ProtectionCand>
              Local AS: 5673.873 Peer AS: 5673.873
              Age: 2:25:04      Metric: 1      Metric2: 100
              Validation State: unverified
              Task: BGP_5673.873.10.50.250.10
              Announcement bits (2): 0-VPNA-OSPF 2-KRT
              AS path: I (Originator)
              Cluster list: 1.1.1.1
              Originator ID: 10.50.250.2
              Communities: target:5673.873:100 origin:5673.873:4 rte-
type:0.0.0.0:1:0
              Import Accepted
              VPN Label: 16
              Localpref: 100
              Router ID: 10.50.250.10
              Primary Routing Table bgp.13vpn.0
              Indirect next hops: 1
                Protocol next hop: 10.50.250.2 Metric: 100
                Label operation: Push 16
                Label TTL action: prop-ttl
                Load balance label: Label 16: None;
                Indirect next hop: 0xa8fbd20 1048597 INH Session ID:
0x183
                Indirect path forwarding next hops: 1
                  Next hop type: Router
                  Next hop: 10.50.0.2 via ge-0/0/4.12
                  Session Id: 0x163

```

```

10.50.250.2/32 Originating RIB: inet.3
Metric: 100 Node path count: 1
Forwarding nexthops: 1
Nexthop: 10.50.0.2 via ge-0/0/4.12

```

Solution - Task 6.2: VPNB configuration

- 4) Configure a Layer 3 VPNB connecting sites: VPNB-CE1, VPNB-CE2 and VPNB-CE3. Site CE3 is connected through a 3rd party ISP1 that is connected with your AS at R1. Use BGP as PE-CE protocol for both sites (use AS 65100). Make sure you use your local AS number to configuring the route-target communities and the loopback for the route-distinguisher. Ensure you receive prefixes from each site.

a. R5

VPNB is connecting three different sites.

```

[edit]
lab@R5# set interfaces ge-0/0/5 unit 110 description "=== connection to VPNB-CE1
=== "

```

```

[edit]
lab@R5# set interfaces ge-0/0/5 unit 110 vlan-id 110

```

```

[edit]
lab@R5# set interfaces ge-0/0/5 unit 110 family inet address 172.17.0.1/30

```

The VPNB is configured in a new routing instance of type vrf.

```

[edit]
lab@R5# set routing-instances VPNB instance-type vrf

[edit]
lab@R5# set routing-instances VPNB interface ge-0/0/5.110

```

The task requires the route-distinguisher to be based on the Loopback IP address.

```

[edit]
lab@R5# set routing-instances VPNB route-distinguisher 10.50.250.5:200

```

The vrf import and export policy are assigned to the VPNB routing instance.

```

[edit]
lab@R5# set routing-instances VPNB vrf-import VPNB-import.target

[edit]
lab@R5# set routing-instances VPNB vrf-export VPNB-export.target

[edit]
lab@R5# set routing-instances VPNB vrf-table-label

[edit]
lab@R5# set routing-instances VPNB protocols bgp group VPNB-CE1 peer-as 65100

```

The as-override is needed to replace the AS 65001 with MPLS core AS number.

```

[edit]
lab@R5# set routing-instances VPNB protocols bgp group VPNB-CE1 as-override

[edit]

```



```
lab@R5# set routing-instances VPNB protocols bgp group VPNB-CE1 neighbor
172.17.0.2
```

Below the vrf import and export policy are configured used for controlling the routes sent or received in VPNB.

```
[edit]
lab@R5# set policy-options policy-statement VPNB-export.target term accept.rest
then community add VPNB-target

[edit]
lab@R5# set policy-options policy-statement VPNB-export.target term accept.rest
then accept

[edit]
lab@R5# set policy-options policy-statement VPNB-import.target term 1 from
protocol bgp

[edit]
lab@R5# set policy-options policy-statement VPNB-import.target term 1 from
community VPNB-target

[edit]
lab@R5# set policy-options policy-statement VPNB-import.target term 1 then accept
```

Communities have to be defined under the `policy-option` stanza.

```
[edit]
lab@R5# set policy-options community VPNB-target members target:371786601L:200
```

b. R8

Similar configuration is applied to router R8 for connection site VPNB-CE2.

```
[edit]
lab@R8# set interfaces ge-0/0/5 unit 111 description "=== connection to VPNB-CE2
=== "

[edit]
lab@R8# set interfaces ge-0/0/5 unit 111 vlan-id 111

[edit]
lab@R8# set interfaces ge-0/0/5 unit 111 family inet address 172.17.0.5/30

[edit]
lab@R8# set routing-instances VPNB instance-type vrf

[edit]
lab@R8# set routing-instances VPNB interface ge-0/0/5.111

[edit]
lab@R8# set routing-instances VPNB route-distinguisher 10.50.250.8:200

[edit]
lab@R8# set routing-instances VPNB vrf-import VPNB-import.target

[edit]
lab@R8# set routing-instances VPNB vrf-export VPNB-export.target

[edit]
lab@R8# set routing-instances VPNB vrf-table-label

[edit]
lab@R8# set routing-instances VPNB protocols bgp group VPNB-CE2 peer-as 65100

[edit]
```

```

lab@R8# set routing-instances VPNB protocols bgp group VPNB-CE2 as-override

[edit]
lab@R8# set routing-instances VPNB protocols bgp group VPNB-CE2 neighbor
172.17.0.6

[edit]
lab@R8# set policy-options policy-statement VPNB-export.target term accept.rest
then community add VPNB-target

[edit]
lab@R8# set policy-options policy-statement VPNB-export.target term accept.rest
then accept

[edit]
lab@R8# set policy-options policy-statement VPNB-import.target term 1 from
protocol bgp

[edit]
lab@R8# set policy-options policy-statement VPNB-import.target term 1 from
community VPNB-target

[edit]
lab@R8# set policy-options policy-statement VPNB-import.target term 1 then accept

[edit]
lab@R8# set policy-options community VPNB-target members target:371786601L:200

```

c. Verify the configuration

```

[edit]
lab@R5# run show bgp summary
Groups: 4 Peers: 4 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0
      842      504      0      0      0      0
inet6.0
      16       0      0      0      0      0
bgp.l3vpn.0
      21      21      0      0      0      0
bgp.l2vpn.0
       0       0      0      0      0      0
Peer      AS      InPkt      OutPkt      OutQ      Flaps      Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.17.0.2 65100      461      544      0      1      3:57:05
Establ
  VPNB.inet.0: 6/6/6/0

```

```

[edit]
lab@R5# run show route receive-protocol bgp 172.17.0.2

inet.0: 563 destinations, 907 routes (559 active, 0 holddown, 4 hidden)
inet.1: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
inet.3: 5 destinations, 9 routes (3 active, 0 holddown, 4 hidden)
VPNA.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
VPNB.inet.0: 491 destinations, 744 routes (491 active, 0 holddown, 0 hidden)

```

Prefix	Nexthop	MED	Lclpref	AS_path
* 172.17.1.0/24	172.17.0.2			65100 I
* 172.17.2.0/24	172.17.0.2			65100 I
* 172.17.3.0/24	172.17.0.2			65100 I
* 172.17.4.0/24	172.17.0.2			65100 I
* 172.17.5.0/24	172.17.0.2			65100 I
* 172.17.6.0/24	172.17.0.2			65100 I

```
[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.0 511 511 0 0 0 0
inet6.0 16 0 0 0 0 0
bgp.l3vpn.0 30 30 0 0 0 0
bgp.l2vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
172.17.0.6 65100 448 480 0 1 3:26:03
Establ
VPNB.inet.0: 5/5/5/0
```

```
[edit]
lab@R8# run show route receive-protocol bgp 172.17.0.6

inet.0: 529 destinations, 529 routes (529 active, 0 holddown, 0 hidden)

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

VPNA.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)

VPNB.inet.0: 28 destinations, 28 routes (28 active, 0 holddown, 0 hidden)
Prefix Nexthop MED Lclpref AS path
* 172.17.21.0/24 172.17.0.6 65100 I
* 172.17.22.0/24 172.17.0.6 65100 I
* 172.17.23.0/24 172.17.0.6 65100 I
* 172.17.24.0/24 172.17.0.6 65100 I
* 172.17.25.0/24 172.17.0.6
```

5) You have to provide connectivity to and from AS 87.109 for site CE.. Make sure you do not advertise additional routes to the other VPN sites.

a. R5

```
[edit]
lab@R5# set policy-options policy-statement VPNB-export.target term
filter.AS87.109 from protocol bgp
```

```
[edit]
lab@R5# set policy-options policy-statement VPNB-export.target term
filter.AS87.109 from as-path 87.109.origin
```

```
[edit]
lab@R5# set policy-options policy-statement VPNB-export.target term
filter.AS87.109 then reject
```

The configuration lines below are explicitly advertising the routes received from AS 87.109 to CE1.

```
[edit]
lab@R5# set policy-options policy-statement to.VPNB-CE1 term accept.AS87.109 from
protocol bgp
```

```
[edit]
lab@R5# set policy-options policy-statement to.VPNB-CE1 term accept.AS87.109 from
as-path 87.109.origin
```

```
[edit]
lab@R5# set policy-options policy-statement to.VPNB-CE1 term accept.AS87.109 then
accept
```

```
[edit]
lab@R5# set policy-options policy-statement AS87.109.origin term 1 from as-path
87.109.origin
```

```
[edit]
lab@R5# set policy-options policy-statement AS87.109.origin term 1 then accept
```

```
[edit]
lab@R5# set policy-options policy-statement AS87.109.origin term 2 then reject
```

```
[edit]
lab@R5# set policy-options as-path 87.109.origin "87\109 ."
```

The `rib-groups` is a method to copy routes between different routing tables. Using routing policy together with `rib-groups` provide greater flexibility and control.

Copy the direct interface routes is needed for successful resolving the routes copied between the routing tables.

```
[edit]
lab@R5# set routing-options rib-groups local-interfaces import-rib inet.0
```

```
[edit]
lab@R5# set routing-options rib-groups local-interfaces import-rib VPNB.inet.0
```

Unlike with the interface routes `rib-group`, the order of the different RIBs is important when copying routes learned by routing protocols.

The `rib-group` *bgp.to.VPNB* is intended to copy the routes from the master `inet.0` routing table to the VPNB routing instance.

```
[edit]
lab@R5# set routing-options rib-groups bgp.to.VPNB import-rib inet.0
```

```
[edit]
lab@R5# set routing-options rib-groups bgp.to.VPNB import-rib VPNB.inet.0
```

Applying policy to the `rib-group` allows for further filtering of the copied routes.

```
[edit]
lab@R5# set routing-options rib-groups bgp.to.VPNB import-policy AS87.109.origin
```

Another `rib-group` is needed to copy the routes from the VPNB routing table to the master `inet.0` routing table. Basically, the order of the RIBs is the opposite of the `rib-group` *bgp.to.VPNB*.

```
[edit]
lab@R5# set routing-options rib-groups VPNB.to.bgp import-rib VPNB.inet.0
```

```
[edit]
lab@R5# set routing-options rib-groups VPNB.to.bgp import-rib inet.0
```

The `rib-groups` only work if they are applied to the interface routes or a particular routing protocol.

```
[edit]
lab@R5# set routing-options interface-routes rib-group inet local-interfaces
```

The rib-group *bgp.to.VPNB* is associated to the internal BGP group. You might make the mistake here and apply the rib-group to the external BGP group to peer U2. Although that will accomplish what is asked, you have to take into account that the connection to U2 could fail.

```
[edit]
lab@R5# set protocols bgp group internal family inet unicast rib-group
bgp.to.VPNB
```

```
[edit]
lab@R5# set routing-instances VPNB routing-options interface-routes rib-group
inet local-interfaces
```

```
[edit]
lab@R5# set routing-instances VPNB protocols bgp group VPNB-CE1 family inet
unicast rib-group VPNB.to.bgp
```

b. Verify the configuration

After the commit, all the routes received by AS 87.109 are copied to the VPNB table.

```
[edit]
lab@R5# run show route table VPNB.inet.0
```

```
VPNB.inet.0: 492 destinations, 745 routes (491 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.64.0.0/10      *[BGP/170] 02:27:24, MED 50, localpref 100, from 10.50.250.10
                 AS path: 87.109 1620 61671 I, validation-state: unverified
                 > to 10.50.0.85 via ge-0/0/4.35
                 [BGP/170] 03:44:01, MED 150, localpref 100
                 AS path: 87.109 1620 61671 I, validation-state: unverified
                 > to 15.16.18.85 via ge-0/0/5.2003
1.84.160.0/20    *[BGP/170] 03:44:01, MED 60, localpref 100
                 AS path: 87.109 1620 33112 I, validation-state: unverified
                 > to 15.16.18.85 via ge-0/0/5.2003
1.96.0.0/11      *[BGP/170] 02:27:24, MED 50, localpref 100, from 10.50.250.10
                 AS path: 87.109 1620 33112 63164 40776 51777 I, validation-
state: unverified
                 > to 10.50.0.85 via ge-0/0/4.35
                 [BGP/170] 03:44:01, MED 150, localpref 100
                 AS path: 87.109 1620 33112 63164 40776 51777 I, validation-
state: unverified
                 > to 15.16.18.85 via ge-0/0/5.2003
1.161.192.0/21   *[BGP/170] 03:44:01, MED 60, localpref 100
                 AS path: 87.109 1620 33112 30404 32138 45045 I, validation-
state: unverified
                 > to 15.16.18.85 via ge-0/0/5.2003
1.176.0.0/12     *[BGP/170] 02:27:24, MED 50, localpref 100, from 10.50.250.10
                 AS path: 87.109 1620 33112 49129 16320 52954 I, validation-
state: unverified
                 > to 10.50.0.85 via ge-0/0/4.35
                 [BGP/170] 03:44:01, MED 150, localpref 100
                 AS path: 87.109 1620 33112 49129 16320 52954 I, validation-
state: unverified
                 > to 15.16.18.85 via ge-0/0/5.2003
```

Conversely, the routes received from VPNB-CE1 are copied to master routing table.

```
[edit]
lab@R5# run show route receive-protocol bgp 172.17.0.2
```

```
inet.0: 563 destinations, 907 routes (559 active, 0 holddown, 4 hidden)
  Prefix                Nexthop              MED      Lclpref    AS path
* 172.17.1.0/24         172.17.0.2          65100 I
* 172.17.2.0/24         172.17.0.2          65100 I
* 172.17.3.0/24         172.17.0.2          65100 I
* 172.17.4.0/24         172.17.0.2          65100 I
* 172.17.5.0/24         172.17.0.2          65100 I
* 172.17.6.0/24         172.17.0.2          65100 I

inet.1: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

inet.3: 5 destinations, 9 routes (3 active, 0 holddown, 4 hidden)

VPNA.inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)

VPNB.inet.0: 491 destinations, 744 routes (491 active, 0 holddown, 0 hidden)
  Prefix                Nexthop              MED      Lclpref    AS path
* 172.17.1.0/24         172.17.0.2          65100 I
* 172.17.2.0/24         172.17.0.2          65100 I
* 172.17.3.0/24         172.17.0.2          65100 I
* 172.17.4.0/24         172.17.0.2          65100 I
* 172.17.5.0/24         172.17.0.2          65100 I
* 172.17.6.0/24         172.17.0.2          65100 I
```

The AS 87.109 routes are not sent to router R8, hence they are not advertised to VPNB-CE2.

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.17.0.6

VPNB.inet.0: 28 destinations, 28 routes (28 active, 0 holddown, 0 hidden)
  Prefix                Nexthop              MED      Lclpref    AS path
* 10.50.0.32/30         Self                  I
* 10.50.0.44/30         Self                  I
* 10.50.0.48/30         Self                  I
* 10.50.0.84/30         Self                  I
* 10.50.0.100/30        Self                  I
* 10.50.250.5/32        Self                  I
* 15.16.18.84/30        Self                  I
* 172.16.5.4/30         Self                  4356 I
* 172.17.0.0/30         Self                  I
* 172.17.1.0/24         Self                  5673.873 I
* 172.17.2.0/24         Self                  5673.873 I
* 172.17.3.0/24         Self                  5673.873 I
* 172.17.4.0/24         Self                  5673.873 I
* 172.17.5.0/24         Self                  5673.873 I
* 172.17.6.0/24         Self                  5673.873 I
```

- 6) Since site CE3 is connected to 3rd party provider ISP1, you must extend the VPNB using an inter-provider method that will provide you with a way to filter routes received from ISP-1 based on target-communities. Do not use routing policy on R1 for providing connectivity to CE3. The remote PE is using route-target community *target:4356:500*.

Three options can be used to extend a L3VPN via other provider network.

- Option A – within each AS, the routes are announced by the MP-BGP protocol used for all L3VPNs. On the AS border routers a VRF for each Inter-AS L3VPN is configured and a pure IP connection is established between the ASBRs. This is the least scalable solution and does not fit in

the requirements of the Task. Because, the routes exchanged between the ASBRs are not vpnv4 routes, they do not carry route-target communities.

- Option B – All routes exchanged should be stored locally on ASBR. Since, the exchanged routes are vpnv4 carrying VPN labels with route-target communities, there is a way to further filter or manipulate the VPN routes.
- Option C – provides the most scalable method, using Labeled BGP sessions to create an MPLS path end-to-end. On top of that, a multihop external MP-BGP session exchanges the vpnv4 routes between the ASs. Since, the ASBRs do not have visibility of the overlay MP-BGP routes, there is not way to control the exchange of the routes.

From the three methods the Option B is one that should be used to accomplish the Task.

a. R1

First, the interface between the two networks must be configured. Inter-AS option B exchanges labeled routes between the ASBRs, this means the interconnect interfaces must be also configured with `family mpls`.

```
[edit]
lab@R1# set interfaces ge-0/0/5 unit 112 description "=== connection to ISP1 ==="

[edit]
lab@R1# set interfaces ge-0/0/5 unit 112 vlan-id 112

[edit]
lab@R1# set interfaces ge-0/0/5 unit 112 family inet address 172.17.0.9/30

[edit]
lab@R1# set interfaces ge-0/0/5 unit 112 family mpls
```

Next the external MP-BGP session is configured between router R1 and the ISP1 peer.

```
[edit]
lab@R1# set protocols bgp group ISP1 log-updown

[edit]
lab@R1# set protocols bgp group ISP1 family inet-vpn unicast

[edit]
lab@R1# set protocols bgp group ISP1 peer-as 4356

[edit]
lab@R1# set protocols bgp group ISP1 neighbor 172.17.0.10
```

b. R5

Because using a routing policy on router R1 to manipulate the routes exchanged between the ASBRs is not allowed, you can accept the remote route-target community on router R5 and R8. This will provide successful communication between the local and remote sites of VPNB.

```
[edit]
lab@R5# set policy-options policy-statement VPNB-export.target term accept.rest
then community add VPNB-CE3

[edit]
lab@R5# set policy-options policy-statement VPNB-import.target term 1 from
community VPNB-CE3
```

```
[edit]
lab@R5# set policy-options community VPNB-CE3 members target:4356:500
```

c. R8

```
[edit]
lab@R8# set policy-options policy-statement VPNB-export.target term accept.rest
then community add VPNB-CE3
```

```
[edit]
lab@R8# set policy-options policy-statement VPNB-import.target term 1 from
community VPNB-CE3
```

```
[edit]
lab@R8# set policy-options community VPNB-CE3 members target:4356:500
```

d. Verify the configuration

Router R1 receives vpnv4 routes from the ISP1 peer.

```
[edit]
lab@R1# run show route table bgp.l3vpn.0

bgp.l3vpn.0: 39 destinations, 39 routes (39 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

4356:500:172.16.5.4/30
    *[BGP/170] 04:41:35, localpref 100
        AS path: 4356 I, validation-state: unverified
        > to 172.17.0.10 via ge-0/0/5.112, Push 299786
4356:500:172.17.31.0/24
    *[BGP/170] 04:41:35, localpref 100
        AS path: 4356 65100 I, validation-state: unverified
        > to 172.17.0.10 via ge-0/0/5.112, Push 299786
4356:500:172.17.32.0/24
    *[BGP/170] 04:41:35, localpref 100
        AS path: 4356 65100 I, validation-state: unverified
        > to 172.17.0.10 via ge-0/0/5.112, Push 299786
4356:500:172.17.33.0/24
    *[BGP/170] 04:41:35, localpref 100
        AS path: 4356 65100 I, validation-state: unverified
        > to 172.17.0.10 via ge-0/0/5.112, Push 299786
4356:500:172.17.34.0/24
    *[BGP/170] 04:41:35, localpref 100
        AS path: 4356 65100 I, validation-state: unverified
        > to 172.17.0.10 via ge-0/0/5.112, Push 299786
4356:500:172.17.35.0/24
    *[BGP/170] 04:41:35, localpref 100
        AS path: 4356 65100 I, validation-state: unverified
        > to 172.17.0.10 via ge-0/0/5.112, Push 299786
```

Using the remote route-target community in the import vrf policy, router R5 imports the remote routes into the vrf routing table.

```
[edit]
lab@R5# run show route 172.17.31.0/24

VPNB.inet.0: 492 destinations, 745 routes (492 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.17.31.0/24    *[BGP/170] 02:48:55, localpref 100, from 10.50.250.10
                  AS path: 4356 65100 I, validation-state: unverified
                  > to 10.50.0.46 via ge-0/0/4.56, label-switched-path r5-to-r1
```

```
bgp.l3vpn.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
```


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

```
4356:500:172.17.31.0/24
      * [BGP/170] 02:48:55, localpref 100, from 10.50.250.10
        AS path: 4356 65100 I, validation-state: unverified
        > to 10.50.0.46 via ge-0/0/4.56, label-switched-path r5-to-r1
```

Similarly, router R8 successfully imports the routes into the VPNB routing table.

```
[edit]
lab@R8# run show route advertising-protocol bgp 172.17.0.6

VPNB.inet.0: 28 destinations, 28 routes (28 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 172.17.31.0/24        Self              0
* 172.17.32.0/24        Self              0
* 172.17.33.0/24        Self              0
* 172.17.34.0/24        Self              0
* 172.17.35.0/24        Self              0
```

- 7) Ensure that all BGP sessions in VPNB can send and receive packets larger than the default 512-byte maximum segment size.

- a. R5 and R8

By default, BGP uses packets with a maximum size of 512-bytes. This way it is ensured that even if the MTU is lower on some links, BGP packets will not be fragmented. To force BGP to use bigger packets, you have to enable the MTU discovery functionality. This will automatically detect the maximum MTU between the peers and will use it for the size of the BGP packets.

```
[edit]
lab@R5# set routing-instances VPNB protocols bgp mtu-discovery

[edit]
lab@R8# set routing-instances VPNB protocols bgp mtu-discovery
```

- 8) Make sure that traffic coming from VPNB with precedence bits *010*, *011*, *100* and *100* is classified as *assured-forwarding*. Ensure 20% high priority bandwidth reservation for that traffic in your network, while keep the default reservation for rest of the traffic.

The last task requires providing QoS should be for VPNB only.

- a. R1, R2, R4, R5, R6 and R8

First, behavior aggregate classifiers are configured that will be later used for classifying the traffic.

The EXP classifier is needed to map the traffic coming from the MPLS core interfaces and associate it to forwarding-classes. To inherit the default mapping use the import function to use the default EXP classifier as a template. Next, change the mapping for the assured-forwarding classes as required by the task.

```
[edit]
lab@R1# set class-of-service classifiers exp custom-exp import default

[edit]
```

```
lab@R1# set class-of-service classifiers exp custom-exp forwarding-class assured-forwarding loss-priority low code-points 010
```

```
[edit]
```

```
lab@R1# set class-of-service classifiers exp custom-exp forwarding-class assured-forwarding loss-priority low code-points 011
```

```
[edit]
```

```
lab@R1# set class-of-service classifiers exp custom-exp forwarding-class assured-forwarding loss-priority low code-points 100
```

```
[edit]
```

```
lab@R1# set class-of-service classifiers exp custom-exp forwarding-class assured-forwarding loss-priority low code-points 101
```

An IP precedence classifier is intended to map the IP traffic coming from the CE sites to forwarding-classes. Again, the default classifier mappings are used as a template.

```
[edit]
```

```
lab@R1# set class-of-service classifiers inet-precedence VPNB-precedence import default
```

```
[edit]
```

```
lab@R1# set class-of-service classifiers inet-precedence VPNB-precedence forwarding-class assured-forwarding loss-priority low code-points 010
```

```
[edit]
```

```
lab@R1# set class-of-service classifiers inet-precedence VPNB-precedence forwarding-class assured-forwarding loss-priority low code-points 011
```

```
[edit]
```

```
lab@R1# set class-of-service classifiers inet-precedence VPNB-precedence forwarding-class assured-forwarding loss-priority low code-points 100
```

```
[edit]
```

```
lab@R1# set class-of-service classifiers inet-precedence VPNB-precedence forwarding-class assured-forwarding loss-priority low code-points 101
```

A custom scheduler-map and classifier are assigned to the MPLS core interfaces.

```
[edit]
```

```
lab@R1# set class-of-service interfaces ge-0/0/4 scheduler-map custom-map
```

```
[edit]
```

```
lab@R1# set class-of-service interfaces ge-0/0/4 unit * classifiers exp custom-exp
```

The custom scheduler-map is used to map the forwarding-classes to the schedulers.

```
[edit]
```

```
lab@R1# set class-of-service scheduler-maps custom-map forwarding-class assured-forwarding scheduler af
```

```
[edit]
```

```
lab@R1# set class-of-service scheduler-maps custom-map forwarding-class best-effort scheduler be
```

```
[edit]
```

```
lab@R1# set class-of-service scheduler-maps custom-map forwarding-class network-control scheduler nc
```

Custom schedulers are used to define the properties of output queues. The task requires modifying the assured-forwarding only, but it is a good practice to ensure the rest of the traffic is properly handled.

```
[edit]
lab@R1# set class-of-service schedulers af transmit-rate percent 20

[edit]
lab@R1# set class-of-service schedulers af priority high

[edit]
lab@R1# set class-of-service schedulers nc transmit-rate percent 5

[edit]
lab@R1# set class-of-service schedulers nc priority low

[edit]
lab@R1# set class-of-service schedulers be transmit-rate remainder

[edit]
lab@R1# set class-of-service schedulers be priority low
```

b. R5

By default, associating scheduler-map can be done under the physical interfaces under the class-of-service stanza. However, since you are not allowed to apply custom CoS policies to traffic other than VPNB, you have to enable `per-unit-scheduler` for the interface to the CE sites. This allows you to apply the custom scheduler only for a specific unit.

```
[edit]
lab@R5# set interfaces ge-0/0/5 per-unit-scheduler

[edit]
lab@R5# set class-of-service interfaces ge-0/0/5 unit 110 scheduler-map custom-map

[edit]
lab@R5# set class-of-service interfaces ge-0/0/5 unit 110 classifiers inet-precedence VPNB-precedence
```

When `vrf-table-label` is used for the routing instances, a default classifier is applied to the routing-instance logical interface. You have to replace the default EXP classifier with the custom one, to map the traffic coming from the MPLS core. The purpose of this classifier is to match on the vpn label after the transport label is stripped off.

```
[edit]
lab@R5# set class-of-service routing-instances VPNB classifiers exp custom-exp
```

c. R8

The same configuration is applied to router R8 as well.

```
[edit]
lab@R8# set interfaces ge-0/0/5 per-unit-scheduler

[edit]
lab@R8# set class-of-service interfaces ge-0/0/5 unit 111 scheduler-map custom-map

[edit]
lab@R8# set class-of-service interfaces ge-0/0/5 unit 111 classifiers inet-precedence VPNB-precedence

[edit]
lab@R8# set class-of-service routing-instances VPNB classifiers exp custom-exp
```

d. Verify the configuraiton

```
[edit]
lab@R5# run show route instance VPNB extensive
VPNB:
  Router ID: 172.17.0.1
  Type: vrf                               State: Active
  Interfaces:
    ge-0/0/5.110
    lsi.1
  Route-distinguisher: 10.50.250.5:200
  Vrf-import: [ VPNB-import.target ]
  Vrf-export: [ VPNB-export.target ]
  Fast-reroute-priority: low
  Tables:
    VPNB.inet.0      : 744 routes (491 active, 0 holddown, 0 hidden)
    VPNB.iso.0       : 0 routes (0 active, 0 holddown, 0 hidden)
    VPNB.inet6.0     : 0 routes (0 active, 0 holddown, 0 hidden)
    VPNB.mdt.0       : 0 routes (0 active, 0 holddown, 0 hidden)

[edit]
lab@R5# run show class-of-service routing-instance VPNB
Routing instance: VPNB

Logical interface: lsi.1      Index: 335
Object                          Name                               Type      Index
Classifier                       custom-exp                       exp        13646
```

Solution - Task 6.3: VPLS configuration

- 9) Configure VPLS L2 VPN between sites VPLS-1 and VPLS-2, using the interfaces as shown in the physical diagram. Since site VPLS-2 has connectivity to both R6 and R8, dedicate R8 as primary PE and R6 as backup PE. You have to use LDP signaled VPLS with automatic site discovery, hence you are not allowed to configure the sites explicitly.

The task description suggests that LDP signaled VPLS should connect the two sites. However, automatic site discovery must be used. Enable MP-BGP protocol for that purpose. This is known as FEC 129 LDP signaled VPLS.

- a. All devices including the route-reflector

A new MP-BGP family is added to BGP running in the network.

```
[edit]
lab@R1# set protocols bgp group internal family l2vpn auto-discovery-only
```

- b. R1

Configuring a new physical interface to connecting site VPLS-1. You have to exclude the apply-group to enable family mpls on logical interfaces. Otherwise, the commit will fail due to an error.

```
[edit]
lab@R1# set interfaces ge-0/0/3 apply-groups-except enable_mpls

[edit]
lab@R1# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R1# set interfaces ge-0/0/3 encapsulation vlan-vpls
```

```
[edit]
lab@R1# set interfaces ge-0/0/3 unit 3001 description "=== connection to VPLS-1
==="
```

```
[edit]
lab@R1# set interfaces ge-0/0/3 unit 3001 encapsulation vlan-vpls
```

```
[edit]
lab@R1# set interfaces ge-0/0/3 unit 3001 vlan-id 3001
```

```
[edit]
lab@R1# set interfaces ge-0/0/3 unit 3001 family vpls
```

Make sure the routing instance type is set to `vpls`. Because BGP is used for auto-discovery, you have to assign route-distinguisher and route-target to the routing instance configuration.

```
[edit]
lab@R1# set routing-instances VPLS instance-type vpls
```

```
[edit]
lab@R1# set routing-instances VPLS interface ge-0/0/3.3001
```

```
[edit]
lab@R1# set routing-instances VPLS route-distinguisher 10.50.250.1:300
```

The auto-discovery process require a unique extended L2 VPN community identifying the particular instance.

```
[edit]
lab@R1# set routing-instances VPLS l2vpn-id l2vpn-id:300:300
```

```
[edit]
lab@R1# set routing-instances VPLS vrf-target target:371786601L:300
```

The `no-tunnel-service` eliminates the need of hardware tunnel interface.

```
[edit]
lab@R1# set routing-instances VPLS protocols vpls interface ge-0/0/3.3001
```

```
[edit]
lab@R1# set routing-instances VPLS protocols vpls no-tunnel-services
```

c. R8

Configuring a new physical interface to connect site VPLS-2. You have to exclude the `apply-group` again.

```
[edit]
lab@R8# set interfaces ge-0/0/3 apply-groups-except enable_mpls
```

```
[edit]
lab@R8# set interfaces ge-0/0/3 description "=== connection to VPLS-2 ==="
```

```
[edit]
lab@R8# set interfaces ge-0/0/3 vlan-tagging
```

```
[edit]
lab@R8# set interfaces ge-0/0/3 encapsulation vlan-vpls
```

```
[edit]
lab@R8# set interfaces ge-0/0/3 unit 3002 encapsulation vlan-vpls
```

```
[edit]
```

```
lab@R8# set interfaces ge-0/0/3 unit 3002 vlan-id 3002
```

```
[edit]
```

```
lab@R8# set interfaces ge-0/0/3 unit 3002 family vpls
```

The route-target community must match on all routing-instances that are part of the same VPLS service.

```
[edit]
```

```
lab@R8# set routing-instances VPLS instance-type vpls
```

```
[edit]
```

```
lab@R8# set routing-instances VPLS interface ge-0/0/3.3002
```

```
[edit]
```

```
lab@R8# set routing-instances VPLS route-distinguisher 10.50.250.8:300
```

The extended L2 VPN community must also match from the same VPLS service.

```
[edit]
```

```
lab@R8# set routing-instances VPLS l2vpn-id l2vpn-id:300:300
```

```
[edit]
```

```
lab@R8# set routing-instances VPLS vrf-target target:371786601L:300
```

You have to configure multihoming to routers R8 and R6. JUNOS provide two methods of configuring multihoming with LDP signaled VPLS, FEC 128 and FEC 129. BGP auto-discovery is required for FEC 129 operation, therefore this method is used for multihoming site VPLS-2.

```
[edit]
```

```
lab@R8# set routing-instances VPLS protocols vpls no-tunnel-services
```

The vpls site configuration is moved under `multi-homing` section. That way the site interfaces are designated as part of the multihomed site. The identifier must match on both PEs connected to the same VPLS CE site.

```
[edit]
```

```
lab@R8# set routing-instances VPLS protocols vpls multi-homing site VPLS-2
identifier 2
```

The preference value defines the primary and backup PEs. A higher preference will take the role of designated forwarder (DF) for the multihomed site. Therefore 200 is set to router R8 and router R6 is set as backup. If the preference statement is omitted, the default value is 100.

```
[edit]
```

```
lab@R8# set routing-instances VPLS protocols vpls multi-homing site VPLS-2
preference 200
```

```
[edit]
```

```
lab@R8# set routing-instances VPLS protocols vpls multi-homing site VPLS-2
interface ge-0/0/3.3002
```

d. R6

Similar configuration is applied to router R6 .

```
[edit]
```

```
lab@R6# set interfaces ge-0/0/3 apply-groups-except enable_mpls
```

```

[edit]
lab@R6# set interfaces ge-0/0/3 description "=== connection to VPLS-2 ==="

[edit]
lab@R6# set interfaces ge-0/0/3 vlan-tagging

[edit]
lab@R6# set interfaces ge-0/0/3 encapsulation vlan-vpls

[edit]
lab@R6# set interfaces ge-0/0/3 unit 3002 encapsulation vlan-vpls

[edit]
lab@R6# set interfaces ge-0/0/3 unit 3002 vlan-id 3002

[edit]
lab@R6# set interfaces ge-0/0/3 unit 3002 family vpls

[edit]
lab@R6# set routing-instances VPLS instance-type vpls

[edit]
lab@R6# set routing-instances VPLS interface ge-0/0/3.3002

[edit]
lab@R6# set routing-instances VPLS route-distinguisher 10.50.250.6:300

[edit]
lab@R6# set routing-instances VPLS l2vpn-id l2vpn-id:300:300

[edit]
lab@R6# set routing-instances VPLS vrf-target target:371786601L:300

[edit]
lab@R6# set routing-instances VPLS protocols vpls no-tunnel-services

[edit]
lab@R6# set routing-instances VPLS protocols vpls multi-homing site VPLS-2
identifier 2

```

Router R6 is configured with preference backup, which means that it is less likely to become a DF.

```

[edit]
lab@R6# set routing-instances VPLS protocols vpls multi-homing site VPLS-2
preference backup

[edit]
lab@R6# set routing-instances VPLS protocols vpls multi-homing site VPLS-2
interface ge-0/0/3.3002

```

e. Verify the configuration

The VPLS is established.

```

[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
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

L2vpn-id: 300:300

Local-id: 10.50.250.1

```
Remote-id      Type  St      Time last up      # Up trans
10.50.250.6    rmt  Up      Dec 3 17:16:19 2016      1
Remote PE: 10.50.250.6, Negotiated control-word: No
Incoming label: 262156, Outgoing label: 262154
Negotiated PW status TLV: No
Local interface: lsi.1048587, Status: Up, Encapsulation: ETHERNET
Description: Intf - vpls VPLS local-id 10.50.250.1 remote-id 10.50.250.6
neighbor 10.50.250.6
Flow Label Transmit: No, Flow Label Receive: No
10.50.250.8    rmt  Up      Dec 3 17:16:44 2016      1
Remote PE: 10.50.250.8, Negotiated control-word: No
Incoming label: 262157, Outgoing label: 262154
Negotiated PW status TLV: No
Local interface: lsi.1048588, Status: Up, Encapsulation: ETHERNET
Description: Intf - vpls VPLS local-id 10.50.250.1 remote-id 10.50.250.8
neighbor 10.50.250.8
Flow Label Transmit: No, Flow Label Receive: No
```

Router R8 is designated as DF router for the VPLS-2 site.

[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
```

Legend for interface status

```
Up -- operational
Dn -- down
```

Instance: VPLS


```

L2vpn-id: 300:300
Local-id: 10.50.250.8
  Remote-id      Type  St      Time last up      # Up trans
  10.50.250.1    rmt   Up      Dec 5 12:23:28 2016      1
    Remote PE: 10.50.250.1, Negotiated control-word: No
    Incoming label: 262402, Outgoing label: 262146
    Negotiated PW status TLV: No
    Local interface: lsi.1048579, Status: Up, Encapsulation: ETHERNET
    Description: Intf - vpls VPLS local-id 10.50.250.8 remote-id 10.50.250.1
neighbor 10.50.250.1
  Flow Label Transmit: No, Flow Label Receive: No
  10.50.250.6    rmt   Up      Dec 5 12:25:17 2016      1
    Remote PE: 10.50.250.6, Negotiated control-word: No
    Incoming label: 262401, Outgoing label: 262145
    Negotiated PW status TLV: No
    Local interface: lsi.1048580, Status: Up, Encapsulation: ETHERNET
    Description: Intf - vpls VPLS local-id 10.50.250.8 remote-id 10.50.250.6
neighbor 10.50.250.6
  Flow Label Transmit: No, Flow Label Receive: No
Multi-home:
  Local-site      Id      Pref   State
  VPLS-2          2      200    Up

```

[edit]

lab@R6# **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	

Legend for interface status

Up -- operational
Dn -- down

Instance: VPLS

```

L2vpn-id: 300:300
Local-id: 10.50.250.6
  Remote-id      Type  St      Time last up      # Up trans
  10.50.250.1    rmt   Up      Dec 3 14:58:50 2016      1
    Remote PE: 10.50.250.1, Negotiated control-word: No
    Incoming label: 262154, Outgoing label: 262156
    Negotiated PW status TLV: No
    Local interface: lsi.1048586, Status: Up, Encapsulation: ETHERNET
    Description: Intf - vpls VPLS local-id 10.50.250.6 remote-id 10.50.250.1
neighbor 10.50.250.1
  Flow Label Transmit: No, Flow Label Receive: No
  10.50.250.8    rmt   Up      Dec 3 14:59:09 2016      1
    Remote PE: 10.50.250.8, Negotiated control-word: No
    Incoming label: 262155, Outgoing label: 262155
    Negotiated PW status TLV: No
    Local interface: lsi.1048587, Status: Up, Encapsulation: ETHERNET

```

Description: Intf - vpls VPLS local-id 10.50.250.6 remote-id 10.50.250.8 neighbor 10.50.250.8

Flow Label Transmit: No, Flow Label Receive: No

Multi-home:

Local-site	Id	Pref	State
VPLS-2	2	1	Up

The route-reflector receives BGP auto-discovery routes.

```
[edit]
lab@route-reflector# run show route table bgp.l2vpn.0

bgp.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.50.250.1:300:10.50.250.1/96 AD
    *[BGP/170] 02:45:57, localpref 100, from 10.50.250.1
    AS path: I, validation-state: unverified
    > to 10.50.0.97 via ge-0/0/1.49
10.50.250.6:300:10.50.250.6/96 AD
    *[BGP/170] 02:45:59, localpref 100, from 10.50.250.6
    AS path: I, validation-state: unverified
    > to 10.50.0.97 via ge-0/0/1.49
    to 10.50.0.101 via ge-0/0/1.59
10.50.250.8:300:10.50.250.8/96 AD
    *[BGP/170] 02:45:43, localpref 100, from 10.50.250.8
    AS path: I, validation-state: unverified
    > to 10.50.0.101 via ge-0/0/1.59
```

- 10) One more site attached to R2 should be connected to the VPLS VPN. Assume that R2 supports only L2 circuit VPNs. Interconnect the L2VPN-3 site with the VPLS VPN at R8 as primary and R6 as a back-up connection.

L2 circuit VPN uses LDP for signaling; therefor a targeted LDP session is required between the PEs connecting the two sites. However, LDP targeted sessions already exist due to the LDP tunneling configuration.

a. R2

Configure the interface to connect site L2VPN-3 to the network. Again, the apply group must be excluded.

```
[edit]
lab@R2# set interfaces apply-groups enable_mpls

[edit]
lab@R2# set interfaces ge-0/0/3 apply-groups-except enable_mpls

[edit]
lab@R2# set interfaces ge-0/0/3 description "=== connection to L2VPN-3 ==="
```

The encapsulation used on the interface must be ccc, as well as the family on logical interface.

```
[edit]
lab@R2# set interfaces ge-0/0/3 encapsulation ethernet-ccc

[edit]
lab@R2# set interfaces ge-0/0/3 unit 0 family ccc
```

The interface is added to the l2circuit neighbor configuration, together with the backup neighbor. The virtual circuit id must match on both ends of the connection.

```
[edit]
lab@R2# set protocols l2circuit neighbor 10.50.250.8 interface ge-0/0/3.0
virtual-circuit-id 3

[edit]
lab@R2# set protocols l2circuit neighbor 10.50.250.8 interface ge-0/0/3.0 backup-
neighbor 10.50.250.6 virtual-circuit-id 3
```

b. R8

At the other end, you have to 'stitch' the VPLS configuration to the L2 circuit VPN. To do that you have to use a logical-tunnel interface. The peering of two lt interfaces units provides a point-to-point logical connection. To enable the use of lt interfaces on MX routers you have to enable tunnel-services under at least one of the FPC cards available in the chassis.

```
[edit]
lab@R8# set chassis fpc 0 pic 0 tunnel-services

[edit]
lab@R8# set interfaces lt-0/0/0 apply-groups-except enable_mpls
```

Unit 1 is assigned to the VPLS routing instance, therefore VPLS encapsulation should be used.

```
[edit]
lab@R8# set interfaces lt-0/0/0 unit 1 encapsulation ethernet-vpls
```

Unit 1 peers with unit 2 to establish a logical connection.

```
[edit]
lab@R8# set interfaces lt-0/0/0 unit 1 peer-unit 2
```

Similarly, unit 2 is configured with encapsulation ccc, since it is attached to the L2 circuit neighbor at the local end of the L2 circuit configuration.

```
[edit]
lab@R8# set interfaces lt-0/0/0 unit 2 encapsulation ethernet-ccc
```

Unit 2 peers with unit 1 to establish a logical connection.

```
[edit]
lab@R8# set interfaces lt-0/0/0 unit 2 peer-unit 1
```

Unit 1 is assigned to the VPLS routing instance configured already.

```
[edit]
lab@R8# set routing-instances VPLS interface lt-0/0/0.1
```

Unit 2 is used at local end of the L2 circuit definition.

```
[edit]
lab@R8# set protocols l2circuit neighbor 10.50.250.2 interface lt-0/0/0.2
virtual-circuit-id 3
```

c. R6

Similar configuration is applied to router R6.

```
[edit]
lab@R6# set chassis fpc 0 pic 0 tunnel-services

[edit]
lab@R6# set interfaces lt-0/0/0 apply-groups-except enable_mpls

[edit]
lab@R6# set interfaces lt-0/0/0 unit 1 encapsulation ethernet-vpls

[edit]
lab@R6# set interfaces lt-0/0/0 unit 1 peer-unit 2

[edit]
lab@R6# set interfaces lt-0/0/0 unit 2 encapsulation ethernet-ccc

[edit]
lab@R6# set interfaces lt-0/0/0 unit 2 peer-unit 1

[edit]
lab@R6# set routing-instances VPLS interface lt-0/0/0.1
```

Router R2 defines the backup role. The L2 circuit configuration on routers R6 and R8 is exactly the same.

```
[edit]
lab@R6# set protocols l2circuit neighbor 10.50.250.2 interface lt-0/0/0.2
virtual-circuit-id 3
```

d. Verify the configuration

The L2 circuit connection from router R2 to router R8 is in state UP, whereas the connection to R6 is in state Backup. Router R2 does not signal a connection to R6 until the connection to router R8 is DOWN.

```
[edit]
lab@R2# 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: 10.50.250.6

Interface	Type	St	Time last up	# Up trans
ge-0/0/3.0(vc 3)	rmt	BK		

Neighbor: 10.50.250.8

Interface	Type	St	Time last up	# Up trans
ge-0/0/3.0(vc 3)	rmt	Up	Dec 5 12:29:25 2016	1

Remote PE: 10.50.250.8, Negotiated control-word: Yes (Null)

Incoming label: 299793, Outgoing label: 299787

Negotiated PW status TLV: No

Local interface: ge-0/0/3.0, Status: Up, Encapsulation: ETHERNET

Flow Label Transmit: No, Flow Label Receive: No

```
[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: 10.50.250.2
```

Interface	Type	St	Time last up	# Up trans
lt-0/0/0.2(vc 3)	rmt	Up	Dec 5 12:19:57 2016	1

```
Remote PE: 10.50.250.2, Negotiated control-word: Yes (Null)
Incoming label: 299782, Outgoing label: 299782
Negotiated PW status TLV: No
Local interface: lt-0/0/0.2, Status: Up, Encapsulation: ETHERNET
Flow Label Transmit: No, Flow Label Receive: No
```

```
[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: 10.50.250.2
```

Interface	Type	St	Time last up	# Up trans
lt-0/0/0.2(vc 3)	rmt	OL		

- 11) Protect the VPLS VPN from excessive broadcast and multicast traffic send by the users attached behind the VPLS sites. Ensure that not more than 30Mbps of such traffic can be send by a single site.

a. R1, R6 and R8

First, a policer needs to be configured according to the requirements. A firewall filter is configured with a policer. Since, the filter is intended to limit VPLS traffic, it is configured under family vpls.

```
[edit]
lab@R1# set firewall policer VPLS-limit.30 if-exceeding bandwidth-limit 30m

[edit]
lab@R1# set firewall policer VPLS-limit.30 if-exceeding burst-size-limit 750k

[edit]
lab@R1# set firewall policer VPLS-limit.30 then discard

[edit]
lab@R1# set firewall family vpls filter VPLS.police term 1 then policer VPLS-limit.30
```

The place where that firewall filter is applied is important here. The task asks for restricting the broadcast and multicast traffic. Broadcast and multicast traffic is handled in a different way than the unicast traffic by the VPLS engine. It is flooded to more than one or all other PEs part of the particular VPLS. To filter the flooded traffic, you have to apply the filter to the routing instance forwarding table.

```
[edit]
lab@R1# set routing-instances VPLS forwarding-options family vpls flood input VPLS.police
```

- 12) Make sure all 3 sites are able to communicate. However, you are not allowed to manipulate the VLAN tags directly on the interfaces.

The task requests fixing connectivity between all 3 sites: VPLS-1, VPLS-2 and L2VPN-3. This is needed because the two VPLS CE sites are sending the packets tagged using different VLANs. Moreover, the third site connected via L2 circuit is sending untagged packets. To complete the task you have to make sure that the frames traveling through the MPLS core are untagged. To do that you have to configure the PEs connected to the VPLS to remove the tag when frames are sent through the core and insert the correct tag when frames are sent to the CE devices. Since, using vlan maps is not allowed, the other option JUNOS provides is to use VLAN normalization for VPLS instances.

- a. R1, R6 and R8

Configuring `vlan-id none` at the vpls routing instance will remove the vlan tag from the incoming frames and put the correct frame on the other end.

```
[edit]
lab@R1# set routing-instances VPLS vlan-id none
```

- b. Verify the configuration

The output below shows the vlan operation will be performed in ingress and egress direction.

```
[edit]
lab@R1# run show interfaces ge-0/0/3.3001
Logical interface ge-0/0/3.3001 (Index 332) (SNMP ifIndex 540)
Description: === connection to VPLS-1 ===
Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.3001 ] In(pop) Out(push
0x8100.3001) Encapsulation: VLAN-VPLS
```

```

Input packets : 0
Output packets: 0
Protocol vpls, MTU: 1518
Flags: Is-Primary

```

```

[edit]
root@vr-device# run ping 192.168.5.2 routing-instance L2VPN-3
PING 192.168.5.2 (192.168.5.2): 56 data bytes
64 bytes from 192.168.5.2: icmp_seq=0 ttl=63 time=27.131 ms
64 bytes from 192.168.5.2: icmp_seq=1 ttl=64 time=2.601 ms
64 bytes from 192.168.5.2: icmp_seq=2 ttl=64 time=2.362 ms
64 bytes from 192.168.5.2: icmp_seq=3 ttl=64 time=2.420 ms
^C
--- 192.168.5.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 2.362/8.629/27.131/10.683 ms

[edit]
root@vr-device# run ping 192.168.5.1 routing-instance L2VPN-3
PING 192.168.5.1 (192.168.5.1): 56 data bytes
64 bytes from 192.168.5.1: icmp_seq=0 ttl=64 time=3.873 ms
64 bytes from 192.168.5.1: icmp_seq=1 ttl=64 time=3.336 ms
64 bytes from 192.168.5.1: icmp_seq=2 ttl=64 time=22.139 ms
^C
--- 192.168.5.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.336/9.783/22.139/8.740 ms

[edit]
root@vr-device# run ping 192.168.5.1 routing-instance VPLS-2-L3
PING 192.168.5.3 (192.168.5.3): 56 data bytes
64 bytes from 192.168.5.3: icmp_seq=0 ttl=64 time=2.449 ms
64 bytes from 192.168.5.3: icmp_seq=1 ttl=64 time=2.513 ms
64 bytes from 192.168.5.3: icmp_seq=2 ttl=64 time=2.330 ms
64 bytes from 192.168.5.3: icmp_seq=3 ttl=64 time=2.339 ms
64 bytes from 192.168.5.3: icmp_seq=4 ttl=64 time=2.330 ms
^C
--- 192.168.5.3 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 2.330/2.392/2.513/0.075 ms

```