Juniper Networks | Education Services

Advanced Junos Service Provider Routing

LAB GUIDE

Revision V19A

Education Services Courseware

Advanced Junos Service Provider Routing

V19A

Lab Guide
Volume 1 of 2

Education Services

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

Course Number: EDU-JUN-AJSPR

This document is produced by Juniper Networks, Inc.

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

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Advanced Junos Service Provider Routing Lab Guide, Revision V19A

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

Printed in USA.

Revision History:

Revision 10.a — March 2011

Revision 10.b—September 2011

Revision 11.a—January 2012

Revision 12.a—September 2013

Revision V-17.a—June 2017

Revision V19A—February 2020

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

The information in this document has been carefully verified and is believed to be accurate for software Release 19.3. Juniper Networks assumes no responsibilities for any inaccuracies that may appear in this document. In no event will Juniper Networks be liable for direct, indirect, special, exemplary, incidental, or consequential damages resulting from any defect or omission in this document, even if advised of the possibility of such damages.

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

YEAR 2000 NOTICE

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

SOFTWARE LICENSE

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

Contents

Lab 1: Configuring and Monitoring OSPF 1-1

Part 1: Load Starting Configuration 1-2
Part 2: Creating the Network 1-15
Part 3: Configuring the OSPF Link Costs 1-24
Part 4: Configuring Overload 1-29
Part 5: Performing Authentication 1-32

Lab 2: Advanced OSPF 2-1

Part 1: Load Starting Configuration 2-2
Part 2: Creating OSPF Stub Areas 2-17
Part 3: Creating Stub No Summaries Areas 2-21
Part 4: Creating OSPF Not-So-Stubby-Areas 2-26
Part 5: Creating NSSA No Summaries 2-30

Lab 3: Advanced OSPF Options 3-1

Part 1: Load Starting Configuration 3-2
Part 2: Establishing the OSPF Adjacencies and Creating a Virtual Tunnel 3-19
Part 3: Configuring OSPF Multiarea 3-25
Part 4: Configuring External Reachability 3-30

Lab 4: Troubleshooting OSPF 4-1

Part 1: Load Starting Configuration 4-2
Part 2: Troubleshooting OSPF 4-4

Lab 5: IS-IS Configuration and Monitoring 5-1

Part 1: Load Starting Configuration 5-2
Part 2: Configuring the Transit Interfaces to Support ISO Packets 5-16
Part 3: Configuring the IS-IS Network Entity Title 5-21
Part 4: Configuring Interfaces as Part of the IS-IS Protocol 5-25
Part 5: Examining the IS-IS Database 5-35

Lab 6: Advanced IS-IS Configuration and Routing Policy 6-1

Part 1: Load Starting Configuration 6-2
Part 2: Configuring IS-IS Authentication 6-23
Part 3: Manipulating IS-IS Metrics 6-27
Part 4: Configuring IS-IS External Reachability 6-36

Lab 7: Configuring a Multilevel IS-IS Network 7-1

Part 1: Load Starting Configuration 7-2
Part 2: Examining the IS-IS Multilevel Flooding Scope 7-23
Part 3: Modifying the Default Flooding Scope 7-28

Lab 8: Troubleshooting IS-IS 8-1

Part 1: Load Starting Configuration 8-2
Part 2: Troubleshooting IS-IS 8-4

Lab 9: Configuring BGP 9-1
Part 1: Load Starting Configuration 9-2
Part 2: Establishing an IBGP Peering Session 9-29
Part 3: Configuring the P1 and P2 EBGP Peers 9-35
Part 4: Configuring the EBGP Session with the P3 Router 9-43
Part 5: Summarizing the Internal Routes to the Peer Routers 9-46

Course Overview

This five-day course is designed to provide students with detailed coverage of OSPF, IS-IS, BGP, and routing policy. Through demonstrations and hands-on labs, students will gain experience in configuring, monitoring, and troubleshooting the Junos operating system and in monitoring device and protocol operations.

This course uses Juniper Networks vMX Series Routers for the hands-on component, but the lab environment does not preclude the course from being applicable to other Juniper hardware platforms running the Junos OS. This course is based on the Junos OS Release 19.3.

Course Level

Advanced Junos Service Provider Routing (AJSPR) is an advanced-level course.

Intended Audience

This course benefits individuals responsible for implementing, monitoring, and troubleshooting Layer 3 components of a service provider's network.

Prerequisites

Students should have intermediate-level networking knowledge and an understanding of the Open Systems Interconnection (OSI) model and the TCP/IP protocol suite. Students should also attend the *Introduction to the Junos Operating System* (IJOS) and *Junos Intermediate Routing* (JIR) courses prior to attending this class.

Objectives

After successfully completing this course, you should be able to:

- Describe the various OSPF link-state advertisement (LSA) types.
- Explain the flooding of LSAs in an OSPF network.
- Describe the shortest-path-first (SPF) algorithm.
- Explain OSPF link metrics.
- Describe the various OSPF authentication methods.
- Explain the key differences between OSPFv2 and OSPFv3.
- Describe OSPF area types and operations.
- Configure various OSPF area types.
- Summarize and restrict routes.
- Configure OSPF multi-area adjacencies
- Configure OSPF virtual links.
- Explain OSPF external reachability.
- List useful commands that are used to troubleshoot and verify OSPF.
- Isolate different OSPF issues.
- Explain the concepts and operation of IS-IS.
- Describe various IS-IS link-state
- protocol data unit (LSP) types.
- List IS-IS adjacency rules and troubleshoot common adjacency issues.
- Configure and monitor IS-IS.
- Display and interpret the link-state database (LSDB).
- Perform advanced IS-IS configuration options.
- Implement IS-IS routing policy.
- Explain the default operation in multilevel IS-IS.

- Describe address summarization methods used in IS-IS.
- Configure and monitor a multilevel IS-IS network.
- List useful commands to troubleshoot and verify IS-IS problems.
- Troubleshoot and isolate different IS-IS issues.
- Describe basic BGP operation.
- List common BGP attributes.
- Explain the route selection process for BGP.
- Describe how to alter the route selection process.
- Configure some advanced options for BGP peers.
- Explain how policies function in BGP.
- Describe BGP attributes and explain how these attributes can be used to manipulate traffic.
- Show how Regex can be used in policies to manipulate AS-path.
- Describe the BGP attributes Origin, MED and communities in detail and explain the operation of those attributes.
- Manipulate these BGP attributes using routing policy.
- Describe the operation of BGP route reflection.
- Configure a route reflector.
- Describe the virtual route reflector.
- Describe the operation of optimal route reflection.
- Describe the operation of a BGP confederation.
- Configure confederations.
- Describe peering relationships in a confederation.
- Describe DDoS attacks.
- Describe DDoS mitigation techniques.
- Describe FlowSpec Operations.

Course Agenda

Day 1

Chapter 1: Course Introduction

Chapter 2: OSPF

Lab 1: Configuring OSPF

Chapter 3: OSPF Areas

Lab 2: Advanced OSPF

Chapter 4: Advanced OSPF Options

Lab 3: Advanced OSPF Options

Day 2

Chapter 5: Troubleshooting OSPF

Lab 4: Troubleshooting OSPF

Chapter 6: IS-IS

Lab 5: IS-IS Configuration and Monitoring

Chapter 7: Advanced IS-IS Operations

Lab 6: Advanced IS-IS Configuration Options and Routing Policy

Day 3

Chapter 8: Multilevel IS-IS Networks

Lab 7: Configuring a Multilevel IS-IS Network

Chapter 9: Troubleshooting IS-IS

Lab 8: Troubleshooting IS-IS

Chapter 10: BGP

Lab 9: Configuring BGP

Day 4

Chapter 11: BGP Attributes and Policy—Part 1

Lab 10: BGP Attributes – Part 1

Chapter 12: BGP Attributes and Policy—Part 2

Lab 11: BGP Attributes – Part 2

Chapter 13: Route Reflection and Confederations

Lab 12: Scaling BGP

Day 5

Chapter 14: BGP FlowSpec

Lab 13: BGP FlowSpec

Chapter 15: Troubleshooting BGP

Lab 14: Troubleshooting BGP

Chapter 16: Troubleshooting Policy

Lab 15: Troubleshooting Routing Policy

Appendix A: BGP Route Damping

Lab 16: BGP Route Damping (Optional)

Document Conventions

CLI and GUI Text

Frequently throughout this course, we refer to text that appears in a command-line interface (CLI) or a graphical user interface (GUI). To make the language of these documents easier to read, we distinguish GUI and CLI text from chapter text according to the following table.

Style	Description	Usage Example
Franklin Gothic	Normal text.	Most of what you read in the Lab Guide and Student Guide.

`Courier New`	Console text:
	• Screen captures
	• Noncommand-related syntax
	GUI text elements:
	• Menu names
	• Text field entry
		
		`commit complete`
		`Exiting configuration mode`
		
		Select `File > Open`, and then click `Configuration.conf` in the `Filename` text box.

Input Text Versus Output Text

You will also frequently see cases where you must enter input text yourself. Often these instances will be shown in the context of where you must enter them. We use bold style to distinguish text that is input versus text that is simply displayed.

Style	Description	Usage Example
`Normal CLI`	No distinguishing variant.	`Physical interface:fxp0, Enabled`
`Normal GUI`		View configuration history by clicking `Configuration > History`.
`CLI Input`	Text that you must enter.	`lab@San_Jose>` **`show route`**
`GUI Input`		Select `File > Save`, and type **`config.ini`** in the `Filename` field.

Undefined Syntax Variables

Finally, this course distinguishes syntax variables, where you must assign the value (undefined variables). Note that these styles can be combined with the input style as well.

Style	Description	Usage Example
`CLI Undefined` *`GUI Undefined`*	Text where the variable's value is the user's discretion or text where the variable's value as shown in the lab guide might differ from the value the user must input according to the lab topology.	Type **`set policy`** ***`policy-name`***. **`ping 10.0.`*****`x.y`*** Select `File > Save`, and type ***`filename`*** in the `Filename` field.

Additional Information

Education Services Offerings

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

About This Publication

This course was developed and tested using the software release listed on the copyright page. Previous and later versions of software might behave differently so you should always consult the documentation and release notes for the version of code you are running before reporting errors.

This document is written and maintained by the Juniper Networks Education Services development team. Please send questions and suggestions for improvement to training@juniper.net.

Technical Publications

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

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

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

Juniper Networks Support

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

Lab

Configuring and Monitoring OSPF

Overview

In this lab, you use the lab diagram titled "Lab Network Diagram: Configuring OSPF" to establish a multi-area OSPF routing domain. You explore the operation of the network focusing on **show** commands and the link-state database (LSDB). You then explore configuration options, such as reference bandwidth, overload, and authentication.

By completing this lab, you will perform the following tasks:

- Verify the router's existing configuration.
- Verify the router's interface status.
- Build a multiarea OSPF network.
- Change OSPF costs on links.
- Configure a router for overload.
- Perform authentication on OSPF packets.

Part 1: Load Starting Configuration

In this lab part, you verify the initial configuration of the routers. You then verify that the interfaces are operational.

Note

The instructor will tell you the nature of your access and will provide you with the necessary details to access your assigned device.

Note

The lab topology requires you to display information in the different logical systems of mxD. You will be changing the perspective of the CLI by issuing the **`set cli logical-system`** ***`ls-name`*** command. Refer to the lab diagram for the correct logical system name.

By changing the perspective of the CLI, you will be able to perform network commands such as **`ping`** or **`traceroute`** from the perspective of the pertinent logical system.

Step 1.1

Ensure that you know to which student device you have been assigned. Check with your instructor if you are not certain. Consult the management network diagram to determine the management address of the student devices.

Access the CLI on your mxD device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab1-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.4
Last login: Tue Nov 19 18:29:49 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab1-start.config
load complete

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

lab@mxD>
```

Step 1.2

Issue the **show configuration** command. Use the lab diagram to verify that the mxD router has two logical systems, R3-1 and R3-2, that have the correct interface configuration.

```
lab@mxD> show configuration
## Last commit: 2019-11-19 18:32:50 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxD;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    R3-1 {
        interfaces {
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 10.0.1.2/24;
                    }
                }
            }
            lo0 {
                unit 1 {
```

```
                    family inet {
                        address 172.16.1.2/32;
                    }
                }
            }
        }
        routing-options {
            autonomous-system 65512;
        }
    }
    R3-2 {
        interfaces {
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 10.0.2.2/24;
                    }
                }
            }
            lo0 {
                unit 2 {
                    family inet {
                        address 172.16.2.2/32;
                    }
                }
            }
        }
        routing-options {
            autonomous-system 65512;
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.4/24;
            }
        }
    }
}

lab@mxD>
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxD router. If the configuration is incorrect, please notify your instructor.

Step 1.3

Access the CLI on your mxA device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123**.

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab1-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.1
Last login: Tue Nov 19 18:27:02 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# load override ajspr/lab1-start.config
load complete

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

lab@mxA>
```

Step 1.4

Issue the **show configuration** command. Use the lab diagram to verify that the mxA router has the correct interface configuration. Verify that there are four 20.20/24 static routes and a policy statement.

```
lab@mxA> show configuration
## Last commit: 2019-12-26 10:12:54 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxA;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
```

```
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.22.121.1/24;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.1.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.1/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.1.1/32;
            }
        }
    }
}
policy-options {
    policy-statement static-to-ospf {
        term 1 {
            from protocol static;
            then accept;
        }
    }
}
routing-options {
    static {
        route 20.20.0.0/24 reject;
```

```
        route 20.20.1.0/24 reject;
        route 20.20.2.0/24 reject;
        route 20.20.3.0/24 reject;
    }
    autonomous-system 65512;
}
lab@mxA>
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxA router. If the configuration is incorrect, please notify your instructor.

Step 1.5

Use the **`ping 10.0.1.2 count 5`** command to verify that you can ping the physical interface on the R3-1 router.

```
lab@mxA> ping 10.0.1.2 count 5
PING 10.0.1.2 (10.0.1.2): 56 data bytes
64 bytes from 10.0.1.2: icmp_seq=0 ttl=64 time=4.347 ms
64 bytes from 10.0.1.2: icmp_seq=1 ttl=64 time=1.434 ms
64 bytes from 10.0.1.2: icmp_seq=2 ttl=64 time=18.736 ms
64 bytes from 10.0.1.2: icmp_seq=3 ttl=64 time=1.416 ms
64 bytes from 10.0.1.2: icmp_seq=4 ttl=64 time=1.491 ms

--- 10.0.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.416/5.485/18.736/6.720 ms

lab@mxA>
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.6

To aid in completing the labs in a timely manner, some routing information has been preconfigured on your router. Issue the **`show route protocol static table inet.0`** command to ensure that the correct routing information is present.

```
lab@mxA> show route protocol static table inet.0

inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
20.20.0.0/24       *[Static/5] 02:50:32
                      Reject
20.20.1.0/24       *[Static/5] 02:50:32
                      Reject
20.20.2.0/24       *[Static/5] 02:50:32
                      Reject
20.20.3.0/24       *[Static/5] 02:50:32
                      Reject

lab@mxA>
```

Question: How many static route entries are in your routing table?

_____

Answer: There should be four static route entries. Please notify the instructor if your router does not currently have this configuration.

Step 1.7

Access the CLI on your mxB device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123**.

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab1-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.2
Last login: Tue Nov 19 18:28:56 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxB> configure
Entering configuration mode

[edit]
lab@mxB# load override ajspr/lab1-start.config
load complete

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

lab@mxB>
```

Step 1.8

Issue the **show configuration** command. Use the lab diagram to verify that the mxB router has the correct interface configuration. Verify that OSPF has been enabled (all interfaces in area 0).

```
lab@mxB> show configuration
## Last commit: 2019-11-19 18:37:23 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxB;
```

```
 root-authentication {
     encrypted-password
"$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
.3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
 }
 login {
     user lab {
         uid 2000;
         class super-user;
         authentication {
             encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
         }
     }
 }
 services {
     ssh;
     netconf {
         ssh;
     }
 }
 syslog {
     user * {
         any emergency;
     }
     file messages {
         any notice;
         authorization info;
     }
     file interactive-commands {
         interactive-commands any;
     }
 }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.22.121.2/24;
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.22.122.2/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.2/24;
            }
        }
```

```
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.31.100.1/32;
            }
        }
    }
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-0/0/0.0;
            interface ge-0/0/1.0;
        }
    }
}

lab@mxB>
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxB router. If the configuration is incorrect, please notify your instructor.

Step 1.9

Use the **`ping 172.22.121.1 count 5`** command to verify that you can ping the physical interface on the mxA router.

```
lab@mxB> ping 172.22.121.1 count 5
PING 172.22.121.1 (172.22.121.1): 56 data bytes
64 bytes from 172.22.121.1: icmp_seq=0 ttl=64 time=27.643 ms
64 bytes from 172.22.121.1: icmp_seq=1 ttl=64 time=4.170 ms
64 bytes from 172.22.121.1: icmp_seq=2 ttl=64 time=1.454 ms
64 bytes from 172.22.121.1: icmp_seq=3 ttl=64 time=106.790 ms
64 bytes from 172.22.121.1: icmp_seq=4 ttl=64 time=1.608 ms

--- 172.22.121.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.454/28.333/106.790/40.439 ms

lab@mxB>
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.10

Access the CLI on your mxC device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab1-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode

```
[lab@desktop ~]$ ssh lab@172.25.11.3
Last login: Tue Nov 19 18:29:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# load override ajspr/lab1-start.config
load complete

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

lab@mxC>
```

Step 1.11

Issue the **show configuration** command. Use the lab diagram to verify that the mxC router has the correct interface configuration. Verify that there are four 20.20/24 static routes and a policy statement.

```
lab@mxC> show configuration
## Last commit: 2019-12-26 10:13:07 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxC;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
```

```
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.22.122.1/24;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.2.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.3/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.2.1/32;
            }
        }
    }
}
policy-options {
    policy-statement static-to-ospf {
        term 1 {
            from protocol static;
            then accept;
        }
    }
}
routing-options {
    static {
        route 20.20.4.0/24 reject;
```

```
        route 20.20.5.0/24 reject;
        route 20.20.6.0/24 reject;
        route 20.20.7.0/24 reject;
    }
    autonomous-system 65512;
}
lab@mxC>
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxC router. If the configuration is incorrect, please notify your instructor.

Step 1.12

Use the **`ping 10.0.2.2 count 5`** command to verify that you can ping the physical interface on the R3-2 router.

```
lab@mxC> ping 10.0.2.2 count 5
PING 10.0.2.2 (10.0.2.2): 56 data bytes
64 bytes from 10.0.2.2: icmp_seq=0 ttl=64 time=1.476 ms
64 bytes from 10.0.2.2: icmp_seq=1 ttl=64 time=1.162 ms
64 bytes from 10.0.2.2: icmp_seq=2 ttl=64 time=1.538 ms
64 bytes from 10.0.2.2: icmp_seq=3 ttl=64 time=1.594 ms
64 bytes from 10.0.2.2: icmp_seq=4 ttl=64 time=1.588 ms

--- 10.0.2.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.162/1.472/1.594/0.161 ms

lab@mxC>
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.13

Use the **ping 172.22.122.2 count 5** command to verify that you can ping the physical interface on the mxB router.

```
lab@mxC> ping 172.22.122.2 count 5
PING 172.22.122.2 (172.22.122.2): 56 data bytes
64 bytes from 172.22.122.2: icmp_seq=0 ttl=64 time=1.745 ms
64 bytes from 172.22.122.2: icmp_seq=1 ttl=64 time=151.563 ms
64 bytes from 172.22.122.2: icmp_seq=2 ttl=64 time=1.497 ms
64 bytes from 172.22.122.2: icmp_seq=3 ttl=64 time=1.592 ms
64 bytes from 172.22.122.2: icmp_seq=4 ttl=64 time=1.845 ms

--- 172.22.122.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.497/31.648/151.563/59.957 ms

lab@mxC>
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.14

To aid in completing the labs in a timely manner, some routing information has been preconfigured on your router. Issue the **show route protocol static table inet.0** command to ensure that the correct routing information is present.

```
lab@mxC> show route protocol static table inet.0

inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.4.0/24       *[Static/5] 03:00:37
                      Reject
20.20.5.0/24       *[Static/5] 03:00:37
                      Reject
20.20.6.0/24       *[Static/5] 03:00:37
                      Reject
20.20.7.0/24       *[Static/5] 03:00:37
                      Reject

lab@mxC>
```

Question: How many static route entries are in your routing table?

_____

Answer: There should be four static route entries. Please notify the instructor if your router does not currently have this configuration.

Part 2: Creating the Network

In this lab part, you configure and monitor a multiarea OSPF network. First, you configure the interfaces participating in OSPF. You then configure your device to participate in a multiarea OSPF network and verify operations using command-line interface (CLI) operational mode commands.

Step 2.1

Refer to the network diagram in your lab topology handout. Write down the interfaces that will run OSPF and to which area each is attached.

Interfaces and OSPF areas:

_____

_____

_____

_____

Question: Which routers are area border routers (ABRs)?

_____

Answer: The mxA and mxC routers are the ABRs.

Question: Which routers are internal backbone routers?

_____

Answer: The mxB router is a backbone router.

Question: Which routers are internal non-backbone routers?

_____

Answer: The R3 logical systems on the mxD router are internal non-backbone routers.

Step 2.2

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Configure mxA's interfaces and unit numbers in the correct OSPF area. Do not forget the loopback interface. Commit your configuration when completed.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit protocols ospf

[edit protocols ospf]
lab@mxA# set area 0 interface lo0.0
```

```
[edit protocols ospf]
lab@mxA# set area 0 interface ge-0/0/0.0

[edit protocols ospf]
lab@mxA# set area 10 interface ge-0/0/3.0

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 2.3

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxD> set cli logical-system R3-1
Logical system: R3-1

lab@mxD:R3-1>
```

Step 2.4

Enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Configure R3-1's interfaces and unit numbers in the correct OSPF area. Do not forget the loopback interface. Commit your configuration and exit out of configuration mode when completed.

```
lab@mxD:R3-1> configure
Entering configuration mode

[edit]
lab@mxD:R3-1# edit protocols ospf

[edit protocols ospf]
lab@mxD:R3-1# set area 10 interface ge-0/0/1.0

[edit protocols ospf]
lab@mxD:R3-1# set area 10 interface lo0.1

[edit protocols ospf]
lab@mxD:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R3-1>
```

Step 2.5

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, verify the OSPF adjacencies with the neighboring routers by issuing the **`run show ospf neighbor`** command.

```
[edit protocols ospf]
lab@mxA# run show ospf neighbor
```

```
Address          Interface              State     ID               Pri  Dead
172.22.121.2     ge-0/0/0.0             Full      172.31.100.1     128    38
10.0.1.2         ge-0/0/3.0             Full      172.16.1.2       128    37
```

Question: Does the neighbor adjacency state show `Full` for both OSPF neighbors?

_____

Answer: Yes, the adjacency state should show `Full` for both OSPF neighbors.

Step 2.6

Issue the **`run show route`** command to look at the routing table of the mxA router.

```
[edit protocols ospf]
lab@mxA# run show route

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.1.0/24        *[Direct/0] 03:38:09
                    > via ge-0/0/3.0
10.0.1.1/32        *[Local/0] 03:38:09
                      Local via ge-0/0/3.0
20.20.0.0/24       *[Static/5] 03:16:44
                      Reject
20.20.1.0/24       *[Static/5] 03:16:44
                      Reject
20.20.2.0/24       *[Static/5] 03:16:44
                      Reject
20.20.3.0/24       *[Static/5] 03:16:44
                      Reject
172.16.1.1/32      *[Direct/0] 03:38:09
                    > via lo0.0
172.16.1.2/32      *[OSPF/10] 00:02:20, metric 1
                    > to 10.0.1.2 via ge-0/0/3.0
172.22.121.0/24    *[Direct/0] 03:38:09
                    > via ge-0/0/0.0
172.22.121.1/32    *[Local/0] 03:38:09
                      Local via ge-0/0/0.0
172.22.122.0/24    *[OSPF/10] 00:08:41, metric 2
                    > to 172.22.121.2 via ge-0/0/0.0
172.25.11.0/24     *[Direct/0] 5d 17:25:40
                    > via fxp0.0
172.25.11.1/32     *[Local/0] 5d 17:25:40
                      Local via fxp0.0
172.31.100.1/32    *[OSPF/10] 00:08:41, metric 1
                    > to 172.22.121.2 via ge-0/0/0.0
224.0.0.5/32       *[OSPF/10] 00:08:47, metric 1
                      MultiRecv

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
ff02::2/128          *[INET6/0] 5d 17:25:41
                        MultiRecv
```

Question: Do all routes show as active? Why or why not?

_____

Answer: Yes. Each route should be preceded by the asterisk (*), indicating that the route is active.

Step 2.7

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Configure mxC's interfaces and unit numbers in the correct OSPF area. Do not forget the loopback interface. Commit your configuration when completed.

```
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# edit protocols ospf

[edit]
lab@mxC# set area 0 interface lo0.0

[edit protocols ospf]
lab@mxC# set area 0 interface ge-0/0/0.0

[edit protocols ospf]
lab@mxC# set area 20 interface ge-0/0/3.0

[edit protocols ospf]
lab@mxC# commit
commit complete

[edit protocols ospf]
lab@mxC#
```

Step 2.8

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, change the perspective of the CLI to that of the R3-2 logical system by issuing the **`set cli logical-system R3-2`** command.

```
lab@mxD:R3-1> set cli logical-system R3-2
Logical system: R3-2

lab@mxD:R3-2>
```

Step 2.9

Enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Configure R3-2's interfaces and unit numbers in the correct OSPF area. Do not forget the loopback interface. Commit your configuration and exit out of configuration mode when completed.

```
lab@mxD:R3-2> configure
```

```
Entering configuration mode

[edit]
lab@mxD:R3-2# edit protocols ospf

[edit protocols ospf]
lab@mxD:R3-2# set area 20 interface ge-0/0/4.0

[edit protocols ospf]
lab@mxD:R3-2# set area 20 interface lo0.2

[edit protocols ospf]
lab@mxD:R3-2# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R3-2>
```

Step 2.10

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, verify the OSPF adjacencies with the neighboring routers by issuing the **`run show ospf neighbor`** command.

```
[edit protocols ospf]
lab@mxC# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.122.2     ge-0/0/0.0             Full      172.31.100.1     128    36
10.0.2.2         ge-0/0/3.0             Full      172.16.2.2       128    32
```

Question: Does the neighbor adjacency state show `Full` for both OSPF neighbors?

_____

Answer: Yes, the adjacency state should show `Full` for both OSPF neighbors.

Step 2.11

Issue the **`run show route`** command to look at the routing table of the mxC router.

```
[edit protocols ospf]
lab@mxC# run show route

inet.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.1.0/24        *[OSPF/10] 00:05:12, metric 3
                    > to 172.22.122.2 via ge-0/0/0.0
10.0.2.0/24        *[Direct/0] 03:43:55
                    > via ge-0/0/3.0
10.0.2.1/32        *[Local/0] 03:43:55
                       Local via ge-0/0/3.0
```

```
20.20.4.0/24        *[Static/5] 03:22:31
                       Reject
20.20.5.0/24        *[Static/5] 03:22:31
                       Reject
20.20.6.0/24        *[Static/5] 03:22:31
                       Reject
20.20.7.0/24        *[Static/5] 03:22:31
                       Reject
172.16.1.1/32       *[OSPF/10] 00:05:12, metric 2
                     > to 172.22.122.2 via ge-0/0/0.0
172.16.1.2/32       *[OSPF/10] 00:05:12, metric 3
                     > to 172.22.122.2 via ge-0/0/0.0
172.16.2.1/32       *[Direct/0] 03:43:55
                     > via lo0.0
172.16.2.2/32       *[OSPF/10] 00:00:32, metric 1
                     > to 10.0.2.2 via ge-0/0/3.0
172.22.121.0/24     *[OSPF/10] 00:05:12, metric 2
                     > to 172.22.122.2 via ge-0/0/0.0
172.22.122.0/24     *[Direct/0] 03:43:55
                     > via ge-0/0/0.0
172.22.122.1/32     *[Local/0] 03:43:55
                       Local via ge-0/0/0.0
172.25.11.0/24      *[Direct/0] 5d 18:00:54
                     > via fxp0.0
172.25.11.3/32      *[Local/0] 5d 18:00:54
                       Local via fxp0.0
172.31.100.1/32     *[OSPF/10] 00:05:12, metric 1
                     > to 172.22.122.2 via ge-0/0/0.0
224.0.0.5/32        *[OSPF/10] 00:05:18, metric 1
                       MultiRecv

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

ff02::2/128         *[INET6/0] 5d 18:00:54
                       MultiRecv
```

Question: Do all routes show as active? Why or why not?

_____

_____

Answer: Yes. Each route should be preceded by the asterisk (*), indicating that the route is active.

Step 2.12

Four static routes were defined in the mxC router's starting configuration. Issue the **`top show routing-options static`** command to view the static routes.

```
[edit protocols ospf]
lab@mxC# top show routing-options static
route 20.20.4.0/24 reject;
route 20.20.5.0/24 reject;
route 20.20.6.0/24 reject;
```

```
route 20.20.7.0/24 reject;

[edit protocols ospf]
lab@mxC#
```

Question: Do the configured static routes match the external routes listed on the lab diagram?

_____

Answer: Yes, the configured routes should be the same as the external routes on the lab diagram.

Step 2.13

A policy to redistribute the static routes into OSPF has been preconfigured for you. Issue the **`top show policy-options`** command to view the policy.

```
[edit protocols ospf]
lab@mxC# top show policy-options
policy-statement static-to-ospf {
    term 1 {
        from protocol static;
        then accept;
    }
}

[edit protocols ospf]
lab@mxC#
```

Question: How can you apply this policy so that any configured static routes will be redistributed into OSPF?

_____

_____

Answer: To redistribute the static routes into OSPF, you must apply the *`static-to-ospf`* policy as an export policy to the mxC router's OSPF neighbors.

Step 2.14

Apply the *`static-to-ospf`* policy as an export policy to export the static routes into OSPF and commit your configuration.

```
[edit protocols ospf]
lab@mxC# show
area 0.0.0.0 {
    interface lo0.0;
    interface ge-0/0/0.0;
}
area 0.0.0.20 {
    interface ge-0/0/3.0;
}

[edit protocols ospf]
```

```
lab@mxC# set export static-to-ospf

[edit protocols ospf]
lab@mxC# commit
commit complete

[edit protocols ospf]
lab@mxC#
```

Step 2.15

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, issue the **`top show routing-options static`** command to view the preconfigured static routes.

```
[edit protocols ospf]
lab@mxA# top show routing-options static
route 20.20.0.0/24 reject;
route 20.20.1.0/24 reject;
route 20.20.2.0/24 reject;
route 20.20.3.0/24 reject;

[edit protocols ospf]
lab@mxA#
```

Question: Do the configured static routes match the external routes listed on the lab diagram?

_____

Answer: Yes, the configured routes should be the same as the external routes on the lab diagram.

Step 2.16

A policy to redistribute the static routes into OSPF has been preconfigured for you. Issue the **`top show policy-options`** command to view the policy.

```
[edit protocols ospf]
lab@mxA# top show policy-options
policy-statement static-to-ospf {
    term 1 {
        from protocol static;
        then accept;
    }
}

[edit protocols ospf]
lab@mxA#
```

Question: How can you apply this policy so that any configured static routes will be redistributed into OSPF?

_____

_____

Answer: To redistribute the static routes into OSPF, you must apply the *static-to-ospf* policy as an export policy to the mxC router's OSPF neighbors.

Step 2.17

Apply the *static-to-ospf* policy as an export policy to export the static routes into OSPF and commit your configuration.

```
[edit protocols ospf]
lab@mxA# show
area 0.0.0.0 {
    interface lo0.0;
    interface ge-0/0/0.0;
}
area 0.0.0.10 {
    interface ge-0/0/3.0;
}

[edit protocols ospf]
lab@mxA# set export static-to-ospf

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 2.18

Use the **run show ospf database** command to examine the link-state database (LSDB) on the mxA router, which is the ABR and an ASBR. Notice that the output is organized by areas, Area 0.0.0.0 first and then the nonbackbone areas in numerical order, followed by the external routes (Type 5 exported static routes) labeled as the "OSPF AS SCOPE link state database".

Notice the Router LSAs (Type 1) which represent each router in an Area. Network LSAs (Type 2) are the Gigabit Ethernet links. The Summary LSAs (Type 3) contain the contents of either Router LSAs or Network LSAs converted by the ABR and injected into the other area. The ASBRSum LSA (Type 4) are listed in the non-backbone area.

```
[edit protocols ospf]
lab@mxA# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   *172.16.1.1       172.16.1.1       0x80000005    85  0x22 0xcf94  48
Router    172.16.2.1       172.16.2.1       0x80000004   317  0x22 0xea75  48
Router    172.31.100.1     172.31.100.1     0x8000000a  1349  0x22 0x5c1f  60
Network   172.22.121.2     172.31.100.1     0x80000002   255  0x22 0x3e9b  32
```

```
Network  172.22.122.2      172.31.100.1      0x80000001  1349  0x22 0x4296  32
Summary *10.0.1.0          172.16.1.1        0x80000003  1682  0x22 0x620e  28
Summary  10.0.2.0          172.16.2.1        0x80000003  1069  0x22 0x501e  28
Summary *172.16.1.2        172.16.1.1        0x80000001  1675  0x22 0x4f6e  28
Summary  172.16.2.2        172.16.2.1        0x80000001  1064  0x22 0x3d7e  28

    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Router  *172.16.1.1        172.16.1.1        0x80000006    85  0x22 0x9302  36
Router   172.16.1.2        172.16.1.2        0x80000003  1643  0x22 0x5177  48
Network *10.0.1.1          172.16.1.1        0x80000001  1682  0x22 0xae6   32
Summary *10.0.2.0          172.16.1.1        0x80000001  1341  0x22 0x6fff  28
Summary *172.16.1.1        172.16.1.1        0x80000003  1152  0x22 0x4b72  28
Summary *172.16.2.1        172.16.1.1        0x80000001  1341  0x22 0x5864  28
Summary *172.16.2.2        172.16.1.1        0x80000001  1060  0x22 0x5862  28
Summary *172.22.121.0      172.16.1.1        0x80000004   612  0x22 0xe756  28
Summary *172.22.122.0      172.16.1.1        0x80000002  1348  0x22 0xea53  28
Summary *172.31.100.1      172.16.1.1        0x80000002   871  0x22 0x5ded  28
ASBRSum *172.16.2.1        172.16.1.1        0x80000001   315  0x22 0x4a71  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq       Age  Opt  Cksum  Len
Extern  *20.20.0.0         172.16.1.1        0x80000001    85  0x22 0x6b60  36
Extern  *20.20.1.0         172.16.1.1        0x80000001    85  0x22 0x606a  36
Extern  *20.20.2.0         172.16.1.1        0x80000001    85  0x22 0x5574  36
Extern  *20.20.3.0         172.16.1.1        0x80000001    85  0x22 0x4a7e  36
Extern   20.20.4.0         172.16.2.1        0x80000001   317  0x22 0x388e  36
Extern   20.20.5.0         172.16.2.1        0x80000001   317  0x22 0x2d98  36
Extern   20.20.6.0         172.16.2.1        0x80000001   317  0x22 0x22a2  36
Extern   20.20.7.0         172.16.2.1        0x80000001   317  0x22 0x17ac  36
```

Question: How many and what types of link-state advertisements (LSAs) exist in OSPF database for Area 0?

_____

_____

Answer: You should see three Router, two Network and four Summary LSAs for Area 0.

Part 3: Configuring the OSPF Link Costs

In this lab part, you configure the OSPF link costs.

Step 3.1

Issue the **`run show ospf interface detail`** command and answer the following question.

```
[edit protocols ospf]
lab@mxA# run show ospf interface detail
Interface            State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0           BDR     0.0.0.0         172.31.100.1    172.16.1.1         1
  Type: LAN, Address: 172.22.121.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  DR addr: 172.22.121.2, BDR addr: 172.22.121.1, Priority: 128
  Adj count: 1
```

```
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
lo0.0               DR      0.0.0.0         172.16.1.1      0.0.0.0             0
  Type: LAN, Address: 172.16.1.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  DR addr: 172.16.1.1, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 0
ge-0/0/3.0          DR      0.0.0.10        172.16.1.1      172.16.1.2          1
  Type: LAN, Address: 10.0.1.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  DR addr: 10.0.1.1, BDR addr: 10.0.1.2, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
```

Question: What are the current costs for the Gigabit Ethernet links and why?

_____

_____

Answer: Both of the Gigabit Ethernet links should have the same cost of 1, because at the default `reference-bandwidth` setting of 100 Mbps, the calculation of a Gigabit Ethernet link is actually lower than one and must be rounded up to the nearest integer according to the RFC.

Step 3.2

At this point, we will better represent the link bandwidths in the network. Using the **`reference-bandwidth`** command, alter the metric calculation such that the bandwidth of a 10-Gigabit Ethernet link becomes the basis for the formula. Commit your configuration when completed.

```
[edit protocols ospf]
lab@mxA# set reference-bandwidth 10g

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 3.3

Issue the **`run show ospf interface detail`** command to see if the link costs changed.

```
[edit protocols ospf]
lab@mxA# run show ospf interface detail
Interface            State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0           BDR     0.0.0.0         172.31.100.1    172.16.1.1         1
  Type: LAN, Address: 172.22.121.1, Mask: 255.255.255.0, MTU: 1500, Cost: 10
  DR addr: 172.22.121.2, BDR addr: 172.22.121.1, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 10
lo0.0                DR      0.0.0.0         172.16.1.1      0.0.0.0            0
  Type: LAN, Address: 172.16.1.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  DR addr: 172.16.1.1, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 0
ge-0/0/3.0           DR      0.0.0.10        172.16.1.1      172.16.1.2         1
  Type: LAN, Address: 10.0.1.1, Mask: 255.255.255.0, MTU: 1500, Cost: 10
  DR addr: 10.0.1.1, BDR addr: 10.0.1.2, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 10
```

Question: Did the cost of the links in the network change? If so, what are some of the new costs?

_____

_____

Answer: If you used a reference bandwidth value of 10 Gbps or higher, the Gigabit Ethernet link costs should have changed. If you used a reference bandwidth value of 10 Gbps, the Gigabit Ethernet link costs should be 10.

Step 3.4

Change the metric on your Area 0 loopback interface to be ***10000*** and commit your configuration.

```
[edit protocols ospf]
lab@mxA# set area 0 interface lo0.0 metric 10000

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 3.5

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxD:R3-2> set cli logical-system R3-1
Logical system: R3-1

lab@mxD:R3-1>
```

Step 3.6

Enter the **`show route 172.16/16`** command and examine the routing table on your student device at this point.

```
lab@mxD:R3-1> show route 172.16/16

inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.1/32       *[OSPF/10] 00:03:08, metric 10001
                     > to 10.0.1.1 via ge-0/0/1.0
172.16.1.2/32       *[Direct/0] 2d 01:17:59
                     > via lo0.1
172.16.2.1/32       *[OSPF/10] 00:04:15, metric 12
                     > to 10.0.1.1 via ge-0/0/1.0
172.16.2.2/32       *[OSPF/10] 00:04:15, metric 13
                     > to 10.0.1.1 via ge-0/0/1.0
```

Question: Did any of the metric costs increase as a result of this configuration change? If so, which routes changed metric values?

_____

_____

Answer: Yes. The metric cost changed for routes to all of the other routers' loopback addresses.

Note

If no metric values have changed on your router, then **STOP** until some networks have changed. If, however, some networks in your routing table do have increased metrics, then proceed to the next step.

Step 3.7

You will now alter the metrics for the interfaces in the OSPF area 10. The mxA router should use a metric of *5000*, and the R3-1 router should use a metric of *2500*.

Enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Configure the ge-0/0/1 interface with a metric of *2500*. Commit your configuration when completed.

```
lab@mxD:R3-1> configure
Entering configuration mode
```

```
[edit]
lab@mxD:R3-1# edit protocols ospf

[edit protocols ospf]
lab@mxD:R3-1# set area 10 interface ge-0/0/1.0 metric 2500

[edit protocols ospf]
lab@mxD:R3-1# commit
commit complete

[edit protocols ospf]
lab@mxD:R3-1#
```

Step 3.8

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, configure the ge-0/0/3 interface with a metric of *5000*. Commit your configuration when completed.

```
[edit protocols ospf]
lab@mxA# set area 10 interface ge-0/0/3.0 metric 5000

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 3.9

Examine the routing table of mxA. Specifically, use the **`run show route 172.16/16`** command to look at the cost to reach the loopback address of the R3-1 router on the other end of the link.

```
[edit protocols ospf]
lab@mxA# run show route 172.16/16

inet.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.1/32       *[Direct/0] 2d 01:32:01
                     > via lo0.0
172.16.1.2/32       *[OSPF/10] 00:00:51, metric 5000
                     > to 10.0.1.2 via ge-0/0/3.0
172.16.2.1/32       *[OSPF/10] 00:13:32, metric 11
                     > to 172.22.121.2 via ge-0/0/0.0
172.16.2.2/32       *[OSPF/10] 00:13:32, metric 12
                     > to 172.22.121.2 via ge-0/0/0.0
```

Question: What is the OSPF cost to reach the loopback of the R3-1 router?

_____

_____

Answer: The cost is 5000 to reach the loopback of the R3-1 router from mxA's perspective.

Step 3.10

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, examine the routing tables of R3-1. Specifically, use the **`run show route 172.16/16`** command to look at the cost to reach the loopback address of the mxA router on the other end of the link.

```
[edit protocols ospf]
lab@mxD:R3-1# run show route 172.16/16

inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.1/32       *[OSPF/10] 00:05:10, metric 12500
                     > to 10.0.1.1 via ge-0/0/1.0
172.16.1.2/32       *[Direct/0] 2d 01:29:51
                     > via lo0.1
172.16.2.1/32       *[OSPF/10] 00:05:10, metric 2511
                     > to 10.0.1.1 via ge-0/0/1.0
172.16.2.2/32       *[OSPF/10] 00:05:10, metric 2512
                     > to 10.0.1.1 via ge-0/0/1.0
```

Question: What is the OSPF cost to reach the loopback of the mxA router?

_____

_____

Answer: The cost is 12500 (10000 + 2500) to reach the loopback of the mxA router from R3-1's perspective.

Part 4: Configuring Overload

In this lab part, you configure the R3-1 router to be in overload mode.

Step 4.1

Enter the **`run show route`** to examine the routing table for the R3-1 router and look at the metric for the transit links. You should also see the 20.20/16 static routes.

```
[edit protocols ospf]
lab@mxD:R3-1# run show route

inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.1.0/24        *[Direct/0] 2d 01:33:04
                    > via ge-0/0/1.0
10.0.1.2/32        *[Local/0] 2d 01:33:04
                      Local via ge-0/0/1.0
10.0.2.0/24        *[OSPF/10] 00:08:23, metric 2512
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.0.0/24       *[OSPF/150] 1d 21:35:30, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.1.0/24       *[OSPF/150] 1d 21:35:30, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.2.0/24       *[OSPF/150] 1d 21:35:30, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.3.0/24       *[OSPF/150] 1d 21:35:30, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.4.0/24       *[OSPF/150] 1d 21:39:20, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.5.0/24       *[OSPF/150] 1d 21:39:20, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.6.0/24       *[OSPF/150] 1d 21:39:20, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.7.0/24       *[OSPF/150] 1d 21:39:20, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
172.16.1.1/32      *[OSPF/10] 00:08:23, metric 12500
                    > to 10.0.1.1 via ge-0/0/1.0
172.16.1.2/32      *[Direct/0] 2d 01:33:04
                    > via lo0.1
172.16.2.1/32      *[OSPF/10] 00:08:23, metric 2511
                    > to 10.0.1.1 via ge-0/0/1.0
172.16.2.2/32      *[OSPF/10] 00:08:23, metric 2512
                    > to 10.0.1.1 via ge-0/0/1.0
172.22.121.0/24    *[OSPF/10] 00:08:23, metric 2510
                    > to 10.0.1.1 via ge-0/0/1.0
172.22.122.0/24    *[OSPF/10] 00:08:23, metric 2511
                    > to 10.0.1.1 via ge-0/0/1.0
172.31.100.1/32    *[OSPF/10] 00:08:23, metric 2510
                    > to 10.0.1.1 via ge-0/0/1.0
224.0.0.5/32       *[OSPF/10] 1d 22:02:07, metric 1
                      MultiRecv

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

ff02::2/128        *[INET6/0] 2d 01:33:05
                      MultiRecv
```

Question: What is the cost that is being used on the R3-1 router's ge-0/0/1 to help calculate the total cost to reach remote networks?

_____

Answer: The R3-1 router is using a cost of 2500 on its ge-0/0/1.0 interface.

Step 4.2

Configure the R3-1 router to be in overload mode and commit your configuration when completed.

```
[edit protocols ospf]
lab@mxD:R3-1# set overload

[edit protocols ospf]
lab@mxD:R3-1# commit
commit complete

[edit protocols ospf]
lab@mxD:R3-1#
```

Step 4.3

Enter the **run show route** command to examine R3-1's route table again.

```
[edit protocols ospf]
lab@mxD:R3-1# run show route

inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.1.0/24        *[Direct/0] 2d 01:38:16
                    > via ge-0/0/1.0
10.0.1.2/32        *[Local/0] 2d 01:38:16
                      Local via ge-0/0/1.0
10.0.2.0/24        *[OSPF/10] 00:00:39, metric 65547
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.0.0/24       *[OSPF/150] 1d 21:40:42, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.1.0/24       *[OSPF/150] 1d 21:40:42, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.2.0/24       *[OSPF/150] 1d 21:40:42, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.3.0/24       *[OSPF/150] 1d 21:40:42, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.4.0/24       *[OSPF/150] 1d 21:44:32, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.5.0/24       *[OSPF/150] 1d 21:44:32, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.6.0/24       *[OSPF/150] 1d 21:44:32, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.7.0/24       *[OSPF/150] 1d 21:44:32, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
172.16.1.1/32      *[OSPF/10] 00:00:39, metric 75535
                    > to 10.0.1.1 via ge-0/0/1.0
172.16.1.2/32      *[Direct/0] 2d 01:38:16
                    > via lo0.1
172.16.2.1/32      *[OSPF/10] 00:00:39, metric 65546
                    > to 10.0.1.1 via ge-0/0/1.0
172.16.2.2/32      *[OSPF/10] 00:00:39, metric 65547
                    > to 10.0.1.1 via ge-0/0/1.0
172.22.121.0/24    *[OSPF/10] 00:00:39, metric 65545
                    > to 10.0.1.1 via ge-0/0/1.0
```

```
172.22.122.0/24     *[OSPF/10] 00:00:39, metric 65546
                     > to 10.0.1.1 via ge-0/0/1.0
172.31.100.1/32     *[OSPF/10] 00:00:39, metric 65545
                     > to 10.0.1.1 via ge-0/0/1.0
224.0.0.5/32        *[OSPF/10] 1d 22:07:19, metric 1
                       MultiRecv

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

ff02::2/128         *[INET6/0] 2d 01:38:17
                       MultiRecv

[edit protocols ospf]
lab@mxD:R3-1#
```

Question: Did the metrics change?

_____

_____

Answer: Yes. The metric should be greater than 65535.

Question: Are all OSPF neighbors still fully adjacent?

_____

_____

Answer: Yes. Configuring a router as overloaded does not flap the adjacencies.

Part 5: Performing Authentication

In this lab part, you perform authentication in OSPF area 10 between the mxA router and the R3-1 router.

Step 5.1

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, configure OSPF traceoptions such that all detailed OSPF logs are sent to a file name *ospf.log*.

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

[edit protocols ospf]
lab@mxA# set traceoptions flag all detail
```

Step 5.2

Configure area 10 to use an authentication key of *juniper* using the Message Digest 5 (MD5) algorithm. Use a key-id of *10* and commit your configuration when completed.

```
[edit protocols ospf]
lab@mxA# set area 10 interface ge-0/0/3.0 authentication md5 10 key juniper
```

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

[edit protocols ospf]
lab@mxA#
```

Step 5.3

Use the **run show ospf neighbor** command a few times to verify the OSPF neighbor state is no longer in a `Full` state. Note that, because of the dead timer, it might take up to forty seconds for the neighbor to disappear from the output.

```
[edit protocols ospf]
lab@mxA# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.121.2     ge-0/0/0.0             Full      172.31.100.1     128    31
10.0.1.2         ge-0/0/3.0             Full      172.16.1.2       128     0

[edit protocols ospf]
lab@mxA# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.121.2     ge-0/0/0.0             Full      172.31.100.1     128    34

[edit protocols ospf]
lab@mxA#
```

Step 5.4

Issue the **run show log ospf.log | match auth** command.

```
[edit protocols ospf]
lab@mxA# run show log ospf.log | match auth
Apr 26 14:49:27.902213 OSPF packet ignored: authentication type mismatch (0) from
   10.0.1.2
Apr 26 14:49:35.557288 OSPF packet ignored: authentication type mismatch (0) from
   10.0.1.2
Apr 26 14:49:43.231995 OSPF packet ignored: authentication type mismatch (0) from
   10.0.1.2
Apr 26 14:49:51.946490 OSPF packet ignored: authentication type mismatch (0) from
   10.0.1.2
Apr 26 14:49:59.672812 OSPF packet ignored: authentication type mismatch (0) from
   10.0.1.2
Apr 26 14:50:08.668801 OSPF packet ignored: authentication type mismatch (0) from
   10.0.1.2
Apr 26 14:50:17.483160 OSPF packet ignored: authentication type mismatch (0) from
   10.0.1.2
Apr 26 14:50:25.947841 OSPF packet ignored: authentication type mismatch (0) from
   10.0.1.2
Apr 26 14:50:33.696803 OSPF packet ignored: authentication type mismatch (0) from
   10.0.1.2

[edit protocols ospf]
lab@mxA#
```

Question: Are there any issues related to the neighbor relationship with the R3-1 router? Is this expected?

_____

_____

Answer: Yes, there is a problem. All OSPF packets from the R3-1 router are being ignored because of an authentication mismatch. This is the expected behavior.

Step 5.5

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, configure your R3-1 router to support an authentication key of *juniper* using the Message Digest 5 (MD5) algorithm. Use a key-id of *10*, commit your configuration and exit to operational mode when completed.

```
[edit protocols ospf]
lab@mxD:R3-1# set area 10 interface ge-0/0/1.0 authentication md5 10 key juniper

[edit protocols ospf]
lab@mxD:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R3-1>
```

Step 5.6

Issue the **show ospf neighbor** command to verify that the neighbor adjacencies have returned to the Full state.

```
lab@mxD:R3-1> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.0.1.1         ge-0/0/1.0             Full      172.16.1.1       128    39

lab@mxD:R3-1>
```

Question: Has the neighbor relationship between R3-1 and mxA been restored to a full state?

_____

Answer: Yes, the neighbor relationship has been restored to a full state.

Question: What happens if one router enters a higher key-id value while using the same password?

_____

_____

Answer: The highest key value is used by default. However, the key-id values must match in order to form a Full adjacency.

Step 5.7

Log out of each of your assigned devices using the **exit** command.

```
lab@mxD:R3-1> exit

mxD (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

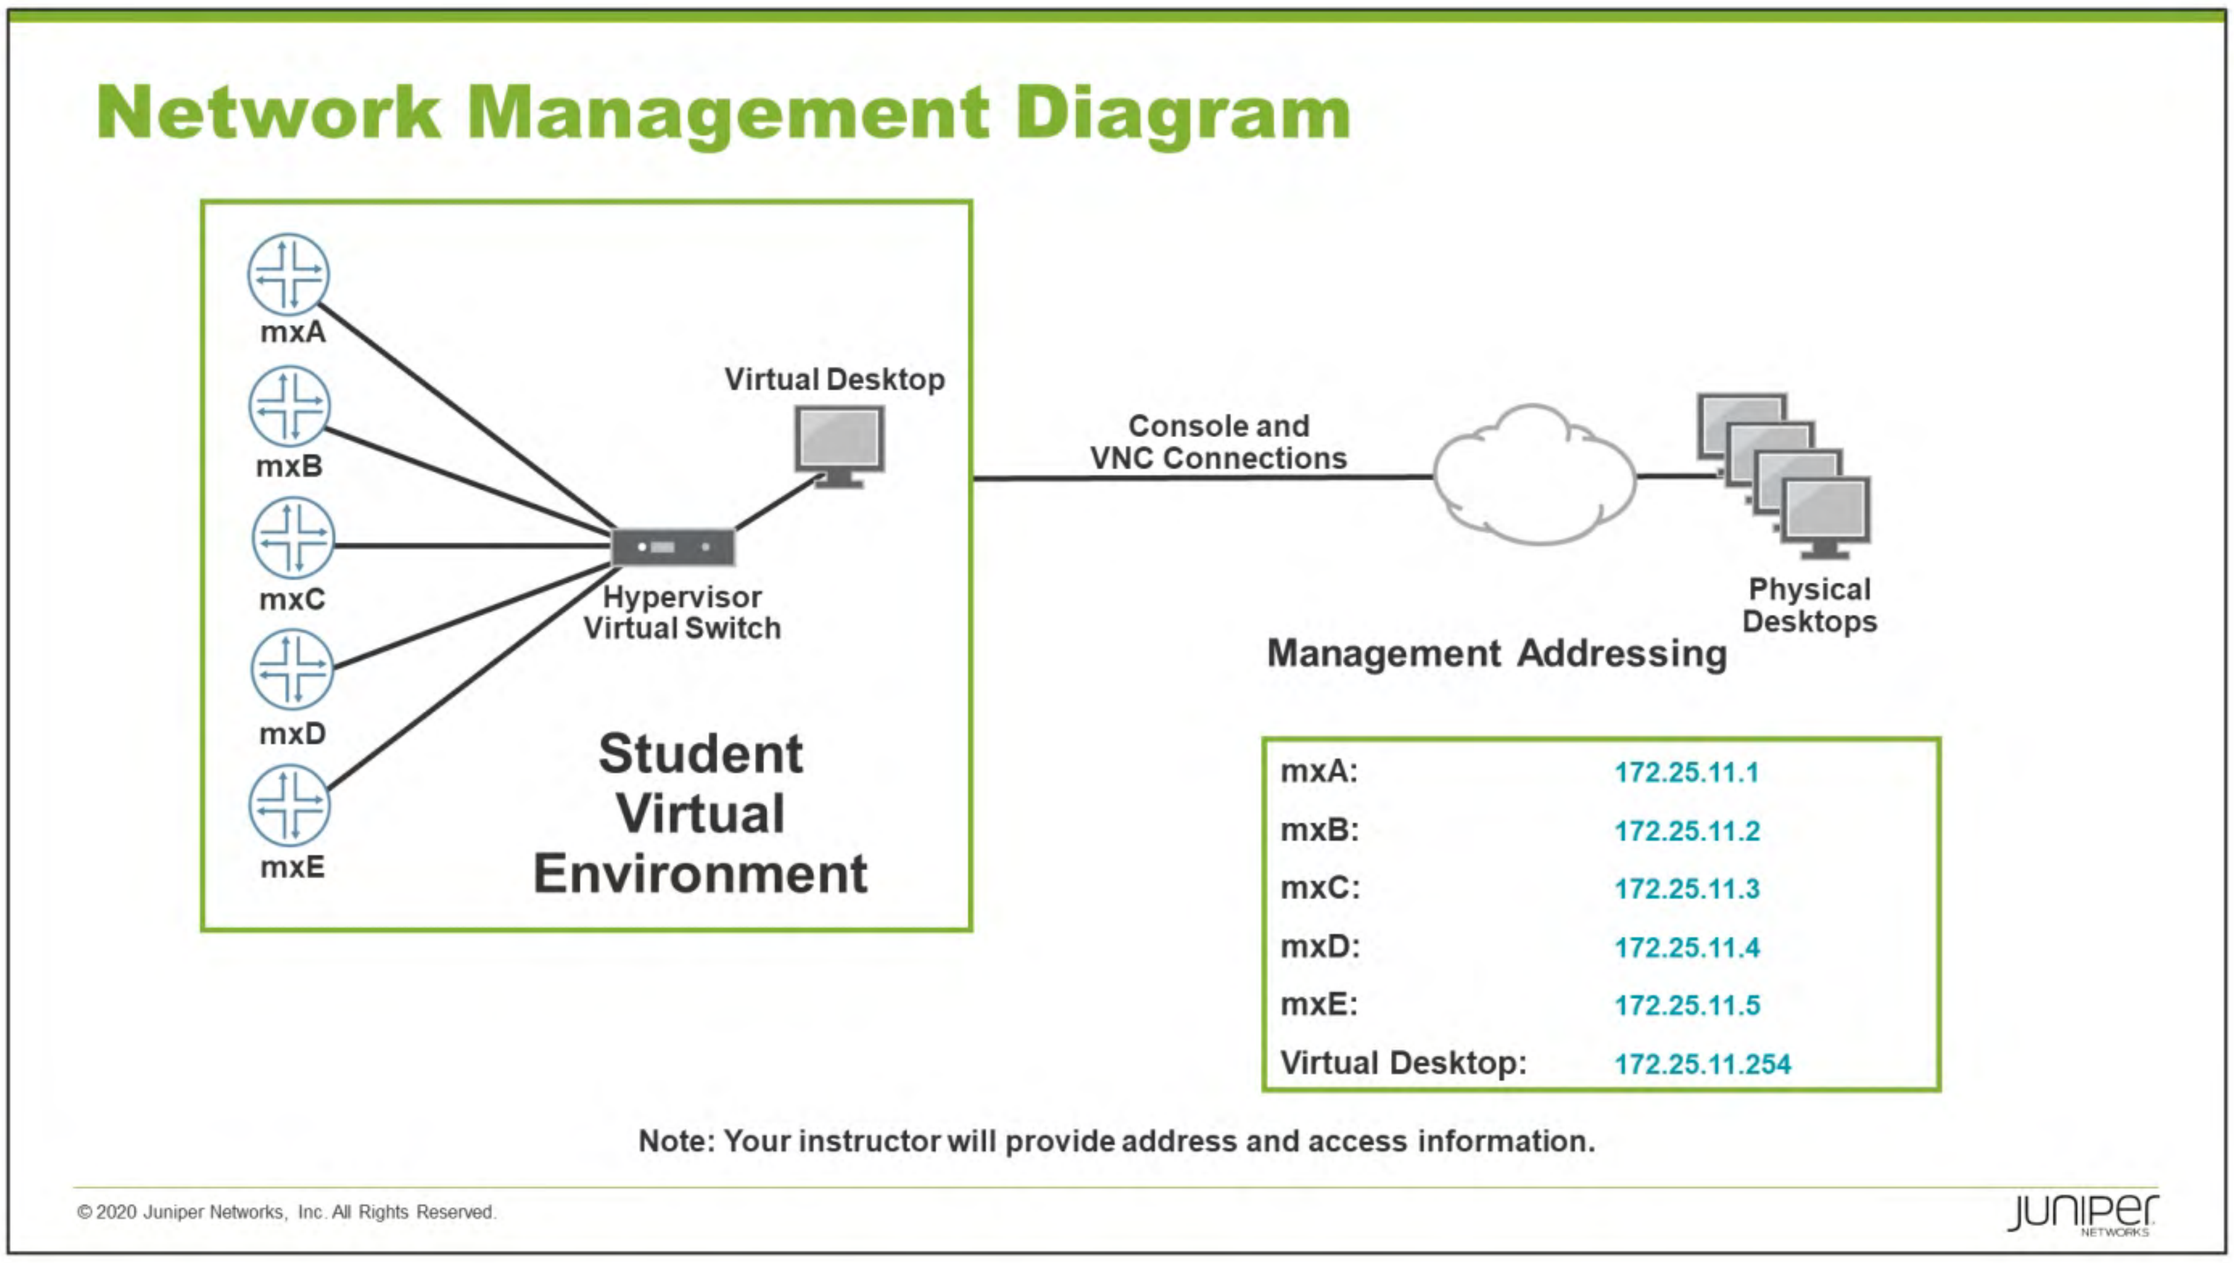
Network Management Diagram
mxA
mxB
mxC
mxD
mxE
Virtual Desktop
Hypervisor
Virtual Switch
Student
Virtual
Environment
Console and
VNC Connections
Physical
Desktops
Management Addressing
mxA: 172.25.11.1
mxB: 172.25.11.2
mxC: 172.25.11.3
mxD: 172.25.11.4
mxE: 172.25.11.5
Virtual Desktop: 172.25.11.254
Note: Your instructor will provide address and access information.
© 2020 Juniper Networks, Inc. All Rights Reserved.
juniper NETWORKS

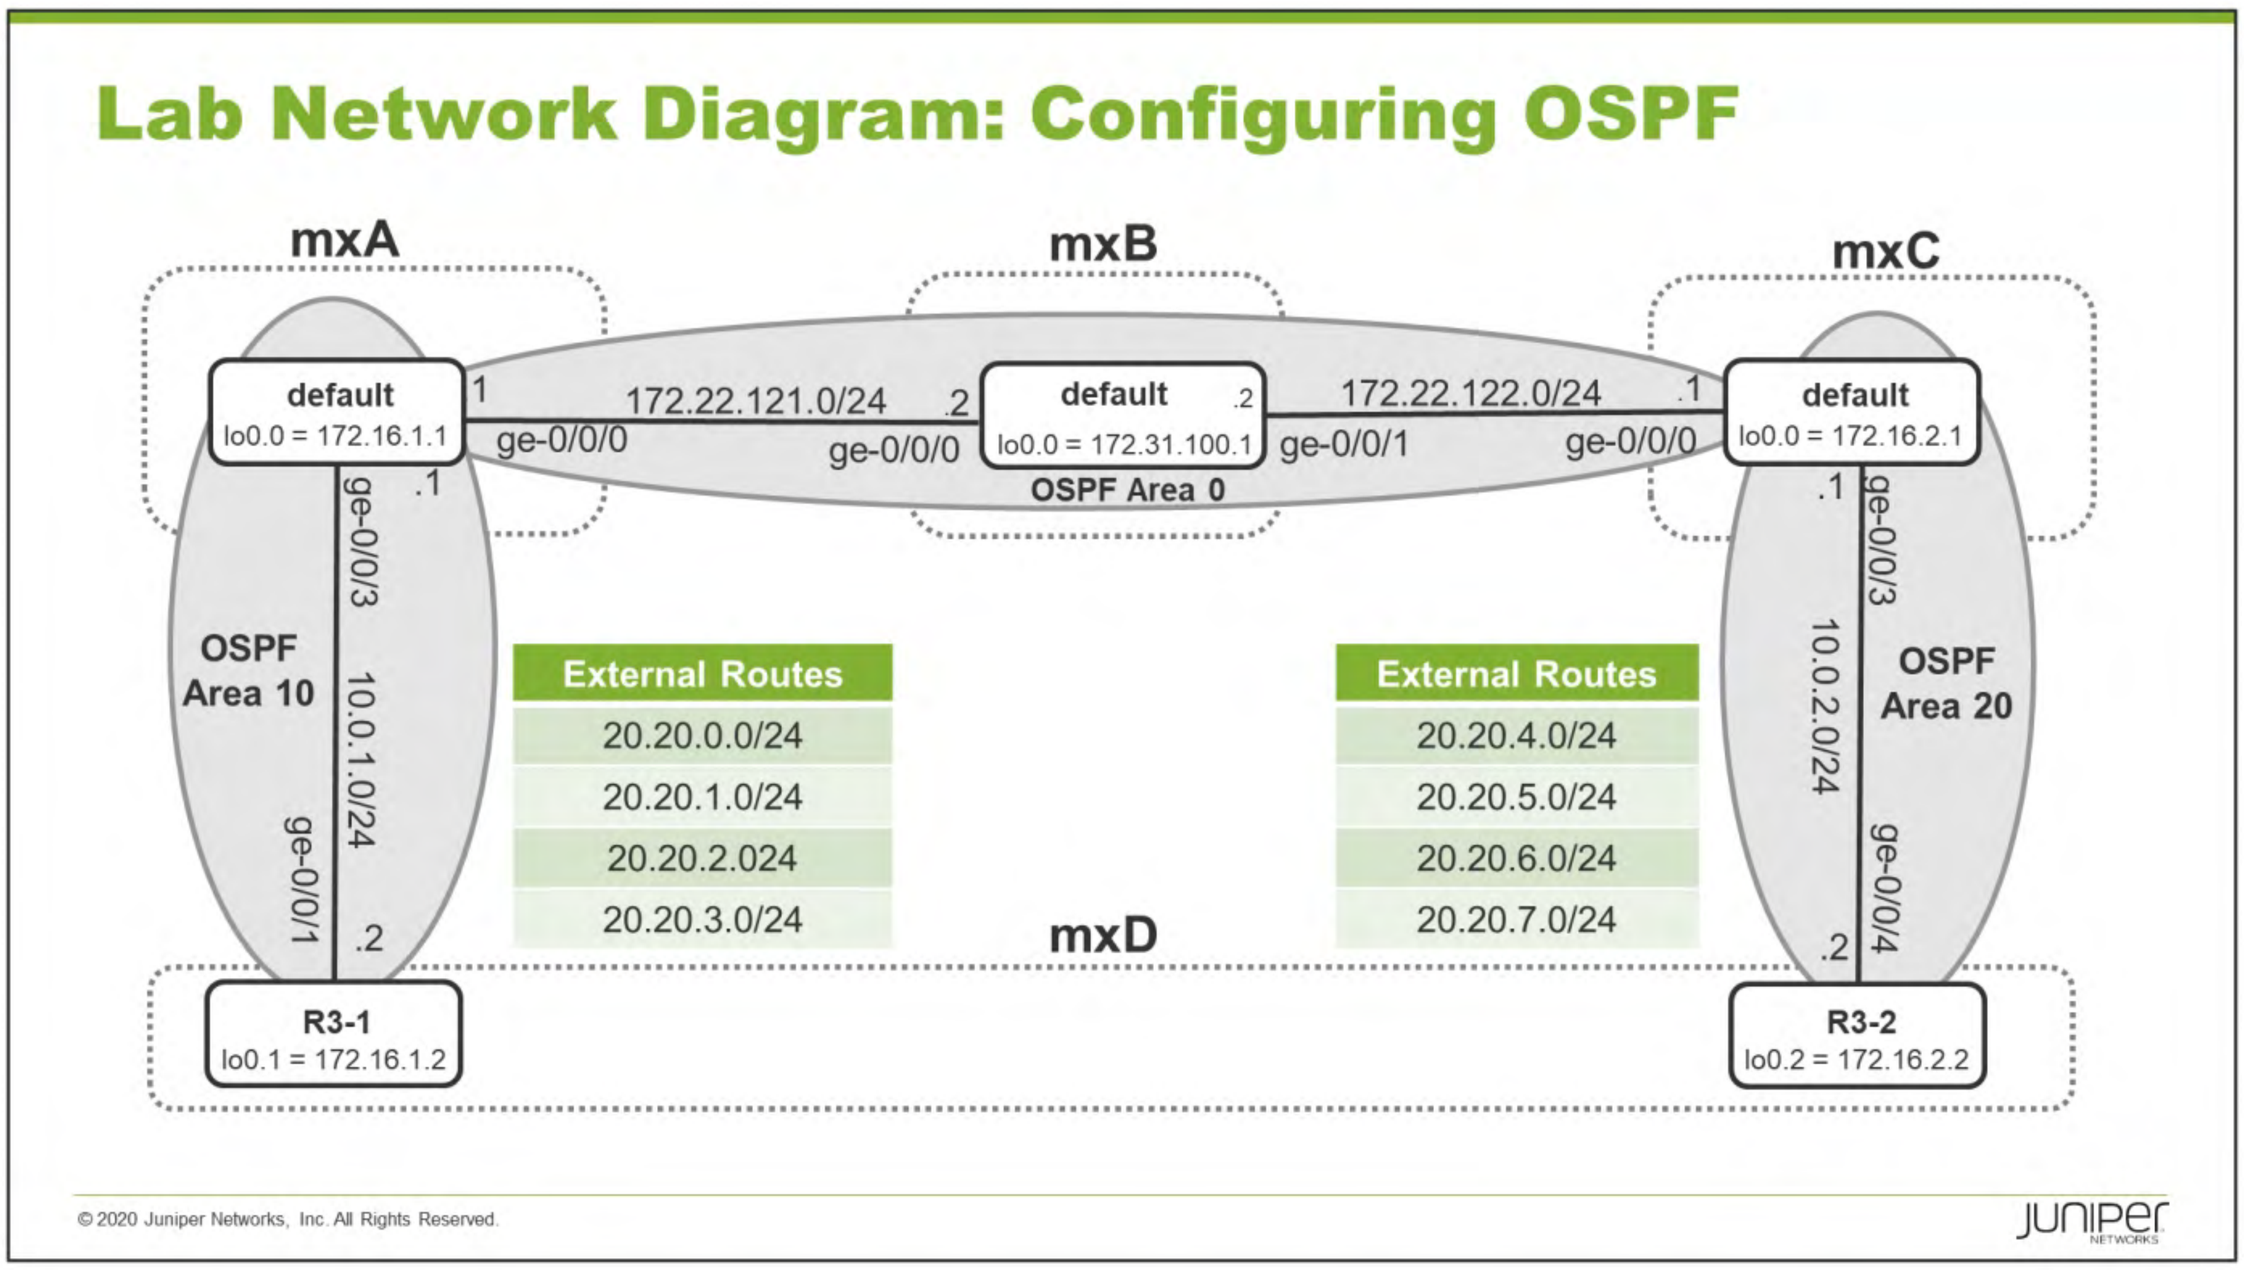
Lab Network Diagram: Configuring OSPF
mxA
mxB
mxC
mxD
default
lo0.0 = 172.16.1.1
default
lo0.0 = 172.31.100.1
default
lo0.0 = 172.16.2.1
172.22.121.0/24
172.22.122.0/24
ge-0/0/0
ge-0/0/0
ge-0/0/1
ge-0/0/0
OSPF Area 0
ge-0/0/3
10.0.1.0/24
ge-0/0/1
OSPF
Area 10
ge-0/0/3
10.0.2.0/24
ge-0/0/4
OSPF
Area 20
External Routes
20.20.0.0/24
20.20.1.0/24
20.20.2.024
20.20.3.0/24
External Routes
20.20.4.0/24
20.20.5.0/24
20.20.6.0/24
20.20.7.0/24
R3-1
lo0.1 = 172.16.1.2
R3-2
lo0.2 = 172.16.2.2
© 2020 Juniper Networks, Inc. All Rights Reserved.
juniper NETWORKS

Lab

Advanced OSPF

Overview

In this lab, you use the lab diagram titled “Lab Network Diagram: Advanced OSPF” to establish a multiarea OSPF routing domain. You configure and monitor OSPF areas and route summarization to convert some of the areas in the OSPF routing domain into OSPF stub areas. You then convert them into stub no-summaries areas. Finally, you convert the areas into OSPF NSSA areas, as well as NSSA with no summaries areas.

By completing this lab, you perform the following tasks:

- Create a OSPF stub area.
- Create a OSPF stub no-summaries area.
- Create a OSPF not-so-stubby area.
- Create a OSPF not-so-stubby no-summaries area.

Part 1: Load Starting Configuration

In this lab part, you verify the initial configuration of the routers. You then verify that the interfaces are operational and that OSPF neighbor relationships have formed.

Note

The instructor will tell you the nature of your access and will provide you with the necessary details to access your assigned device.

Note

The lab topology requires you to display information in the different logical systems of mxD. You will be changing the perspective of the CLI by issuing the **`set cli logical-system`** ***`ls-name`*** command. Refer to the lab diagram for the correct logical system name.

By changing the perspective of the CLI, you will be able to perform network commands such as **`ping`** or **`traceroute`** from the perspective of the pertinent logical system.

Step 1.1

Ensure that you know to which student device you have been assigned. Consult the management network diagram to determine the management address of the student devices.

Access the CLI on your mxD device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab2-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.4
Password: lab123
Last login: Tue Nov 19 18:31:29 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab2-start.config
load complete

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

lab@mxD>
```

Step 1.2

Issue the **show configuration** command. Use the lab diagram to verify that the mxD router has two logical systems, R3-1 and R3-2, that have the correct interface and OSPF configuration.

```
lab@mxD> show configuration
## Last commit: 2019-11-20 11:24:42 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxD;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    R3-1 {
        interfaces {
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 10.0.1.2/24;
                    }
                }
            }
            lo0 {
                unit 1 {
```

```
                    family inet {
                        address 172.16.1.2/32;
                    }
                }
            }
        }
        protocols {
            ospf {
                area 0.0.0.10 {
                    interface ge-0/0/1.0;
                    interface lo0.1;
                }
            }
        }
        routing-options {
            autonomous-system 65512;
        }
    }
    R3-2 {
        interfaces {
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 10.0.2.2/24;
                    }
                }
            }
            lo0 {
                unit 2 {
                    family inet {
                        address 172.16.2.2/32;
                    }
                }
            }
        }
        protocols {
            ospf {
                area 0.0.0.20 {
                    interface ge-0/0/4.0;
                    interface lo0.2;
                }
            }
        }
        routing-options {
            autonomous-system 65512;
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.4/24;
            }
        }
```

```
        }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications. If the configuration is incorrect, please notify your instructor.

Step 1.3

Access the CLI on your mxA device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab2-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.1
Password: lab123
Last login: Tue Nov 19 18:34:38 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# load override ajspr/lab2-start.config
load complete

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

lab@mxA>
```

Step 1.4

Issue the **show configuration** command. Use the lab diagram to verify that the mxA router has the correct interface and OSPF configuration. Verify that there are four 20.20/24 static routes and a policy statement.

```
lab@mxA> show configuration
## Last commit: 2019-11-20 11:26:46 UTC by lab
version 20190829.221548_builder.r1052644;
## Last commit: 2019-11-20 11:26:46 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxA;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
```

```
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.22.121.1/24;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.1.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.1/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.1.1/32;
```

```
                }
            }
        }
    }
    policy-options {
        policy-statement static-to-ospf {
            term 1 {
                from protocol static;
                then accept;
            }
        }
    }
    routing-options {
        static {
            route 20.20.0.0/24 reject;
            route 20.20.1.0/24 reject;
            route 20.20.2.0/24 reject;
            route 20.20.3.0/24 reject;
        }
        autonomous-system 65512;
    }
    protocols {
        ospf {
            area 0.0.0.0 {
                interface lo0.0;
                interface ge-0/0/0.0;
            }
            area 0.0.0.10 {
                interface ge-0/0/3.0;
            }
            export static-to-ospf;
        }
    }
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxA router. If the configuration is incorrect, please notify your instructor.

Question: An export policy has been applied to the OSPF configuration. What does it accomplish?

_____

_____

Answer: The statement `export static-to-ospf` causes all of the locally configured and active static routes to be advertised as external OSPF routes (using Type 5 LSAs).

Step 1.5

Use the **ping 10.0.1.2 count 5** command to verify that you can ping the physical interface on the R3-1 router.

```
lab@mxA> ping 10.0.1.2 count 5
PING 10.0.1.2 (10.0.1.2): 56 data bytes
64 bytes from 10.0.1.2: icmp_seq=0 ttl=64 time=4.347 ms
64 bytes from 10.0.1.2: icmp_seq=1 ttl=64 time=1.434 ms
64 bytes from 10.0.1.2: icmp_seq=2 ttl=64 time=18.736 ms
64 bytes from 10.0.1.2: icmp_seq=3 ttl=64 time=1.416 ms
64 bytes from 10.0.1.2: icmp_seq=4 ttl=64 time=1.491 ms

--- 10.0.1.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.416/5.485/18.736/6.720 ms

lab@mxA>
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.6

Issue the **show ospf neighbor 10.0.1.2** command to verify that the mxA and R3-1 routers have formed a neighbor relationship in Area 10.

```
lab@mxA> show ospf neighbor 10.0.1.2
Address          Interface              State     ID               Pri  Dead
10.0.1.2         ge-0/0/3.0             Full      172.16.1.2       128    37
```

Question: Have the mxA and R3-1 router formed a neighbor relationship with one another?

_____

Answer: Yes, the two routers are now OSPF neighbors.

Step 1.7

Issue the **show ospf database area *10* advertising-router 172.16.1.1** command to verify that the mxA router is properly redistributing its static routes into OSPF.

```
lab@mxA> show ospf database area 10 advertising-router 172.16.1.1

    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   *172.16.1.1       172.16.1.1       0x8000005a  2681  0x22 0xf44b  36
Summary *10.0.2.0         172.16.1.1       0x80000054  2493  0x22 0xc853  28
Summary *172.16.1.1       172.16.1.1       0x80000055  1181  0x22 0xa6c4  28
Summary *172.16.2.1       172.16.1.1       0x80000053  2118  0x22 0xb3b6  28
Summary *172.16.2.2       172.16.1.1       0x80000053  1931  0x22 0xb3b4  28
```

```
Summary *172.22.121.0     172.16.1.1       0x80000056   993  0x22 0x43a8  28
Summary *172.22.122.0     172.16.1.1       0x80000055  1743  0x22 0x44a6  28
Summary *172.31.100.1     172.16.1.1       0x80000054  1556  0x22 0xb840  28
ASBRSum *172.16.2.1       172.16.1.1       0x80000053   806  0x22 0xa5c3  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *20.20.0.0        172.16.1.1       0x80000052   618  0x22 0xc8b1  36
Extern  *20.20.1.0        172.16.1.1       0x80000052   431  0x22 0xbdbb  36
Extern  *20.20.2.0        172.16.1.1       0x80000052   243  0x22 0xb2c5  36
Extern  *20.20.3.0        172.16.1.1       0x80000052    56  0x22 0xa7cf  36
```

Question: Has the mxA router redistributed its static routes into OSPF?

_____

Answer: Yes, the four static routes have been redistributed by the mxA router into OSPF.

Step 1.8

Access the CLI on your mxB device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab2-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.2
Password: lab123
Last login: Tue Nov 19 18:36:44 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxB> configure
Entering configuration mode

[edit]
lab@mxB# load override ajspr/lab2-start.config
load complete

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

lab@mxB>
```

Step 1.9

Issue the **show configuration** command. Use the lab diagram to verify that the mxB router has the correct interface configuration. Verify that OSPF has been enabled (all interfaces in area 0).

```
lab@mxB> show configuration
## Last commit: 2019-11-20 11:31:44 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxB;
```

```
    root-authentication {
        encrypted-password
"$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
.3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.22.121.2/24;
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.22.122.2/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.2/24;
            }
        }
```

```
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.31.100.1/32;
            }
        }
    }
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-0/0/0.0;
            interface ge-0/0/1.0;
        }
    }
}

lab@mxB>
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxB router. If the configuration is incorrect, please notify your instructor.

Step 1.10

Use the **`ping 172.22.121.1 count 5`** command to verify that you can ping the physical interface on the mxA router.

```
lab@mxB> ping 172.22.121.1 count 5
PING 172.22.121.1 (172.22.121.1): 56 data bytes
64 bytes from 172.22.121.1: icmp_seq=0 ttl=64 time=27.643 ms
64 bytes from 172.22.121.1: icmp_seq=1 ttl=64 time=4.170 ms
64 bytes from 172.22.121.1: icmp_seq=2 ttl=64 time=1.454 ms
64 bytes from 172.22.121.1: icmp_seq=3 ttl=64 time=106.790 ms
64 bytes from 172.22.121.1: icmp_seq=4 ttl=64 time=1.608 ms

--- 172.22.121.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.454/28.333/106.790/40.439 ms

lab@mxB>
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.11

Issue the **show ospf neighbor 172.22.121.1** command to verify that the mxB and mxA routers have formed a neighbor relationship in Area 0.

```
lab@mxB> show ospf neighbor 172.22.121.1
Address          Interface              State     ID               Pri  Dead
172.22.121.1     ge-0/0/0.0             Full      172.16.1.1       128    32
```

Question: Have the mxB and mxA router formed a neighbor relationship with one another?

_____

Answer: Yes, the two routers are now OSPF neighbors.

Step 1.12

Access the CLI on your mxC device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab2-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.3
Password: lab123
Last login: Tue Nov 19 18:38:23 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# load override ajspr/lab2-start.config
load complete

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

lab@mxC>
```

Step 1.13

Issue the **show configuration** command. Use the lab diagram to verify that the mxC router has the correct interface and OSPF configuration. Verify that there are four 20.20/24 static routes and a policy statement.

```
lab@mxC> show configuration
## Last commit: 2019-11-20 11:34:26 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxC;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
```

```
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.22.122.1/24;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.2.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.3/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
```

```
                address 172.16.2.1/32;
            }
        }
    }
}
policy-options {
    policy-statement static-to-ospf {
        term 1 {
            from protocol static;
            then accept;
        }
    }
}
routing-options {
    static {
        route 20.20.4.0/24 reject;
        route 20.20.5.0/24 reject;
        route 20.20.6.0/24 reject;
        route 20.20.7.0/24 reject;
    }
    autonomous-system 65512;
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-0/0/0.0;
        }
        area 0.0.0.20 {
            interface ge-0/0/3.0;
        }
        export static-to-ospf;
    }
}
lab@mxC>
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxC router. If the configuration is incorrect, please notify your instructor.

Question: An export policy has been applied to the OSPF configuration. What does it accomplish?

_____

Answer: The statement `export static-to-ospf` causes all of the locally configured and active static routes to be advertised as external OSPF routes (using Type 5 LSAs).

Step 1.14

Use the **ping 10.0.2.2 count 5** command to verify that you can ping the physical interface on the R3-2 router.

```
lab@mxC> ping 10.0.2.2 count 5
PING 10.0.2.2 (10.0.2.2): 56 data bytes
64 bytes from 10.0.2.2: icmp_seq=0 ttl=64 time=3.625 ms
64 bytes from 10.0.2.2: icmp_seq=1 ttl=64 time=1.662 ms
64 bytes from 10.0.2.2: icmp_seq=2 ttl=64 time=2.081 ms
64 bytes from 10.0.2.2: icmp_seq=3 ttl=64 time=1.616 ms
64 bytes from 10.0.2.2: icmp_seq=4 ttl=64 time=28.670 ms

--- 10.0.2.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.616/7.531/28.670/10.595 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.15

Use the **ping 172.22.122.2 count 5** command to verify that you can ping the physical interface on the mxB router.

```
lab@mxC> ping 172.22.122.2 count 5
PING 172.22.122.2 (172.22.122.2): 56 data bytes
64 bytes from 172.22.122.2: icmp_seq=0 ttl=64 time=49.748 ms
64 bytes from 172.22.122.2: icmp_seq=1 ttl=64 time=1.778 ms
64 bytes from 172.22.122.2: icmp_seq=2 ttl=64 time=3.155 ms
64 bytes from 172.22.122.2: icmp_seq=3 ttl=64 time=2.343 ms
64 bytes from 172.22.122.2: icmp_seq=4 ttl=64 time=2.620 ms

--- 172.22.122.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.778/11.929/49.748/18.915 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.16

Issue the **show ospf neighbor** command to verify that the mxC router has formed a neighbor relationship with its directly attached neighboring routers.

```
lab@mxC> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.122.2     ge-0/0/0.0             Full      172.31.100.1     128    37
10.0.2.2         ge-0/0/3.0             Full      172.16.2.2       128    32
```

Question: Has the mxC router formed a neighbor relationship with the directly attached routers?

_____

Answer: Yes, the mxC router now has two OSPF neighbors.

Step 1.17

Issue the **`show ospf database advertising-router 172.16.2.1`** command to verify that the mxC router is properly redistributing its static routes into OSPF.

```
lab@mxC> show ospf database advertising-router 172.16.2.1

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.2.1       172.16.2.1       0x80000055  1310  0x22 0x48c6  48
Summary *10.0.2.0         172.16.2.1       0x80000054  1840  0x22 0xad6f  28
Summary *172.16.2.2       172.16.2.1       0x80000052  1663  0x22 0x9acf  28

    OSPF database, Area 0.0.0.20
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.2.1       172.16.2.1       0x80000056  1134  0x22 0xf64a  36
Network *10.0.2.1         172.16.2.1       0x80000052  2016  0x22 0x6b30  32
Summary *10.0.1.0         172.16.2.1       0x80000054  2546  0x22 0xcc4f  28
Summary *172.16.1.1       172.16.2.1       0x80000053  2369  0x22 0xb7b2  28
Summary *172.16.1.2       172.16.2.1       0x8000001a  2193  0x22 0x2a77  28
Summary *172.16.2.1       172.16.2.1       0x80000054  1487  0x22 0x96d3  28
Summary *172.22.121.0     172.16.2.1       0x80000053    75  0x22 0x4ca0  28
Summary *172.22.122.0     172.16.2.1       0x80000054  2722  0x22 0x35b6  28
Summary *172.31.100.1     172.16.2.1       0x80000052  2899  0x22 0xb544  28
ASBRSum *172.16.1.1       172.16.2.1       0x80000052   252  0x22 0xabbe  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *20.20.4.0        172.16.2.1       0x80000052   957  0x22 0x95df  36
Extern  *20.20.5.0        172.16.2.1       0x80000052   781  0x22 0x8ae9  36
Extern  *20.20.6.0        172.16.2.1       0x80000052   604  0x22 0x7ff3  36
Extern  *20.20.7.0        172.16.2.1       0x80000052   428  0x22 0x74fd  36
```

Question: Has the mxC router redistributed its static routes into OSPF?

_____

Answer: Yes, the four static routes have been redistributed by the mxC router into OSPF.

Part 2: Creating OSPF Stub Areas

In this lab part, you convert Area 10 into a stub area. You then look at the link-state database (LSDB) of the R3-1 router to verify that external routing information is no longer present.

Step 2.1

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxD> set cli logical-system R3-1
Logical system: R3-1

lab@mxD:R3-1>
```

Step 2.2

Issue the show ospf database command to view the types of LSA that have been received by the R3-1 router.

```
lab@mxD:R3-1> show ospf database

    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   172.16.1.1       172.16.1.1       0x8000005b  1926  0x22 0xf24c  36
Router  *172.16.1.2       172.16.1.2       0x80000059    42  0x22 0xaec2  48
Network *10.0.1.2         172.16.1.2       0x8000001a  2476  0x22 0xc312  32
Summary  10.0.2.0         172.16.1.1       0x80000055  1739  0x22 0xc654  28
Summary  172.16.1.1       172.16.1.1       0x80000056   426  0x22 0xa4c5  28
Summary  172.16.2.1       172.16.1.1       0x80000054  1364  0x22 0xb1b7  28
Summary  172.16.2.2       172.16.1.1       0x80000054  1176  0x22 0xb1b5  28
Summary  172.22.121.0     172.16.1.1       0x80000057   239  0x22 0x41a9  28
Summary  172.22.122.0     172.16.1.1       0x80000056   989  0x22 0x42a7  28
Summary  172.31.100.1     172.16.1.1       0x80000055   801  0x22 0xb641  28
ASBRSum  172.16.2.1       172.16.1.1       0x80000054    51  0x22 0xa3c4  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern   20.20.0.0        172.16.1.1       0x80000052  2863  0x22 0xc8b1  36
Extern   20.20.1.0        172.16.1.1       0x80000052  2676  0x22 0xbdbb  36
Extern   20.20.2.0        172.16.1.1       0x80000052  2488  0x22 0xb2c5  36
Extern   20.20.3.0        172.16.1.1       0x80000052  2301  0x22 0xa7cf  36
Extern   20.20.4.0        172.16.2.1       0x80000052  1944  0x22 0x95df  36
Extern   20.20.5.0        172.16.2.1       0x80000052  1767  0x22 0x8ae9  36
Extern   20.20.6.0        172.16.2.1       0x80000052  1591  0x22 0x7ff3  36
Extern   20.20.7.0        172.16.2.1       0x80000052  1414  0x22 0x74fd  36
```

Question: How many external LSA (i.e. routes) have been received by the R3-1 router? Where did they come from?

_____

_____

Answer: The R3-1 router has received 4 external LSAs from the mxA router and 4 external LSAs from the mxC router.

Step 2.3

Enter configuration mode and navigate to `[edit protocols ospf]`. Configure area 10 as a stub area. Commit your configuration when completed.

```
lab@mxD:R3-1> configure
Entering configuration mode

[edit]
lab@mxD:R3-1# edit protocols ospf

[edit protocols ospf]
lab@mxD:R3-1# set area 10 stub

[edit protocols ospf]
lab@mxD:R3-1# commit
commit complete

[edit protocols ospf]
lab@mxD:R3-1#
```

Step 2.4

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, enter configuration mode and navigate to `[edit protocols ospf]`. Configure area 10 as a stub area. Commit your configuration when completed.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit protocols ospf

[edit protocols ospf]
lab@mxA# set area 10 stub

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 2.5

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, issue the **`run show ospf neighbor`** command to verify that the R3-1 and mxA router have reformed their neighbor relationship.

```
[edit protocols ospf]
lab@mxD:R3-1# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.0.1.1         ge-0/0/1.0             Full      172.16.1.1       128    36

[edit protocols ospf]
lab@mxD:R3-1#
```

Step 2.6

Issue the **`run show route 20.20/16`** command to view R3-1 routing table entries.

```
[edit protocols ospf]
lab@mxD:R3-1# run show route 20.20/16

[edit protocols ospf]
lab@mxD:R3-1#
```

Question: Are the 20.20/16 customer static routes from the ABRs visible on the R3-1 router?

_____

Answer: Although the customer routes should still be visible on routers in area 0 and 20, they should not be visible from the R3-1 router (area 10). The purpose of a stub area is to stop the ABRs from injecting external routing information and therefore reduce the size of the LSDB on routers in the stub area.

Step 2.7

Exit out of configuration mode and change the perspective of the CLI to that of the R3-2 logical system by issuing the **`set cli logical-system R3-2`** command.

```
[edit protocols ospf]
lab@mxD:R3-1# exit configuration-mode
Exiting configuration mode

lab@mxD:R3-1> set cli logical-system R3-2
Logical system: R3-2

lab@mxD:R3-2>
```

Step 2.8

Issue the **`show route 20.20/16`** command to view R3-2 routing table entries.

```
lab@mxD:R3-2> show route 20.20/16

inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/24       *[OSPF/150] 2d 19:48:19, metric 0, tag 0
                    > to 10.0.2.1 via ge-0/0/4.0
20.20.1.0/24       *[OSPF/150] 2d 19:48:19, metric 0, tag 0
                    > to 10.0.2.1 via ge-0/0/4.0
20.20.2.0/24       *[OSPF/150] 2d 19:48:19, metric 0, tag 0
                    > to 10.0.2.1 via ge-0/0/4.0
20.20.3.0/24       *[OSPF/150] 2d 19:48:19, metric 0, tag 0
                    > to 10.0.2.1 via ge-0/0/4.0
20.20.4.0/24       *[OSPF/150] 2d 19:52:09, metric 0, tag 0
                    > to 10.0.2.1 via ge-0/0/4.0
20.20.5.0/24       *[OSPF/150] 2d 19:52:09, metric 0, tag 0
                    > to 10.0.2.1 via ge-0/0/4.0
20.20.6.0/24       *[OSPF/150] 2d 19:52:09, metric 0, tag 0
                    > to 10.0.2.1 via ge-0/0/4.0
```

```
20.20.7.0/24        *[OSPF/150] 2d 19:52:09, metric 0, tag 0
                     > to 10.0.2.1 via ge-0/0/4.0
```

Question: Are the 20.20/16 customer static routes from the ABRs visible on the R3-2 router?

_____

Answer: Yes, the external routes are visible on the R3-2 router because it does not belong to a stub area.

Step 2.9

Change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system`** ***`R3-1`*** command.

```
lab@mxD:R3-2> set cli logical-system R3-1
Logical system: R3-1

lab@mxD:R3-1>
```

Step 2.10

Issue the **`show route 172.16/16`** command.

```
lab@mxD:R3-1> show route 172.16/16

inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.1/32       *[OSPF/10] 00:17:18, metric 1
                     > to 10.0.1.1 via ge-0/0/1.0
172.16.1.2/32       *[Direct/0] 2d 23:53:40
                     > via lo0.1
172.16.2.1/32       *[OSPF/10] 00:17:18, metric 3
                     > to 10.0.1.1 via ge-0/0/1.0
172.16.2.2/32       *[OSPF/10] 00:17:18, metric 4
                     > to 10.0.1.1 via ge-0/0/1.0
```

Question: Even though the R3-1 router is a member of the stub area, are the “internal” OSPF destinations reachable from R3-1?

_____

Answer: Yes, stub area configuration only affects the advertisement of external routing information. All internal OSPF destination should remain reachable by the R3-1 router.

Step 2.11

Attempt to ping the loopback address of the R3-2 router by issuing the **`ping 172.16.2.2 count 5`** command.

```
lab@mxD:R3-1> ping 172.16.2.2 count 5
PING 172.16.2.2 (172.16.2.2): 56 data bytes
64 bytes from 172.16.2.2: icmp_seq=0 ttl=61 time=4.134 ms
64 bytes from 172.16.2.2: icmp_seq=1 ttl=61 time=3.823 ms
```

```
64 bytes from 172.16.2.2: icmp_seq=2 ttl=61 time=95.245 ms
64 bytes from 172.16.2.2: icmp_seq=3 ttl=61 time=3.133 ms
64 bytes from 172.16.2.2: icmp_seq=4 ttl=61 time=3.606 ms

--- 172.16.2.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.133/21.988/95.245/36.630 ms
```

Question: Are the loopback addresses visible from each router in the network? Is the ping successful?

_____

_____

Answer: All loopback addresses should be visible and the ping should be successful.

Part 3: Creating Stub No Summaries Areas

In this lab part, you convert the stub area to a stub no-summaries area.

Step 3.1

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, issue the **`run show ospf database`** command.

```
[edit protocols ospf]
lab@mxA# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x80000059  1440  0x22 0x27e8  48
Router   172.16.2.1       172.16.2.1       0x80000056  1624  0x22 0x46c7  48
Router   172.31.100.1     172.31.100.1     0x8000005c  1623  0x22 0xb771  60
Network  172.22.121.2     172.31.100.1     0x80000053  2624  0x22 0x9bec  32
Network  172.22.122.2     172.31.100.1     0x80000053   624  0x22 0x9de8  32
Summary *10.0.1.0         172.16.1.1       0x80000060     6  0x22 0xa76b  28
Summary  10.0.2.0         172.16.2.1       0x80000055  2153  0x22 0xab70  28
Summary *172.16.1.2       172.16.1.1       0x80000004     6  0x22 0x4971  28
Summary  172.16.2.2       172.16.2.1       0x80000053  1977  0x22 0x98d0  28

    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x80000009     6  0x20 0xafe5  36
Router   172.16.1.2       172.16.1.2       0x8000000b     7  0x20 0x6958  48
Network  10.0.1.2         172.16.1.2       0x80000005     7  0x20 0xce0   32
Summary *10.0.2.0         172.16.1.1       0x80000001     6  0x20 0x8de3  28
Summary *172.16.1.1       172.16.1.1       0x80000001     6  0x20 0x6d54  28
Summary *172.16.2.1       172.16.1.1       0x80000001     6  0x20 0x7648  28
Summary *172.16.2.2       172.16.1.1       0x80000001     6  0x20 0x7646  28
Summary *172.22.121.0     172.16.1.1       0x80000001     6  0x20 0xc37   28
Summary *172.22.122.0     172.16.1.1       0x80000001     6  0x20 0xb36   28
Summary *172.31.100.1     172.16.1.1       0x80000001     6  0x20 0x7dd0  28
    OSPF AS SCOPE link state database
```

```
 Type        ID                Adv Rtr            Seq       Age  Opt  Cksum  Len
Extern  *20.20.0.0         172.16.1.1        0x80000054  1238  0x22 0xc4b3  36
Extern  *20.20.1.0         172.16.1.1        0x80000054  1039  0x22 0xb9bd  36
Extern  *20.20.2.0         172.16.1.1        0x80000053  1816  0x22 0xb0c6  36
Extern  *20.20.3.0         172.16.1.1        0x80000053  1628  0x22 0xa5d0  36
Extern   20.20.4.0         172.16.2.1        0x80000053  1271  0x22 0x93e0  36
Extern   20.20.5.0         172.16.2.1        0x80000053  1094  0x22 0x88ea  36
Extern   20.20.6.0         172.16.2.1        0x80000053   918  0x22 0x7df4  36
Extern   20.20.7.0         172.16.2.1        0x80000053   742  0x22 0x72fe  36
```

Question: Which router is generating the summary LSAs that are being advertised in area 10?

_____

Answer: The mxA router is advertising all of the summary LSAs into area 10.

Question: What is the purpose of the summary LSAs that are being advertised into area 10?

_____

Answer: The summary LSAs in area 10 allow the internal routers in area 10 (i.e. R3-1) to reach the networks in area 0 and area 20.

Step 3.2

Convert the stub area to a stub no-summaries area with the **`set area 10 stub no-summaries`** command. Commit your configuration when completed.

```
[edit protocols ospf]
lab@mxA# set area 10 stub no-summaries

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Question: On which routers must you issue the command?

_____

Answer: The command must be issued on the ABRs since it is the ABR's responsibility to generate Type 3 summary LSAs.

Step 3.3

Issue the **`run show ospf database`** command again. It may take a minute for all entries to appear.

```
[edit protocols ospf]
lab@mxA# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x80000059  1807  0x22 0x27e8  48
Router   172.16.2.1       172.16.2.1       0x80000056  1991  0x22 0x46c7  48
Router   172.31.100.1     172.31.100.1     0x8000005c  1990  0x22 0xb771  60
Network  172.22.121.2     172.31.100.1     0x80000053  2991  0x22 0x9bec  32
Network  172.22.122.2     172.31.100.1     0x80000053   991  0x22 0x9de8  32
Summary *10.0.1.0         172.16.1.1       0x80000062     6  0x22 0xa36d  28
Summary  10.0.2.0         172.16.2.1       0x80000055  2520  0x22 0xab70  28
Summary *172.16.1.2       172.16.1.1       0x80000005     6  0x22 0x4772  28
Summary  172.16.2.2       172.16.2.1       0x80000053  2344  0x22 0x98d0  28

    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x8000000b     7  0x20 0xabe7  36
Router   172.16.1.2       172.16.1.2       0x8000000e    13  0x20 0x635b  48
Network  10.0.1.2         172.16.1.2       0x80000007    13  0x20 0x8e2   32
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *20.20.0.0        172.16.1.1       0x80000054  1605  0x22 0xc4b3  36
Extern  *20.20.1.0        172.16.1.1       0x80000054  1406  0x22 0xb9bd  36
Extern  *20.20.2.0        172.16.1.1       0x80000054   135  0x22 0xaec7  36
Extern  *20.20.3.0        172.16.1.1       0x80000053  1995  0x22 0xa5d0  36
Extern   20.20.4.0        172.16.2.1       0x80000053  1638  0x22 0x93e0  36
Extern   20.20.5.0        172.16.2.1       0x80000053  1461  0x22 0x88ea  36
Extern   20.20.6.0        172.16.2.1       0x80000053  1285  0x22 0x7df4  36
Extern   20.20.7.0        172.16.2.1       0x80000053  1109  0x22 0x72fe  36

[edit protocols ospf]
lab@mxA#
```

Question: Are the summary LSAs being advertised into area 10? What does this mean in terms of the reachability of the R3-1 router?

_____

_____

Answer: No, summary LSAs are not being advertised into area 10. Because of this fact, the R3-1 router can only reach destinations that are within area 10.

Step 3.4

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, issue the **`show route`** command to view the routing table of the R3-1 router.

```
lab@mxD:R3-1> show route

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

```
10.0.1.0/24        *[Direct/0] 3d 00:18:00
                    > via ge-0/0/1.0
10.0.1.2/32        *[Local/0] 3d 00:18:00
                      Local via ge-0/0/1.0
172.16.1.2/32      *[Direct/0] 3d 00:18:00
                    > via lo0.1
224.0.0.5/32       *[OSPF/10] 2d 20:47:03, metric 1
                      MultiRecv

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

ff02::2/128        *[INET6/0] 3d 00:18:01
                      MultiRecv
```

Question: Can the R3-1 router reach destinations outside of area 10?

_____

Answer: No, the R3-1 router can only reach destination within area 10.

Step 3.5

Attempt to ping the loopback address of the R3-2 router by issuing the **`ping 172.16.2.2 count 5`** command.

```
lab@mxD:R3-1> ping 172.16.2.2 count 5
PING 172.16.2.2 (172.16.2.2): 56 data bytes
ping: sendto: No route to host
ping: sendto: No route to host
ping: sendto: No route to host
ping: sendto: No route to host
ping: sendto: No route to host

--- 172.16.2.2 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

Question: Were the pings successful?

_____

Answer: No, the ping between the R3 routers should not be successful.

Note

At this point, you have some connectivity issues reaching routers in other OSPF areas in the network because you removed so much information from the routing table. To restore connectivity to the rest of the network, OSPF stub and stub no-summaries areas can use a default route generated by the ABR. Within the Junos OS, this default route is not automatically generated and must be explicitly configured.

Step 3.6

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, restore connectivity in the network by allowing the ABR to generate a default route into the stub area. Issue the **`set area 10 stub default-metric 10`** command and commit your configuration.

```
[edit protocols ospf]
lab@mxA# set area 10 stub default-metric 10

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 3.7

Issue the **`run show ospf database`** command to view the mxA router's LSB.

```
[edit protocols ospf]
lab@mxA# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x8000005a   229  0x22 0x25e9  48
Router   172.16.2.1       172.16.2.1       0x80000057   423  0x22 0x44c8  48
Router   172.31.100.1     172.31.100.1     0x8000005d   247  0x22 0xb572  60
Network  172.22.121.2     172.31.100.1     0x80000054  1247  0x22 0x99ed  32
Network  172.22.122.2     172.31.100.1     0x80000053  2247  0x22 0x9de8  32
Summary *10.0.1.0         172.16.1.1       0x80000062  1262  0x22 0xa36d  28
Summary  10.0.2.0         172.16.2.1       0x80000056   953  0x22 0xa971  28
Summary *172.16.1.2       172.16.1.1       0x80000005  1262  0x22 0x4772  28
Summary  172.16.2.2       172.16.2.1       0x80000054   776  0x22 0x96d1  28

    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x8000000b  1263  0x20 0xabe7  36
Router   172.16.1.2       172.16.1.2       0x8000000e  1269  0x20 0x635b  48
Network  10.0.1.2         172.16.1.2       0x80000007  1269  0x20 0x8e2   32
Summary *0.0.0.0          172.16.1.1       0x80000001    78  0x20 0x6c0a  28
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *20.20.0.0        172.16.1.1       0x80000055   996  0x22 0xc2b4  36
Extern  *20.20.1.0        172.16.1.1       0x80000055   602  0x22 0xb7be  36
```

```
Extern  *20.20.2.0        172.16.1.1       0x80000054  1391  0x22 0xaec7  36
Extern  *20.20.3.0        172.16.1.1       0x80000054  1193  0x22 0xa3d1  36
Extern   20.20.4.0        172.16.2.1       0x80000054    70  0x22 0x91e1  36
Extern   20.20.5.0        172.16.2.1       0x80000053  2717  0x22 0x88ea  36
Extern   20.20.6.0        172.16.2.1       0x80000053  2541  0x22 0x7df4  36
Extern   20.20.7.0        172.16.2.1       0x80000053  2365  0x22 0x72fe  36
```

Question: Are any summary LSAs being advertised by the mxA router into area 10? What does this mean in terms of the reachability of the R3-1 router?

_____

Answer: Yes, one summary LSA which represents a default route is being advertised by the mxA router into area 10. Because of this fact, the R3-1 router can now reach all internal and external OSPF destination.

Step 3.8

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, attempt to ping the loopback of the R3-2 router from the R3-1 router using the **`ping 172.16.2.2 count 5`** command.

```
lab@mxD:R3-1> ping 172.16.2.2 count 5
PING 172.16.2.2 (172.16.2.2): 56 data bytes
64 bytes from 172.16.2.2: icmp_seq=0 ttl=61 time=5.062 ms
64 bytes from 172.16.2.2: icmp_seq=1 ttl=61 time=5.757 ms
64 bytes from 172.16.2.2: icmp_seq=2 ttl=61 time=20.558 ms
64 bytes from 172.16.2.2: icmp_seq=3 ttl=61 time=222.095 ms
64 bytes from 172.16.2.2: icmp_seq=4 ttl=61 time=3.879 ms

--- 172.16.2.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.879/51.470/222.095/85.530 ms
```

Question: Was the ping successful?

_____

Answer: Yes, the ping now should be successful.

Part 4: Creating OSPF Not-So-Stubby-Areas

In this lab part, you convert the OSPF stub area to a NSSA area.

Step 4.1

Enter configuration mode and navigate to `[edit protocols ospf]`. Configure area 10 as an NSSA on the R3-1 router. Commit your configuration and exit configuration mode when completed.

```
lab@mxD:R3-1> configure
Entering configuration mode
```

```
[edit]
lab@mxD:R3-1# edit protocols ospf

[edit protocols ospf]
lab@mxD:R3-1# set area 10 nssa

[edit protocols ospf]
lab@mxD:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R3-1>
```

Step 4.2

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, configure area 10 as an NSSA. Commit your configuration when completed.

```
[edit protocols ospf]
lab@mxA# set area 10 nssa

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 4.3

Issue **`run show ospf neighbor`**.

Note
It might take a minute for the R3-1 adjacency to go back to `Full`.

```
[edit protocols ospf]
lab@mxA# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.121.2     ge-0/0/0.0             Full      172.31.100.1     128    31
10.0.1.2         ge-0/0/3.0             Full      172.16.1.2       128    37

[edit protocols ospf]
lab@mxA#
```

Question: Are all the OSPF neighbors fully adjacent?

_____

Answer: All the neighbors should be adjacent.

Step 4.4

Issue the **`run show ospf database`** command to view the mxA routers LSB.

```
[edit protocols ospf]
lab@mxA# run show ospf database
```

```
    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x8000005a  1547  0x22 0x25e9  48
Router   172.16.2.1       172.16.2.1       0x80000057  1741  0x22 0x44c8  48
Router   172.31.100.1     172.31.100.1     0x8000005d  1565  0x22 0xb572  60
Network  172.22.121.2     172.31.100.1     0x80000054  2565  0x22 0x99ed  32
Network  172.22.122.2     172.31.100.1     0x80000054   565  0x22 0x9be9  32
Summary *10.0.1.0         172.16.1.1       0x80000068   364  0x22 0x9773  28
Summary  10.0.2.0         172.16.2.1       0x80000056  2271  0x22 0xa971  28
Summary *172.16.1.2       172.16.1.1       0x80000001   364  0x22 0x4f6e  28
Summary  172.16.2.2       172.16.2.1       0x80000054  2094  0x22 0x96d1  28

    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x80000002   364  0x20 0xc3d6  36
Router   172.16.1.2       172.16.1.2       0x80000004   365  0x20 0x7751  48
Network  10.0.1.2         172.16.1.2       0x80000001   365  0x20 0x14dc  32
Summary *10.0.2.0         172.16.1.1       0x80000001   364  0x20 0x8de3  28
Summary *172.16.1.1       172.16.1.1       0x80000001   364  0x20 0x6d54  28
Summary *172.16.2.1       172.16.1.1       0x80000001   364  0x20 0x7648  28
Summary *172.16.2.2       172.16.1.1       0x80000001   364  0x20 0x7646  28
Summary *172.22.121.0     172.16.1.1       0x80000001   364  0x20 0xc37   28
Summary *172.22.122.0     172.16.1.1       0x80000001   364  0x20 0xb36   28
Summary *172.31.100.1     172.16.1.1       0x80000001   364  0x20 0x7dd0  28
NSSA    *20.20.0.0        172.16.1.1       0x80000001   364  0x20 0x6d5e  36
NSSA    *20.20.1.0        172.16.1.1       0x80000001   364  0x20 0x6268  36
NSSA    *20.20.2.0        172.16.1.1       0x80000001   364  0x20 0x5772  36
NSSA    *20.20.3.0        172.16.1.1       0x80000001   364  0x20 0x4c7c  36
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *20.20.0.0        172.16.1.1       0x80000056   448  0x22 0xc0b5  36
Extern  *20.20.1.0        172.16.1.1       0x80000056   258  0x22 0xb5bf  36
Extern  *20.20.2.0        172.16.1.1       0x80000055  1175  0x22 0xacc8  36
Extern  *20.20.3.0        172.16.1.1       0x80000055   638  0x22 0xa1d2  36
Extern   20.20.4.0        172.16.2.1       0x80000054  1388  0x22 0x91e1  36
Extern   20.20.5.0        172.16.2.1       0x80000054  1212  0x22 0x86eb  36
Extern   20.20.6.0        172.16.2.1       0x80000054  1035  0x22 0x7bf5  36
Extern   20.20.7.0        172.16.2.1       0x80000054   859  0x22 0x70ff  36

[edit protocols ospf]
lab@mxA#
```

Question: Are external networks being advertised into area 10? What does this mean in terms of the reachability of the R3-1 router?

_____

_____

Answer: Yes, NSSA LSAs that represent the mxA router's external routes are being advertised into area 10. However, the external routes from the mxC router are not being advertised into area 10. Because of these facts, the R3-1 router can reach the external networks being redistributed by the mxA router but it cannot reach the external networks being redistributed by the mxC router.

Question: Are area 0 and area 20 networks being advertised into area 10? What does this mean in terms of the reachability of the R3-1 router?

_____

_____

Answer: Yes, summary LSAs that represent the networks in area 0 and area 20 are being advertised into area 10 by the mxA router. Because of this fact, the R3-1 router can reach all internal OSPF networks.

Step 4.5

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, issue the **`show route`** command to view the routing table of the R3-1 router.

```
lab@mxD:R3-1> show route

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.1.0/24        *[Direct/0] 3d 01:01:13
                    > via ge-0/0/1.0
10.0.1.2/32        *[Local/0] 3d 01:01:13
                      Local via ge-0/0/1.0
10.0.2.0/24        *[OSPF/10] 00:11:22, metric 4
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.0.0/24       *[OSPF/150] 00:11:22, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.1.0/24       *[OSPF/150] 00:11:22, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.2.0/24       *[OSPF/150] 00:11:22, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
20.20.3.0/24       *[OSPF/150] 00:11:22, metric 0, tag 0
                    > to 10.0.1.1 via ge-0/0/1.0
172.16.1.1/32      *[OSPF/10] 00:11:22, metric 1
                    > to 10.0.1.1 via ge-0/0/1.0
```

```
172.16.1.2/32       *[Direct/0] 3d 01:01:13
                     > via lo0.1
172.16.2.1/32       *[OSPF/10] 00:11:22, metric 3
                     > to 10.0.1.1 via ge-0/0/1.0
172.16.2.2/32       *[OSPF/10] 00:11:22, metric 4
                     > to 10.0.1.1 via ge-0/0/1.0
172.22.121.0/24     *[OSPF/10] 00:11:22, metric 2
                     > to 10.0.1.1 via ge-0/0/1.0
172.22.122.0/24     *[OSPF/10] 00:11:22, metric 3
                     > to 10.0.1.1 via ge-0/0/1.0
172.31.100.1/32     *[OSPF/10] 00:11:22, metric 2
                     > to 10.0.1.1 via ge-0/0/1.0
224.0.0.5/32        *[OSPF/10] 2d 21:30:16, metric 1
                       MultiRecv

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

ff02::2/128         *[INET6/0] 3d 01:01:14
                       MultiRecv
```

Question: Are all of the internal OSPF routes visible in the routing table?

_____

Answer: All of the internal networks are visible in the routing table of the R3-1 router.

Question: Are all of the external OSPF routes visible in the routing table?

_____

Answer: No, only the external routes being redistributed by the mxA router are visible in the routing table. The purpose of an NSSA area is to stop the ABRs from injecting external routing information from other areas and therefore reduce the size of the LSDB. However, an NSSA does allow external routes to be injected by an ASBR in that same area, which is why you can see the customer routes from the mxA router.

Part 5: Creating NSSA No Summaries

In this lab part, you change each of the NSSAs to be a `no-summaries` area.

Step 5.1

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, configure the NSSA area for `no-summaries` and commit your configuration.

```
[edit protocols ospf]
lab@mxA# show
```

```
export static-to-ospf;
area 0.0.0.0 {
    interface lo0.0;
    interface ge-0/0/0.0;
}
area 0.0.0.10 {
    nssa;
    interface ge-0/0/3.0;
}

[edit protocols ospf]
lab@mxA# set area 10 nssa no-summaries

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 5.2

Issue a **`run show ospf database`** command to view the LSDB of the mxA router.

Note

Due to route churn, you might see the count fluctuate some before settling down.

```
[edit protocols ospf]
lab@mxA# run show ospf database

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x8000005b  1153  0x22 0x23ea  48
Router   172.16.2.1       172.16.2.1       0x80000058   439  0x22 0x42c9  48
Router   172.31.100.1     172.31.100.1     0x8000005e    86  0x22 0xb373  60
Network  172.22.121.2     172.31.100.1     0x80000055  1086  0x22 0x97ee  32
Network  172.22.122.2     172.31.100.1     0x80000054  2086  0x22 0x9be9  32
Summary *10.0.1.0         172.16.1.1       0x8000006a    64  0x22 0x9375  28
Summary  10.0.2.0         172.16.2.1       0x80000057   968  0x22 0xa772  28
Summary *172.16.1.2       172.16.1.1       0x80000002    64  0x22 0x4d6f  28
Summary  172.16.2.2       172.16.2.1       0x80000055   792  0x22 0x94d2  28

    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *172.16.1.1       172.16.1.1       0x80000004    64  0x20 0xbfd8  36
Router   172.16.1.2       172.16.1.2       0x80000006    65  0x20 0x7353  48
Network  10.0.1.2         172.16.1.2       0x80000003    65  0x20 0x10de  32
NSSA    *20.20.0.0        172.16.1.1       0x80000003    64  0x20 0x6960  36
NSSA    *20.20.1.0        172.16.1.1       0x80000003    64  0x20 0x5e6a  36
NSSA    *20.20.2.0        172.16.1.1       0x80000003    64  0x20 0x5374  36
NSSA    *20.20.3.0        172.16.1.1       0x80000003    64  0x20 0x487e  36
    OSPF AS SCOPE link state database
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Extern  *20.20.0.0        172.16.1.1       0x80000057  1005  0x22 0xbeb6  36
```

```
Extern  *20.20.1.0        172.16.1.1       0x80000056  1779  0x22 0xb5bf  36
Extern  *20.20.2.0        172.16.1.1       0x80000056  1450  0x22 0xaac9  36
Extern  *20.20.3.0        172.16.1.1       0x80000056  1302  0x22 0x9fd3  36
Extern   20.20.4.0        172.16.2.1       0x80000055    86  0x22 0x8fe2  36
Extern   20.20.5.0        172.16.2.1       0x80000054  2733  0x22 0x86eb  36
Extern   20.20.6.0        172.16.2.1       0x80000054  2556  0x22 0x7bf5  36
Extern   20.20.7.0        172.16.2.1       0x80000054  2380  0x22 0x70ff  36

[edit protocols ospf]
lab@mxA#
```

Question: Did a difference exist in the size of the LSDB for area 10 after configuring the `NSSA` area with `no-summaries`?

_____

Answer: Yes, the LSDB was made smaller.

Question: Which LSAs were dropped from the LSDB after making the change?

_____

Answer: The Summary LSAs are no longer present in the LSDB for area 10.

Question: Is the effect of this command different when used in an NSSA as opposed to a stub area?

_____

Answer: The behavior is similar to a stub area with no-summaries.

Step 5.3

Log out of your assigned devices using the **`exit`** command.

```
[edit protocols ospf]
lab@mxA# exit

[edit]
lab@mxA# exit
Exiting configuration mode

lab@mxA> exit

Connection to 172.25.11.1 closed.
```

Tell your instructor that you have completed this lab.

Network Management Diagram

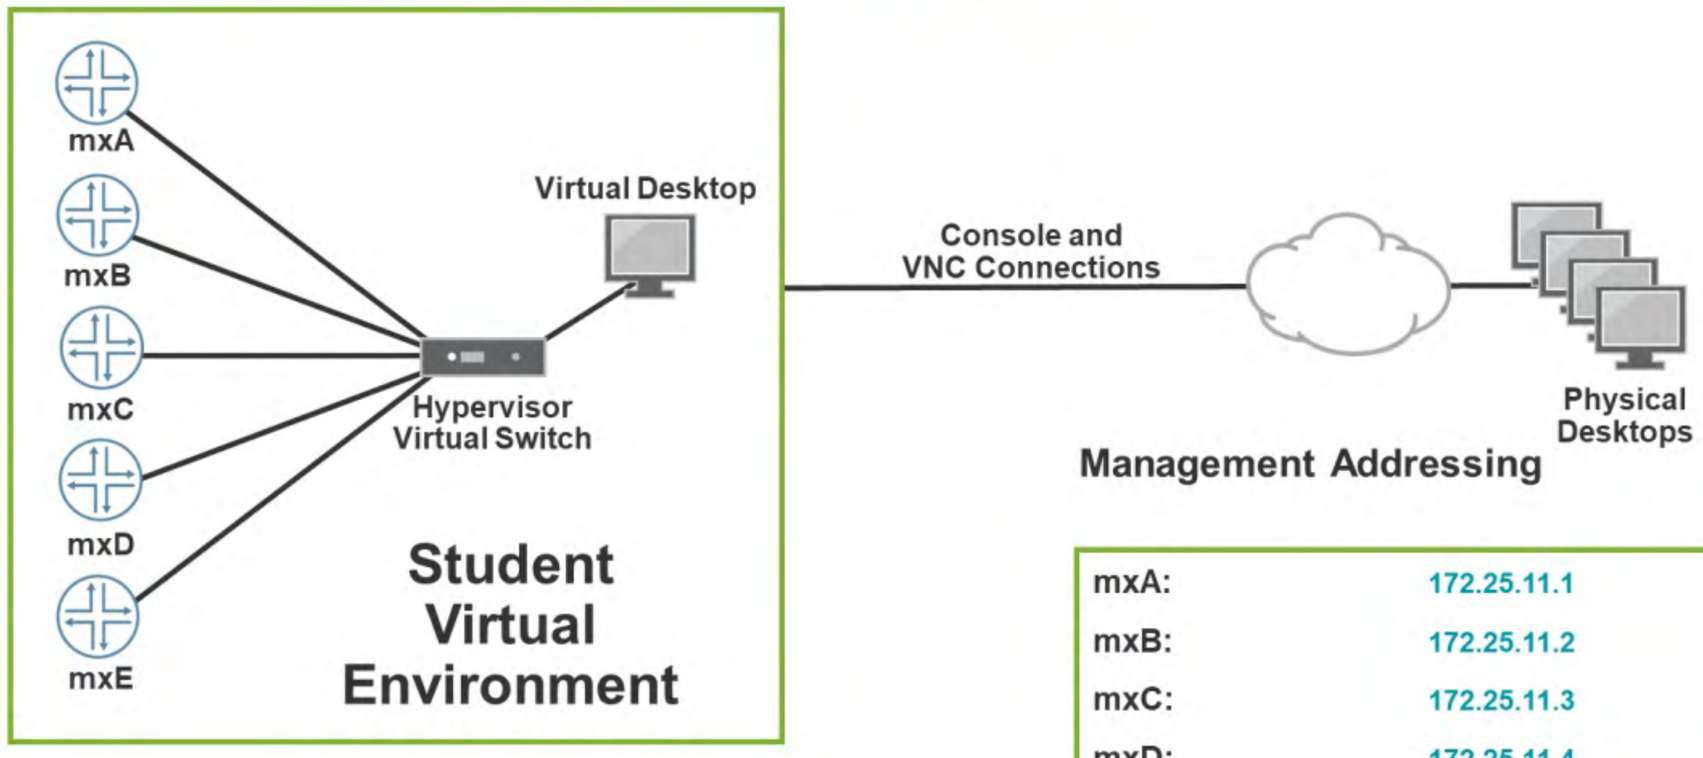

Management Addressing

mxA:	172.25.11.1
mxB:	172.25.11.2
mxC:	172.25.11.3
mxD:	172.25.11.4
mxE:	172.25.11.5
Virtual Desktop:	172.25.11.254

Note: Your instructor will provide address and access information.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab Network Diagram: Advanced OSPF

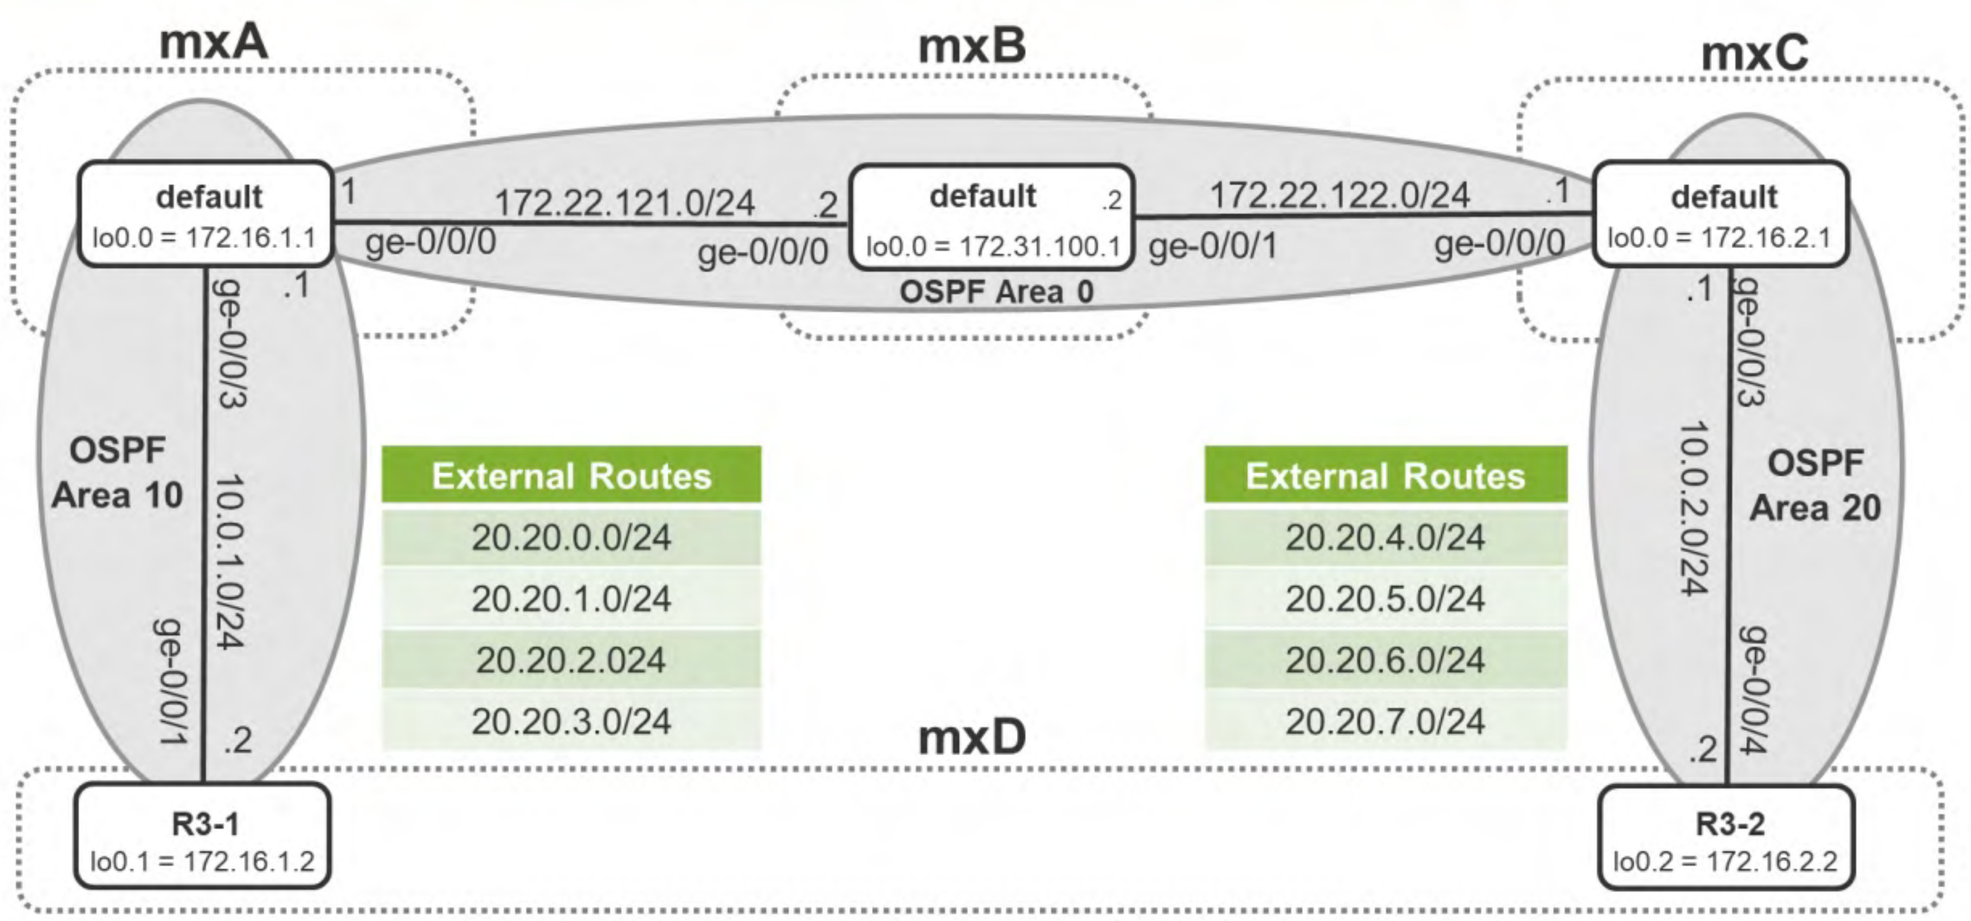

External Routes
20.20.0.0/24
20.20.1.0/24
20.20.2.024
20.20.3.0/24

External Routes
20.20.4.0/24
20.20.5.0/24
20.20.6.0/24
20.20.7.0/24

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab

Advanced OSPF Options

Overview

In this lab, you use the lab diagrams titled “Lab Network Diagram: Advanced OSPF Options -- Parts 1-3” and “Lab Network Diagram: Advanced OSPF Options -- Part 4” to establish a multiarea OSPF routing domain. This lab requires the configuration of a virtual tunnel as backup to the backbone connection and a multiarea adjacency as outlined in RFC 5185. The final part of this lab requires a routing policy to redistribute and advertise routes being received from a RIP network into OSPF external link-state advertisements (LSAs).

By completing this lab, you perform the following tasks:

- Load the starting configuration.
- Establish multiple OSPF adjacencies.
- Configure and verify a virtual tunnel.
- Configure and verify a OSPF multiarea adjacency.
- Establish a RIP neighbor peer sessions.
- Write a routing policy to advertise a default route into RIP.
- Configure prefix-limits in OSPF to prevent excessive external routes.
- Write a routing policy to advertise a RIP summary route into OSPF.
- Write an OSPF import policy to prevent suboptimal routing.

Part 1: Load Starting Configuration

In this lab part, you verify the initial configuration of the routers. You then verify that the interfaces are operational and that OSPF neighbor relationships have formed. Use the lab diagram titled "Lab Network Diagram: Advanced OSPF Options — Parts 1-3" for this lab part.

Note

The instructor will tell you the nature of your access and will provide you with the necessary details to access your assigned device.

Note

The lab topology requires you to display information in the different logical systems of mxD and mxE. You will be changing the perspective of the CLI by issuing the **`set cli logical-system`** ***`ls-name`*** command. Refer to the lab diagram for the correct logical system name.

By changing the perspective of the CLI, you will be able to perform network commands such as **`ping`** or **`traceroute`** from the perspective of the pertinent logical system.

Step 1.1

Ensure that you know to which student login account you have been assigned. Check with your instructor if you are not certain.

Consult the management network diagram to determine the management address of the student devices.

Access the CLI on your mxB device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab3-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.2
Password: lab123
Last login: Wed Nov 20 11:21:00 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxB> configure
Entering configuration mode

[edit]
lab@mxB# load override ajspr/lab3-start.config
load complete

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

lab@mxB>
```

Step 1.2

Issue the **show configuration** command. Use the lab diagram to verify that the mxB router has the correct interface configuration. Verify that OSPF has been enabled (all interfaces in area 0). Note: This is the P1 device in the lab diagram.

```
lab@mxB> show configuration
## Last commit: 2019-11-20 16:53:50 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxB;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    P1 {
        interfaces {
            ge-0/0/0 {
                unit 0 {
                    family inet {
                        address 172.22.121.2/24;
                    }
                }
            }
            ge-0/0/1 {
```

```
                unit 0 {
                    family inet {
                        address 172.22.122.2/24;
                    }
                }
            }
            lo0 {
                unit 0 {
                    family inet {
                        address 172.31.100.1/32;
                    }
                }
            }
        }
        protocols {
            ospf {
                area 0.0.0.0 {
                    interface lo0.0;
                    interface ge-0/0/0.0;
                    interface ge-0/0/1.0;
                }
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.2/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxB router. If the configuration is incorrect, please notify your instructor.

Step 1.3

Access the CLI on your mxE device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab3-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.5
Password: lab123
Last login: Thu Nov 14 21:31:19 2019 from 172.25.11.254
```

```
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxE> configure
Entering configuration mode

[edit]
lab@mxE# load override ajspr/lab3-start.config
load complete

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

lab@mxE>
```

Step 1.4

Issue the **show configuration** command. Use the lab diagram to verify that the mxE router has the correct interface configuration. Verify that OSPF has been enabled (all interfaces in area 0).

```
lab@mxE> show configuration
## Last commit: 2019-11-20 17:00:31 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxE;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
```

```
                interactive-commands any;
            }
        }
    }
    logical-systems {
        P2 {
            interfaces {
                ge-0/0/2 {
                    unit 0 {
                        family inet {
                            address 172.22.124.2/24;
                        }
                    }
                }
                ge-0/0/3 {
                    unit 0 {
                        family inet {
                            address 172.22.123.2/24;
                        }
                    }
                }
                lo0 {
                    unit 1 {
                        family inet {
                            address 172.31.101.1/32;
                        }
                    }
                }
            }
            protocols {
                ospf {
                    area 0.0.0.20 {
                        interface ge-0/0/3.0;
                        interface ge-0/0/2.0;
                        interface lo0.1;
                    }
                }
            }
        }
        P3 {
            interfaces {
                ge-0/0/4 {
                    unit 0 {
                        family inet {
                            address 172.22.126.2/24;
                        }
                    }
                }
                ge-0/0/5 {
                    unit 0 {
                        family inet {
                            address 172.22.125.2/24;
                        }
                    }
                }
```

```
        lo0 {
            unit 2 {
                family inet {
                    address 172.31.102.1/32;
                    address 20.20.1.1/32;
                    address 40.40.1.1/32;
                }
            }
        }
    }
    protocols {
        ospf {
            area 0.0.0.10 {
                nssa;
                interface lo0.2;
                interface ge-0/0/5.0;
                interface ge-0/0/4.0;
            }
        }
        rip {
            group rip {
                export export-rip-policy;
                neighbor ge-0/0/5.0;
                neighbor ge-0/0/4.0;
            }
        }
    }
    policy-options {
        policy-statement export-rip-policy {
            term 1 {
                from {
                    protocol [ static rip aggregate ];
                    route-filter 20.20.0.0/21 orlonger;
                }
                then accept;
            }
            term 2 {
                from {
                    protocol rip;
                    route-filter 0.0.0.0/0 exact;
                }
                then accept;
            }
        }
    }
    routing-options {
        static {
            route 20.20.0.0/24 reject;
            route 20.20.1.0/24 reject;
            route 20.20.2.0/24 reject;
            route 20.20.3.0/24 reject;
            route 20.20.4.0/25 reject;
            route 20.20.4.128/25 reject;
            route 20.20.5.0/26 reject;
            route 20.20.5.64/26 reject;
```

```
                route 20.20.5.128/26 reject;
                route 20.20.5.192/26 reject;
                route 40.40.0.0/24 reject;
                route 40.40.1.0/24 reject;
                route 40.40.2.0/24 reject;
                route 40.40.3.0/24 reject;
                route 172.16.1.2/32 next-hop 172.22.125.1;
                route 172.16.2.2/32 next-hop 172.22.126.1;
                route 172.16.1.4/32 next-hop 172.22.125.1;
                route 172.16.2.4/32 next-hop 172.22.126.1;
                route 10.0.20.0/22 next-hop 172.22.125.1;
                route 10.0.24.0/22 next-hop 172.22.126.1;
            }
            aggregate {
                route 20.20.0.0/21;
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.5/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxE router. If the configuration is incorrect, please notify your instructor.

Note

You might notice that RIP has been configured on the P3 logical system. The RIP configuration will be used in Part 4 of this lab. For now, just ignore the RIP configuration.

Step 1.5

Access the CLI on your mxD device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab3-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.4
Password: lab123
Last login: Wed Nov 20 11:21:18 2019 from 172.25.11.254
```

```
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab3-start.config
load complete

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

lab@mxD>
```

Step 1.6

Issue the **show configuration** command. Use the lab diagram to verify that the mxD router has two logical systems, R3-1 and R3-2, that have the correct interface and OSPF configuration.

```
lab@mxD> show configuration
## Last commit: 2019-11-20 17:02:16 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxD;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
```

```
        }
    }
}
logical-systems {
    R3-1 {
        interfaces {
            ge-0/0/0 {
                unit 0 {
                    family inet {
                        address 172.22.125.1/24;
                    }
                }
            }
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 10.0.10.2/24;
                    }
                }
            }
            lo0 {
                unit 1 {
                    family inet {
                        address 172.16.1.2/32;
                    }
                }
            }
        }
        protocols {
            ospf {
                area 0.0.0.10 {
                    nssa;
                    interface ge-0/0/0.0;
                    interface ge-0/0/1.0;
                    interface lo0.1;
                }
            }
        }
    }
    R3-2 {
        interfaces {
            ge-0/0/3 {
                unit 0 {
                    family inet {
                        address 172.22.126.1/24;
                    }
                }
            }
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 10.0.20.2/24;
                    }
                }
            }
```

```
        lo0 {
            unit 2 {
                family inet {
                    address 172.16.2.2/32;
                }
            }
        }
    }
    protocols {
        ospf {
            area 0.0.0.10 {
                nssa;
                interface lo0.2;
                interface ge-0/0/4.0;
                interface ge-0/0/3.0;
            }
        }
    }
    routing-options {
        autonomous-system 65512;
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.4/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications. If the configuration is incorrect, please notify your instructor.

Step 1.7

Issue the **`show ospf neighbor logical-system R3-1`** command to ensure the R3-1 router has established an OSPF neighbor relationship with the P3 router.

```
lab@mxD> show ospf neighbor logical-system R3-1
Address          Interface              State     ID               Pri  Dead
172.22.125.2     ge-0/0/0.0             Full      20.20.1.1        128    31
```

Question: Have the two routers formed an OSPF neighbor relationship with one another?

_____

Answer: Yes, the two routers should have an OSPF neighbor relationship that is in the `Full` state.

Step 1.8

Issue the **`show ospf neighbor logical-system R3-2`** command to ensure the R3-2 router has established an OSPF neighbor relationship with the P3 router.

```
lab@mxD> show ospf neighbor logical-system R3-2
Address          Interface              State     ID               Pri  Dead
172.22.126.2     ge-0/0/3.0             Full      20.20.1.1        128    33
```

Question: Have the two routers formed an OSPF neighbor relationship with one another?

_____

Answer: Yes, the two routers should have an OSPF neighbor relationship that is in the `Full` state.

Step 1.9

Access the CLI on your mxA device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override ajspr/lab3-start.config`** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.1
Password: lab123
Last login: Thu Nov 21 14:42:41 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# load override ajspr/lab3-start.config
load complete

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

lab@mxA>
```

Step 1.10

Issue the **`show configuration`** command. Use the lab diagram to verify that the mxA router has the correct interface configuration.

```
lab@mxA> show configuration
## Last commit: 2019-11-21 19:19:37 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxA;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.22.121.1/24;
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.22.123.1/24;
            }
        }
    }
    ge-0/0/3 {
```

```
        unit 0 {
            family inet {
                address 10.0.10.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.1/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.1.1/32;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxA router. If the configuration is incorrect, please notify your instructor.

Step 1.11

Use the **`ping 10.0.10.2 count 5`** command to verify that you can ping the physical interface on the R3-1 router.

```
lab@mxA> ping 10.0.10.2 count 5
PING 10.0.10.2 (10.0.10.2): 56 data bytes
64 bytes from 10.0.10.2: icmp_seq=0 ttl=64 time=4.042 ms
64 bytes from 10.0.10.2: icmp_seq=1 ttl=64 time=6.054 ms
64 bytes from 10.0.10.2: icmp_seq=2 ttl=64 time=5.623 ms
64 bytes from 10.0.10.2: icmp_seq=3 ttl=64 time=1.845 ms
64 bytes from 10.0.10.2: icmp_seq=4 ttl=64 time=36.103 ms

--- 10.0.10.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.845/10.733/36.103/12.770 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.12

Use the **ping 172.22.121.2 count 5** command to verify that you can ping the physical interface on the P1 router.

```
lab@mxA> ping 172.22.121.2 count 5
PING 172.22.121.2 (172.22.121.2): 56 data bytes
64 bytes from 172.22.121.2: icmp_seq=0 ttl=64 time=3.489 ms
64 bytes from 172.22.121.2: icmp_seq=1 ttl=64 time=6.187 ms
64 bytes from 172.22.121.2: icmp_seq=2 ttl=64 time=3.745 ms
64 bytes from 172.22.121.2: icmp_seq=3 ttl=64 time=1.832 ms
64 bytes from 172.22.121.2: icmp_seq=4 ttl=64 time=1.905 ms

--- 172.22.121.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.832/3.432/6.187/1.586 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.13

Use the **ping 172.22.123.2 count 5** command to verify that you can ping the physical interface on the P2 router.

```
lab@mxA> ping 172.22.123.2 count 5
PING 172.22.123.2 (172.22.123.2): 56 data bytes
64 bytes from 172.22.123.2: icmp_seq=0 ttl=64 time=3.291 ms
64 bytes from 172.22.123.2: icmp_seq=1 ttl=64 time=1.700 ms
64 bytes from 172.22.123.2: icmp_seq=2 ttl=64 time=5.124 ms
64 bytes from 172.22.123.2: icmp_seq=3 ttl=64 time=2.048 ms
64 bytes from 172.22.123.2: icmp_seq=4 ttl=64 time=1.414 ms

--- 172.22.123.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.414/2.715/5.124/1.364 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.14

Access the CLI on your mxC device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123**.

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab1-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.3
Password: lab123
Last login: Wed Nov 20 16:58:33 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# load override ajspr/lab3-start.config
load complete

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

lab@mxC>
```

Step 1.15

Issue the **show configuration** command. Use the lab diagram to verify that the mxC router has the correct interface configuration.

```
lab@mxC> show configuration
## Last commit: 2019-11-20 16:44:08 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxC;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
```

```
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 172.22.122.1/24;
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.22.124.1/24;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.20.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.3/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.2.1/32;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxC router. If the configuration is incorrect, please notify your instructor.

Step 1.16

Use the **ping 10.0.20.2 count 5** command to verify that you can ping the physical interface on the R3-2 router.

```
lab@mxC> ping 10.0.20.2 count 5
PING 10.0.20.2 (10.0.20.2): 56 data bytes
64 bytes from 10.0.20.2: icmp_seq=0 ttl=64 time=9.943 ms
64 bytes from 10.0.20.2: icmp_seq=1 ttl=64 time=6.318 ms
64 bytes from 10.0.20.2: icmp_seq=2 ttl=64 time=38.288 ms
64 bytes from 10.0.20.2: icmp_seq=3 ttl=64 time=150.987 ms
64 bytes from 10.0.20.2: icmp_seq=4 ttl=64 time=74.104 ms

--- 10.0.20.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 6.318/55.928/150.987/53.392 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.17

Use the **ping 172.22.122.2 count 5** command to verify that you can ping the physical interface on the P1 router.

```
lab@mxC> ping 172.22.122.2 count 5
PING 172.22.122.2 (172.22.122.2): 56 data bytes
64 bytes from 172.22.122.2: icmp_seq=0 ttl=64 time=49.748 ms
64 bytes from 172.22.122.2: icmp_seq=1 ttl=64 time=1.778 ms
64 bytes from 172.22.122.2: icmp_seq=2 ttl=64 time=3.155 ms
64 bytes from 172.22.122.2: icmp_seq=3 ttl=64 time=2.343 ms
64 bytes from 172.22.122.2: icmp_seq=4 ttl=64 time=2.620 ms

--- 172.22.122.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.778/11.929/49.748/18.915 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.18

Use the **`ping 172.22.124.2 count 5`** command to verify that you can ping the physical interface on the P2 router.

```
lab@mxC> ping 172.22.124.2 count 5
PING 172.22.124.2 (172.22.124.2): 56 data bytes
64 bytes from 172.22.124.2: icmp_seq=0 ttl=64 time=4.160 ms
64 bytes from 172.22.124.2: icmp_seq=1 ttl=64 time=1.459 ms
64 bytes from 172.22.124.2: icmp_seq=2 ttl=64 time=203.749 ms
64 bytes from 172.22.124.2: icmp_seq=3 ttl=64 time=30.423 ms
64 bytes from 172.22.124.2: icmp_seq=4 ttl=64 time=2.704 ms

--- 172.22.124.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.459/48.499/203.749/78.365 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Part 2: Establishing the OSPF Adjacencies and Creating a Virtual Tunnel

In this lab part, you will establish the OSPF adjacencies for mxA and mxC. P1 and P2 will provide connectivity among all three OSPF areas. Use the lab diagram labeled “Network Diagram: Advanced OSPF Options and Routing Policy Lab - Parts 1-3” for this part of the lab.

Step 2.1

Enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Establish the OSPF adjacencies, from mxC, with P1, P2, and R3-2. Configure OSPF Area 0 as the backbone area and do not forget the loopback address. Configure OSPF Area 10 as an NSSA and advertise a default route with a metric of 10. At this time, no need exists to configure OSPF on the R3-2, P1, or P2 routers. Their OSPF configuration was preconfigured by the starting configuration that you applied previously. OSPF Area 20 is configured as a normal OSPF area i.e. not a stub or NSSA area. Commit the configuration when completed.

```
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# edit protocols ospf
```

```
[edit protocols ospf]
lab@mxC# set area 0 interface lo0.0

[edit protocols ospf]
lab@mxC# set area 0 interface ge-0/0/0.0

[edit protocols ospf]
lab@mxC# set area 20 interface ge-0/0/1.0

[edit protocols ospf]
lab@mxC# set area 10 nssa default-lsa default-metric 10

[edit protocols ospf]
lab@mxC# set area 10 interface ge-0/0/3.0

[edit protocols ospf]
lab@mxC# commit
commit complete

[edit protocols ospf]
lab@mxC#
```

Step 2.2

Issue a **`run show ospf interface`** command to verify interfaces are running OSPF. It may take a minute for all Nbrs to appear.

```
[edit protocols ospf]
lab@mxC# run show ospf interface
Interface          State   Area          DR ID          BDR ID         Nbrs
ge-0/0/0.0         BDR     0.0.0.0       172.31.100.1   172.16.2.1        1
lo0.0              DR      0.0.0.0       172.16.2.1     0.0.0.0           0
ge-0/0/3.0         BDR     0.0.0.10      172.16.2.2     172.16.2.1        1
ge-0/0/1.0         BDR     0.0.0.20      172.31.101.1   172.16.2.1        1
```

Question: How many interfaces are running OSPF?

_____

Answer: There should be 3 transit interfaces and the loopback interface for a total of 4 interfaces running OSPF.

Step 2.3

Issue a **`run show ospf neighbor`** command to verify the establishment of OSPF adjacencies.

```
[edit protocols ospf]
lab@mxC# run show ospf neighbor
Address          Interface          State     ID             Pri  Dead
172.22.122.2     ge-0/0/0.0         Full      172.31.100.1   128    38
10.0.20.2        ge-0/0/3.0         Full      172.16.2.2     128    32
172.22.124.2     ge-0/0/1.0         Full      172.31.101.1   128    31
```

Question: Are all OSPF adjacencies established and in the `Full` state?

_____

_____

Answer: Yes. There should be three established OSPF adjacencies, one in each OSPF area including Area 0.0.0.10, which is configured as a not-so-stubby-area.

Step 2.4

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Establish the OSPF adjacencies, from mxA, with P1, P2, and R3-1. Configure OSPF Area 0 as the backbone area and do not forget the loopback address. Configure OSPF Area 10 as an NSSA and advertise a default route with a metric of 10. At this time, no need exists to configure adjacencies from within the R3-1, P1, and P2 routers. They were preconfigured by the starting configuration that you applied previously. OSPF Area 20 is configured as a normal OSPF area i.e. not a stub or NSSA area. Commit the configuration when completed.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit protocols ospf

[edit protocols ospf]
lab@mxA# set area 0 interface lo0.0

[edit protocols ospf]
lab@mxA# set area 0 interface ge-0/0/0.0

[edit protocols ospf]
lab@mxA# set area 20 interface ge-0/0/1.0

[edit protocols ospf]
lab@mxA# set area 10 nssa default-lsa default-metric 10

[edit protocols ospf]
lab@mxA# set area 10 interface ge-0/0/3.0

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 2.5

Issue a **`run show ospf interface`** command to verify interfaces are running OSPF. It may take a minute for all Nbrs to appear.

```
[edit protocols ospf]
lab@mxA# run show ospf interface
Interface            State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0           BDR     0.0.0.0         172.31.100.1    172.16.1.1         1
lo0.0                DR      0.0.0.0         0.0.0.0         0.0.0.0            0
ge-0/0/3.0           BDR     0.0.0.10        172.16.1.2      172.16.1.1         1
ge-0/0/1.0           BDR     0.0.0.20        172.31.101.1    172.16.1.1         1
```

Question: How many interfaces are running OSPF?

_____

_____

Answer: There should be 3 transit interfaces and the loopback interface for a total of 4 interfaces running OSPF.

Step 2.6

Issue a **`run show ospf neighbor`** command to verify the establishment of OSPF adjacencies.

```
[edit protocols ospf]
lab@mxA# run show ospf neighbor
Address          Interface              State     ID              Pri  Dead
172.22.121.2     ge-0/0/0.0             Full      172.31.100.1    128    38
10.0.10.2        ge-0/0/3.0             Full      172.16.1.2      128    30
172.22.123.2     ge-0/0/1.0             Full      172.31.101.1    128    33
```

Question: Are all OSPF adjacencies established and in the `Full` state?

_____

_____

Answer: Yes. There should be three established OSPF adjacencies, one in each OSPF area including Area 0.0.0.10, which is configured as a not-so-stubby-area.

Step 2.7

Verify that the routing table has connectivity to all devices in the OSPF domain. Use the **`run show route protocol ospf | match /32`** command to display only the host addresses.

```
[edit protocols ospf]
lab@mxA# run show route protocol ospf | match /32
20.20.1.1/32       *[OSPF/10] 00:02:01, metric 2
40.40.1.1/32       *[OSPF/10] 00:02:01, metric 2
172.16.1.2/32      *[OSPF/10] 00:02:01, metric 1
172.16.2.1/32      *[OSPF/10] 00:02:01, metric 2
172.16.2.2/32      *[OSPF/10] 00:02:01, metric 3
172.31.100.1/32    *[OSPF/10] 00:02:01, metric 1
172.31.101.1/32    *[OSPF/10] 00:02:06, metric 1
172.31.102.1/32    *[OSPF/10] 00:02:01, metric 2
224.0.0.5/32       *[OSPF/10] 00:02:16, metric 1
```

Question: With the exception of the 224.0.0.5/32 OSPF multicast address and your own loopback address, are there entries in the primary routing table (inet.0) for the other eight loopback addresses within the OSPF domain?

_____

_____

Answer: Yes, there should be eight host addresses in the inet.0 routing table, one for each loopback address.

Step 2.8

Using the OSPF **`virtual-link`** command, configure a virtual link in OSPF Area 0 that uses Area 20 as the `transit-area`. The virtual link's `neighbor-id` is the loopback address of the mxC router. The virtual link should be used only as a backup in the event of a P1 failure. This can be accomplished by setting the Area 20 interface to a high metric. Commit this configuration when completed.

```
[edit protocols ospf]
lab@mxA# set area 0 virtual-link transit-area 20 neighbor-id 172.16.2.1

[edit protocols ospf]
lab@mxA# set area 20 interface ge-0/0/1.0 metric 10

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 2.9

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, using the OSPF **`virtual-link`** command, configure a virtual link in OSPF Area 0 that uses Area 20 as the `transit-area`. The virtual link's `neighbor-id` is the loopback address of the mxA router. The virtual link should be used only as a backup in the event of a P1 failure. This can be accomplished by setting the Area 20 interface to a high metric. Commit this configuration when completed.

```
[edit protocols ospf]
lab@mxC# set area 0 virtual-link transit-area 20 neighbor-id 172.16.1.1

[edit protocols ospf]
lab@mxC# set area 20 interface ge-0/0/1.0 metric 10

[edit protocols ospf]
lab@mxC# commit
commit complete

[edit protocols ospf]
lab@mxC#
```

Step 2.10

Verify that the virtual link has been established and that an adjacency has been formed. Use the **`run show ospf interface`** command to display the virtual link interface.

```
[edit protocols ospf]
lab@mxC# run show ospf interface
Interface          State   Area          DR ID          BDR ID         Nbrs
ge-0/0/0.0         BDR     0.0.0.0       172.31.100.1   172.16.2.1        1
lo0.0              DR      0.0.0.0       172.16.2.1     0.0.0.0           0
vl-172.16.1.1      PtToPt  0.0.0.0       0.0.0.0        0.0.0.0           1
ge-0/0/3.0         BDR     0.0.0.10      172.16.2.2     172.16.2.1        1
ge-0/0/1.0         BDR     0.0.0.20      172.31.101.1   172.16.2.1        1
```

Question: What type of interface is created for the virtual link?

_____

Answer: A point-to-point interface is created for the virtual tunnel.

Step 2.11

Verify that the virtual link has an adjacency. Use the **`run show ospf neighbor`** command to display the state of the virtual link interface.

```
[edit protocols ospf]
lab@mxC# run show ospf neighbor
Address          Interface              State     ID              Pri  Dead
172.22.122.2     ge-0/0/0.0             Full      172.31.100.1    128    39
172.22.123.1     vl-172.16.1.1          Full      172.16.1.1        0    36
10.0.20.2        ge-0/0/3.0             Full      172.16.2.2      128    37
172.22.124.2     ge-0/0/1.0             Full      172.31.101.1    128    33
```

Question: What is the OSPF state on the virtual link interface?

_____

Answer: The state of OSPF neighbor on the virtual link interface should be `Full`.

Step 2.12

Use the **`run show route 172.16.1.1`** command to verify that mxC still uses the path through the P1 router to forward packets to the loopback address of mxA and not through the virtual link.

```
[edit protocols ospf]
lab@mxC# run show route 172.16.1.1

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

172.16.1.1/32       *[OSPF/10] 00:08:53, metric 2
                     > to 172.22.122.2 via ge-0/0/0.0
```

Question: Is the mxC router still using the path through the P1 router to forward packets to the loopback address of the mxA router or is mxC using the virtual link?

_____

_____

Answer: The mxC router is using the P1 router and not the virtual link to forward packet to the mxA router. This is because of the metric you set on the interface in Area 20.

Part 3: Configuring OSPF Multiarea

In this lab part, you use the OSPF multiarea adjacency command outlined in RFC 5185 to provide an alternate path for OSPF Area 0.0.0.10. Use the lab diagram labeled "Lab Network Diagram: Advanced OSPF Options – Parts 1-3" for this part of the lab.

Step 3.1

Configure an OSPF Area 10 adjacency through the P1 router as a `secondary` interface with a `metric` of *10*. Adding the Area 0 interface to Area10 with the `secondary` setting provides a backup path for Area 10 in the event of a P3 failure. Commit these changes when completed.

```
[edit protocols ospf]
lab@mxC# set area 10 interface ge-0/0/0.0 secondary metric 10

[edit protocols ospf]
lab@mxC# commit
commit complete

[edit protocols ospf]
lab@mxC#
```

Step 3.2

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, configure an OSPF Area 10 adjacency through the P1 router as a `secondary` interface with a `metric` of *10*. Adding the Area 0 interface to Area10 with the `secondary` setting provides a backup path for Area 10 in the event of a P3 failure. Commit these changes when completed.

```
[edit protocols ospf]
lab@mxA# set area 10 interface ge-0/0/0.0 secondary metric 10

[edit protocols ospf]
lab@mxA# commit
commit complete

[edit protocols ospf]
lab@mxA#
```

Step 3.3

Return to the SSH session that you have open to the mxB router.

From the SSH session with the mxB router, change the perspective of the CLI to that of the P1 logical system by issuing the **set cli logical-system *P1*** command.

```
lab@mxB> set cli logical-system P1
Logical system: P1

lab@mxB:P1>
```

Step 3.4

Enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Configure an OSPF Area 10 adjacency with the mxA and mxC routers specifying a `secondary` interface. Make sure that you configure Area 10 as an NSSA. Commit these changes when completed.

```
lab@mxB:P1> configure
Entering configuration mode

[edit]
lab@mxB:P1# edit protocols ospf

[edit protocols ospf]
lab@mxB:P1# set area 10 interface ge-0/0/0.0 secondary

[edit protocols ospf]
lab@mxB:P1# set area 10 interface ge-0/0/1.0 secondary

[edit protocols ospf]
lab@mxB:P1# set area 10 nssa

[edit protocols ospf]
lab@mxB:P1# commit
commit complete

[edit protocols ospf]
lab@mxB:P1#
```

Step 3.5

Use the **run show ospf interface** command to verify the multiarea interfaces are in OSPF Area 10.

```
[edit protocols ospf]
lab@mxB:P1# run show ospf interface
Interface          State   Area          DR ID          BDR ID         Nbrs
ge-0/0/0.0         DR      0.0.0.0       172.31.100.1   172.16.1.1        1
ge-0/0/1.0         DR      0.0.0.0       172.31.100.1   172.16.2.1        1
lo0.0              DR      0.0.0.0       172.31.100.1   0.0.0.0           0
ge-0/0/0.0         PtToPt  0.0.0.10      0.0.0.0        0.0.0.0           1
ge-0/0/1.0         PtToPt  0.0.0.10      0.0.0.0        0.0.0.0           1
```

Question: The interfaces connected to the mxA and mxC routers have two established states in OSPF. What is the established state for the interfaces in Area 0.0.0.10? Why?

_____

_____

Answer: The established interface state for Area 0.0.0.10 is point-to-point. As outlined in RFC 5185, all secondary multiarea adjacencies are formed using a point-to-point interface.

Step 3.6

Use the **`run show ospf neighbor`** command to verify the establishment of an OSPF Area 10 adjacency with the mxA and mxC routers.

```
[edit protocols ospf]
lab@mxB:P1# run show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.121.1     ge-0/0/0.0             Full      172.16.1.1       128    34
  Area 0.0.0.0
172.22.122.1     ge-0/0/1.0             Full      172.16.2.1       128    32
  Area 0.0.0.0
172.22.121.1     ge-0/0/0.0             Full      172.16.1.1       128    37
  Area 0.0.0.10
172.22.122.1     ge-0/0/1.0             Full      172.16.2.1       128    38
  Area 0.0.0.10
```

Question: How many OSPF adjacencies are there for Area 0.0.0.10?

_____

_____

Answer: Two adjacencies have been formed within OSPF Area 0.0.0.10.

Step 3.7

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, use the **`run show route 172.16.2.2`** command to verify that any packets destined to the loopback address of the R3-2 router will be routed through ge-0/0/3.0.

```
[edit protocols ospf]
lab@mxA# run show route 172.16.2.2

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

172.16.2.2/32       *[OSPF/10] 00:30:11, metric 3
                     > to 10.0.10.2 via ge-0/0/3.0
```

Question: What is the primary path to the R3-2 router's loopback address?

_____

Answer: The primary path to the R3-2 router's loopback address is through ge-0/0/3.0.

Step 3.8

Navigate to the `[edit interfaces]` hierarchy and disable the ge-0/0/3 interface. Commit the configuration when completed.

```
[edit protocols ospf]
lab@mxA# top edit interfaces

[edit interfaces]
lab@mxA# set ge-0/0/3 disable

[edit interfaces]
lab@mxA# commit
commit complete

[edit interfaces]
lab@mxA#
```

Step 3.9

Use the **`run show route 172.16.2.2`** command to verify that the multiarea connectivity for OSPF Area 0.0.0.10 has converged through the P1 router.

```
[edit interfaces]
lab@mxA# run show route 172.16.2.2

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

172.16.2.2/32       *[OSPF/10] 00:00:46, metric 12
                     > to 172.22.121.2 via ge-0/0/0.0
```

Question: What is the path used by the mxA router to send packets to the R3-2 router's loopback address?

_____

Answer: The path to the R3-2 router's loopback address is through ge-0/0/0.0.

Step 3.10

Use the **`run traceroute 172.16.2.2`** command to verify that the traffic is traversing P1.

```
[edit interfaces]
lab@mxA# run traceroute 172.16.2.2
traceroute to 172.16.2.2 (172.16.2.2), 30 hops max, 52 byte packets
 1  172.22.121.2 (172.22.121.2)  2.830 ms  1.770 ms  1.418 ms
 2  172.22.122.1 (172.22.122.1)  5.292 ms  2.354 ms  2.249 ms
 3  172.16.2.2 (172.16.2.2)  11.442 ms  8.471 ms  11.403 ms
```

Question: Did OSPF converge to the multiarea configuration?

_____

_____

Answer: Yes. OSPF converged to the backup multiarea adjacency.

Step 3.11

Navigate to the top of the configuration hierarchy and use the **rollback 1** command to enable ge-0/0/3.0. Commit the configuration when completed.

```
[edit interfaces]
lab@mxA# top

[edit]
lab@mxA# rollback 1
load complete

[edit]
lab@mxA# show | compare
[edit interfaces ge-0/0/3]
-   disable;

[edit]
lab@mxA# commit
commit complete

[edit]
lab@mxA#
```

Step 3.12

Verify that OSPF converged back to the primary path by viewing the route to the R3-2 router's loopback address by using the **run show route 172.16.2.2** command.

```
[edit]
lab@mxA# run show route 172.16.2.2

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

172.16.2.2/32       *[OSPF/10] 00:01:31, metric 3
                     > to 10.0.10.2 via ge-0/0/3.0
```

Question: Has the next hop of the 172.16.2.2/32 route converged back to ge-0/0/3.0?

_____

_____

Answer: Yes. The next hop has changed to ge-0/0/3.0.

Part 4: Configuring External Reachability

In this lab part, you configure an external connection from the R3 routing instance to a RIP network. Once established, the RIP routes will be redistributed into OSPF. Use the lab diagram labeled "Lab Network Diagram: Advanced OSPF Options – Part 4" for this part of the lab.

Step 4.1

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxD> set cli logical-system R3-1
Logical system: R3-1

lab@mxD:R3-1>
```

Step 4.2

Enter configuration mode and remove the ge-0/0/0.0 interface from OSPF Area 10 and configure that interface in protocols RIP. Use a RIP group name of *P3* and commit the configuration when completed.

```
lab@mxD:R3-1> configure
Entering configuration mode

[edit]
lab@mxD:R3-1# delete protocols ospf area 10 interface ge-0/0/0.0

[edit]
lab@mxD:R3-1# set protocols rip group P3 neighbor ge-0/0/0.0

[edit]
lab@mxD:R3-1# commit
commit complete

[edit]
lab@mxD:R3-1#
```

Step 4.3

Use the **`run show route receive-protocol rip 172.22.125.2`** command to verify that RIP routes are being received from the P3 router.

```
[edit]
lab@mxD:R3-1# run show route receive-protocol rip 172.22.125.2

inet.0: 33 destinations, 34 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/21       *[RIP/100] 00:01:07, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.0.0/24       *[RIP/100] 00:01:07, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.1.0/24       *[RIP/100] 00:01:07, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
```

```
20.20.2.0/24        *[RIP/100] 00:01:07, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
20.20.3.0/24        *[RIP/100] 00:01:07, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
20.20.4.0/25        *[RIP/100] 00:01:07, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
20.20.4.128/25      *[RIP/100] 00:01:07, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
20.20.5.0/26        *[RIP/100] 00:01:07, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
20.20.5.64/26       *[RIP/100] 00:01:07, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
20.20.5.128/26      *[RIP/100] 00:01:07, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
20.20.5.192/26      *[RIP/100] 00:01:07, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Step 4.4

Navigate to the `[edit policy-options policy-statement export-default]` hierarchy. Create a policy term to advertise only the OSPF default route to the RIP router.

Do not commit the configuration at this time.

```
[edit]
lab@mxD:R3-1# edit policy-options policy-statement export-default

[edit policy-options policy-statement export-default]
lab@mxD:R3-1# set term 1 from protocol ospf

[edit policy-options policy-statement export-default]
lab@mxD:R3-1# set term 1 from route-filter 0/0 exact

[edit policy-options policy-statement export-default]
lab@mxD:R3-1# set term 1 then accept

[edit policy-options policy-statement export-default]
lab@mxD:R3-1# show
term 1 {
    from {
        protocol ospf;
        route-filter 0.0.0.0/0 exact;
    }
    then accept;
}

[edit policy-options policy-statement export-default]
lab@mxD:R3-1#
```

Step 4.5

Navigate to the top of the configuration hierarchy. Apply the ***export-default*** policy as an export policy in protocol RIP group ***P3***. Commit the configuration and exit to operational mode when completed.

```
[edit policy-options policy-statement export-default]
lab@mxD:R3-1# top

[edit]
lab@mxD:R3-1# set protocols rip group P3 export export-default

[edit]
lab@mxD:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R3-1>
```

Step 4.6

Issue the **`show route advertising-protocol rip 172.22.125.1`** command to verify that the R3-1 router is now redistributing the default route into RIP.

```
lab@mxD:R3-1> show route advertising-protocol rip 172.22.125.1

inet.0: 33 destinations, 34 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/150] 01:44:37, metric 11, tag 0
                    > to 10.0.10.1 via ge-0/0/1.0
```

Question: Is the R3-1 router advertising the default route to P3?

_____

_____

Answer: Yes. The default route is being advertised to P3.

Step 4.7

Change the perspective of the CLI to that of the R3-2 logical system by issuing the **`set cli logical-system R3-2`** command.

```
lab@mxD:R3-1> set cli logical-system R3-2
Logical system: R3-2

lab@mxD:R3-2>
```

Step 4.8

Enter configuration mode and remove the ge-0/0/3.0 interface from OSPF Area 10 and configure that interface in protocols RIP. Use a RIP group name of *P3* and commit the configuration when completed.

```
lab@mxD:R3-2> configure
Entering configuration mode

[edit]
lab@mxD:R3-2# delete protocols ospf area 10 interface ge-0/0/3.0

[edit]
lab@mxD:R3-2# set protocols rip group P3 neighbor ge-0/0/3.0
```

```
[edit]
lab@mxD:R3-2# commit
commit complete

[edit]
lab@mxD:R3-2#
```

Step 4.9

Use the **`run show route receive-protocol rip 172.22.126.2`** command to verify that RIP routes are being received from the P3 router.

```
[edit]
lab@mxD:R3-2# run show route receive-protocol rip 172.22.126.2

inet.0: 29 destinations, 30 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[RIP/100] 00:00:41, metric 3, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.0.0/21       *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.0.0/24       *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.1.0/24       *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.2.0/24       *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.3.0/24       *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.4.0/25       *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.4.128/25     *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.5.0/26       *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.5.64/26      *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.5.128/26     *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
20.20.5.192/26     *[RIP/100] 00:00:41, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Step 4.10

Navigate to the `[edit policy-options policy-statement export-default]` hierarchy. Create a policy term to advertise only the OSPF default route to the RIP router.

Do not commit the configuration at this time.

```
[edit]
lab@mxD:R3-2# edit policy-options policy-statement export-default

[edit policy-options policy-statement export-default]
lab@mxD:R3-2# set term 1 from protocol ospf
```

```
[edit policy-options policy-statement export-default]
lab@mxD:R3-2# set term 1 from route-filter 0/0 exact

[edit policy-options policy-statement export-default]
lab@mxD:R3-2# set term 1 then accept

[edit policy-options policy-statement export-default]
lab@mxD:R3-2# show
term 1 {
    from {
        protocol ospf;
        route-filter 0.0.0.0/0 exact;
    }
    then accept;
}

[edit policy-options policy-statement export-default]
lab@mxD:R3-2#
```

Step 4.11

Navigate to the top of the configuration hierarchy. Apply the ***export-default*** policy as an export policy in protocol RIP group ***P3***. Commit the configuration and exit to operational mode when completed.

```
[edit policy-options policy-statement export-default]
lab@mxD:R3-2# top

[edit]
lab@mxD:R3-2# set protocols rip group P3 export export-default

[edit]
lab@mxD:R3-2# commit
commit complete

[edit]
lab@mxD:R3-2#
```

Step 4.12

Issue the **`run show route advertising-protocol rip 172.22.126.1`** command to verify that the R3-2 router is now redistributing the default route into RIP.

```
[edit]
lab@mxD:R3-2# run show route advertising-protocol rip 172.22.126.1

[edit]
lab@mxD:R3-2#
```

Question: Is the R3-2 router advertising the default route to P3?

_____

_____

Answer: No. The default route is not being advertised to P3.

Question: Why is the default route being advertised to the P3 router by R3-1 but not R3-2?

_____

_____

Answer: The advertisement depends on which router advertised the default route first. The R3-2 router (the second router attempting to advertise the default route) will not advertise the route because its active route for the 0/0 default route is a RIP route, not an OSPF route. This is because the RIP preference of 100 is lower than the OSPF external preference of 150. Recall that the `export-default` policy you just applied uses a `from protocol ospf` in its term. A policy can only act on active routes. Therefore, in the previous output, the R3-2 router cannot advertise the route. We will demonstrate this issue and fix it in subsequent steps.

Step 4.13

Display the default route in the routing table using the **`run show route 0/0 exact`** command.

```
[edit]
lab@mxD:R3-2# run show route 0/0 exact

inet.0: 29 destinations, 30 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[RIP/100] 00:09:43, metric 3, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
                    [OSPF/150] 02:41:52, metric 11, tag 0
                    > to 10.0.20.1 via ge-0/0/4.0
```

Question: What is the active protocol for the default route?

_____

_____

Answer: The active protocol for the default route is RIP, because the preference of RIP (100) is lower than the external preference of OSPF (150).

Step 4.14

Set the OSPF `external-preference` to *90* for the R3-2 router. Doing so will make the OSPF external preference less than the RIP preference of 100. Commit the changes and exit out of configuration mode when completed.

```
[edit]
lab@mxD:R3-2# set protocols ospf external-preference 90

[edit]
lab@mxD:R3-2# commit and-quit
```

```
commit complete
Exiting configuration mode

lab@mxD:R3-2>
```

Step 4.15

Display the default route in the routing table using the **`run show route 0/0 exact`** command.

```
lab@mxD:R3-2> show route 0/0 exact

inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/90] 00:03:56, metric 11, tag 0
                    > to 10.0.20.1 via ge-0/0/4.0
```

Question: What is the active protocol for the default route?

_____

_____

Answer: The active protocol for the default route is OSPF, because the preference of RIP (100) is higher than the external preference of OSPF (90).

Step 4.16

Use the **`show route advertising-protocol rip 172.22.126.1`** command to verify that the R3-2 router is advertising the default route to the RIP router.

```
lab@mxD:R3-2> show route advertising-protocol rip 172.22.126.1

inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[OSPF/90] 00:04:31, metric 11, tag 0
                    > to 10.0.20.1 via ge-0/0/4.0
```

Note

Do not be alarmed if you do not see the RIP route as shown in the previous output. It eventually times out and is removed from the routing table. The key is to see that the active route for 0/0 is now an OSPF route and that the R3-2 router is now advertising it.

Question: Is the route now being advertised to the RIP network?

_____

_____

Answer: Yes. The lower OSPF preference has made the default route active under OSPF, which matches the RIP export policy.

Step 4.17

Change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxD:R3-2> set cli logical-system R3-1
Logical system: R3-1

lab@mxD:R3-1>
```

Step 4.18

Enter configuration mode and set the OSPF `external-preference` to *90* for the R3-1 router. Doing so will make the OSPF external preference less than the RIP preference of 100 and prevent the previous problem (on R3-2) from happening on R3-1 in the event of a failure in the network. Commit the changes when completed.

```
lab@mxD:R3-1> configure
Entering configuration mode

[edit]
lab@mxD:R3-1# set protocols ospf external-preference 90

[edit]
lab@mxD:R3-1# commit
commit complete

[edit]
lab@mxD:R3-1#
```

Step 4.19

Issue the **`run show route protocol rip`** command to view the routes that have been learned from P3.

```
[edit]
lab@mxD:R3-1# run show route protocol rip

inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/21       *[RIP/100] 00:36:28, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.0.0/24       *[RIP/100] 00:36:28, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.1.0/24       *[RIP/100] 00:36:28, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.2.0/24       *[RIP/100] 00:36:28, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.3.0/24       *[RIP/100] 00:36:28, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.4.0/25       *[RIP/100] 00:36:28, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.4.128/25     *[RIP/100] 00:36:28, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.5.0/26       *[RIP/100] 00:36:28, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
```

```
20.20.5.64/26       *[RIP/100] 00:36:28, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
20.20.5.128/26      *[RIP/100] 00:36:28, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
20.20.5.192/26      *[RIP/100] 00:36:28, metric 2, tag 0
                     > to 172.22.125.2 via ge-0/0/0.0
224.0.0.9/32        *[RIP/100] 00:00:50, metric 1
                       MultiRecv

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Is it necessary to receive all of the routes in the 20.20/16 range from P3 or is there a single route being received that could be used to summarize them all?

_____

_____

Answer: It is unnecessary for the R3 routers to receive all 11 routes that the P3 router is advertising because there is a single route, 20.20.0.0/21, that summarizes all of the other 10 routes. In the next few steps you will only allow the R3 routers to receive the single 20.20.0.0/21 route.

Step 4.20

Navigate to the `[edit policy-options policy-statement import-rip-route]` hierarchy. Create a policy to accept only the 20.20/21 summary route from the RIP router.

```
[edit]
lab@mxD:R3-1# edit policy-options policy-statement import-rip-route

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-1# set term 1 from protocol rip

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-1# set term 1 from route-filter 20.20/21 exact

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-1# set term 1 then accept

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-1# set term 2 from protocol rip

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-1# set term 2 from route-filter 20.20/21 longer

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-1# set term 2 then reject

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-1# show
term 1 {
    from {
        protocol rip;
```

```
        route-filter 20.20.0.0/21 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol rip;
        route-filter 20.20.0.0/21 longer;
    }
    then reject;
}

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-1#
```

Step 4.21

Navigate to the top of the configuration hierarchy and apply the policy as an import policy under protocols RIP group ***P3***. Commit the configuration when completed.

```
[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-1# top

[edit]
lab@mxD:R3-1# set protocols rip group P3 import import-rip-route

[edit]
lab@mxD:R3-1# commit
commit complete

[edit]
lab@mxD:R3-1#
```

Step 4.22

Display the routes being received from the P3 router using the **`run show route receive-protocol rip 172.22.125.2`** command. Verify that only the summary route is now being received from the P3 RIP router.

```
[edit]
lab@mxD:R3-1# run show route receive-protocol rip 172.22.125.2

inet.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/21       *[RIP/100] 00:49:10, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Is the RIP import policy working?

_____

_____

Answer: Because only the summary route is being received from the RIP neighbor, yes, the import policy appears to be working.

Step 4.23

Navigate to the `[edit policy-options policy-statement export-rip-route]` hierarchy. Create a routing policy to redistribute the RIP summary route into OSPF.

Do not commit at this point.

```
[edit]
lab@mxD:R3-1# edit policy-options policy-statement export-rip-route

[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-1# set term 1 from protocol rip

[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-1# set term 1 then accept

[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-1# show
term 1 {
    from protocol rip;
    then accept;
}

[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-1#
```

Step 4.24

Navigate to the top of the configuration hierarchy. Before applying the policy as an OSPF export policy, protect the network from unnecessary routes by configuring a prefix-export-limit of *1* within protocols OSPF. Commit the configuration and exit to operational mode when completed.

```
[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-1# top

[edit]
lab@mxD:R3-1# set protocols ospf prefix-export-limit 1

[edit]
lab@mxD:R3-1# set protocols ospf export export-rip-route

[edit]
lab@mxD:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R3-1>
```

Step 4.25

Issue the **show ospf database** command for view the R3-1 router's LSDB.

```
lab@mxD:R3-1> show ospf database

    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   20.20.1.1        20.20.1.1        0x80000023  3388  0x20 0x1beb  84
Router   172.16.1.1       172.16.1.1       0x8000000a  1751  0x20 0x3e6   48
Router  *172.16.1.2       172.16.1.2       0x80000031    63  0x20 0xe99d  48
Router   172.16.2.1       172.16.2.1       0x80000007  1864  0x20 0xc412  48
Router   172.16.2.2       172.16.2.2       0x80000023  2720  0x20 0xdfa0  48
Router   172.31.100.1     172.31.100.1     0x80000005  2176  0x20 0xe1bf  48
Network *10.0.10.2        172.16.1.2       0x80000004   315  0x20 0xaa3a  32
Network  10.0.20.2        172.16.2.2       0x80000004  2852  0x20 0x4b8c  32
Summary  172.16.1.1       172.16.1.1       0x80000007   173  0x20 0x615a  28
Summary  172.16.1.1       172.16.2.1       0x80000005  1464  0x20 0x7248  28
Summary  172.16.1.1       172.31.100.1     0x80000005  1535  0x20 0x3f0b  28
Summary  172.16.2.1       172.16.1.1       0x80000005  2351  0x20 0x6e4c  28
Summary  172.16.2.1       172.16.2.1       0x80000007   464  0x20 0x4f6a  28
Summary  172.16.2.1       172.31.100.1     0x80000005  1396  0x20 0x3415  28
Summary  172.22.121.0     172.16.1.1       0x80000007  1901  0x20 0xff3d  28
Summary  172.22.121.0     172.16.2.1       0x80000006  1264  0x20 0x537   28
Summary  172.22.121.0     172.31.100.1     0x80000005   363  0x20 0xd3f8  28
Summary  172.22.122.0     172.16.1.1       0x80000005  2201  0x20 0x33a   28
Summary  172.22.122.0     172.16.2.1       0x80000007  2242  0x20 0xed4d  28
Summary  172.22.122.0     172.31.100.1     0x80000005   176  0x20 0xc803  28
Summary  172.22.123.0     172.16.1.1       0x80000008   273  0x20 0x42ee  28
Summary  172.22.123.0     172.16.2.1       0x80000006  2767  0x20 0x49e7  28
Summary  172.22.123.0     172.31.100.1     0x80000005   924  0x20 0x229e  28
Summary  172.22.124.0     172.16.1.1       0x80000006   673  0x20 0x45eb  28
Summary  172.22.124.0     172.16.2.1       0x80000007  2467  0x20 0x32fd  28
Summary  172.22.124.0     172.31.100.1     0x80000005   737  0x20 0x17a8  28
Summary  172.31.100.1     172.16.1.1       0x80000005  2051  0x20 0x75d4  28
Summary  172.31.100.1     172.16.2.1       0x80000006   265  0x20 0x6cdb  28
Summary  172.31.100.1     172.31.100.1     0x80000004  2649  0x20 0x3d9c  28
Summary  172.31.101.1     172.16.1.1       0x80000006   473  0x20 0xc27c  28
Summary  172.31.101.1     172.16.2.1       0x80000005  2617  0x20 0xbd81  28
Summary  172.31.101.1     172.31.100.1     0x80000005   550  0x20 0x9e2e  28
NSSA     0.0.0.0          172.16.1.1       0x80000007  1826  0x20 0x382c  36
NSSA     0.0.0.0          172.16.2.1       0x80000006   564  0x20 0x3331  36
NSSA    *20.20.0.0        172.16.1.2       0x80000001    63  0x28 0x35d2  36
```

Question: What type of LSA is the R3-1 router using to advertise the 20.20.0.0/21 route into the OSPF domain? Why?

_____

_____

Answer: The R3-1 router is advertising the 20.20.0.0/21 route into OSPF using a Type 7 NSSA LSA because the R3-1 router is part of an NSSA.

Step 4.26

Change the perspective of the CLI to that of the R3-2 logical system by issuing the **`set cli logical-system R3-2`** command.

```
lab@mxD:R3-1> set cli logical-system R3-2
Logical system: R3-2

lab@mxD:R3-2>
```

Step 4.27

Issue the **`show route protocol rip`** command to view the routes that have been learned from P3.

```
lab@mxD:R3-2> show route protocol rip

inet.0: 29 destinations, 30 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/21        [RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.0.0/24       *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.1.0/24       *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.2.0/24       *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.3.0/24       *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.4.0/25       *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.4.128/25     *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.5.0/26       *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.5.64/26      *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.5.128/26     *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
20.20.5.192/26     *[RIP/100] 00:52:02, metric 2, tag 0
                     > to 172.22.126.2 via ge-0/0/3.0
224.0.0.9/32       *[RIP/100] 00:40:57, metric 1
                       MultiRecv

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Is it necessary to receive all of the routes in the 20.20/16 range from P3 or is there a single route being received that could be used to summarize them all?

_____

_____

Answer: It is unnecessary for the R3 routers to receive all 11 routes that the P3 router is advertising because there is a single route, 20.20.0.0/21, that summarizes all of the other 10 routes. In the next few steps you will only allow the R3 routers to receive the single 20.20.0.0/21 route.

Step 4.28

Enter configuration mode and navigate to the `[edit policy-options policy-statement import-rip-route]` hierarchy. Create a policy to accept only the 20.20/21 summary route from the RIP router.

```
lab@mxD:R3-2> configure
Entering configuration mode

[edit]
lab@mxD:R3-2# edit policy-options policy-statement import-rip-route

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-2# set term 1 from protocol rip

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-2# set term 1 from route-filter 20.20/21 exact

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-2# set term 1 then accept

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-2# set term 2 from protocol rip

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-2# set term 2 from route-filter 20.20/21 longer

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-2# set term 2 then reject

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-2# show
term 1 {
    from {
        protocol rip;
        route-filter 20.20.0.0/21 exact;
    }
    then accept;
}
term 2 {
    from {
        protocol rip;
```

```
        route-filter 20.20.0.0/21 longer;
    }
    then reject;
}

[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-2#
```

Step 4.29

Navigate to the top of the configuration hierarchy and apply the policy as an import policy under protocols RIP group ***P3***. Commit the configuration when completed.

```
[edit policy-options policy-statement import-rip-route]
lab@mxD:R3-2# top

[edit]
lab@mxD:R3-2# set protocols rip group P3 import import-rip-route

[edit]
lab@mxD:R3-2# commit
commit complete

[edit]
lab@mxD:R3-2#
```

Step 4.30

Display the routes being received from the P3 router using the **`run show route receive-protocol rip 172.22.126.2`** command. Verify that only the summary route is now being received from the P3 RIP router.

```
[edit]
lab@mxD:R3-2# run show route receive-protocol rip 172.22.126.2

inet.0: 19 destinations, 20 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/21           [RIP/100] 00:55:40, metric 2, tag 0
                       > to 172.22.126.2 via ge-0/0/3.0

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Is the RIP import policy working?

_____

Answer: Because only the summary route is being received from the RIP neighbor, yes, the import policy appears to be working.

Step 4.31

Navigate to the `[edit policy-options policy-statement export-rip-route]` hierarchy. Create a routing policy to redistribute the RIP summary route into OSPF.

Note

Do not commit at this point.

```
[edit]
lab@mxD:R3-2# edit policy-options policy-statement export-rip-route

[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-2# set term 1 from protocol rip

[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-2# set term 1 then accept

[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-2# show
term 1 {
    from protocol rip;
    then accept;
}

[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-2#
```

Step 4.32

Navigate to the top of the configuration hierarchy. Before applying the policy as an OSPF export policy, protect the network from unnecessary routes by configuring a prefix-export-limit of *1* within protocols OSPF. Commit the configuration when completed.

```
[edit policy-options policy-statement export-rip-route]
lab@mxD:R3-2# top

[edit]
lab@mxD:R3-2# set protocols ospf prefix-export-limit 1

[edit]
lab@mxD:R3-2# set protocols ospf export export-rip-route

[edit]
lab@mxD:R3-2# commit
commit complete

[edit]
lab@mxD:R3-2#
```

Step 4.33

Verify connectivity to the RIP network by performing a trace to the RIP router using the redistributed RIP summary route. Enter the **`run traceroute 20.20.1.1`** command to verify connectivity.

```
[edit]
lab@mxD:R3-2# run traceroute 20.20.1.1
traceroute to 20.20.1.1 (20.20.1.1), 30 hops max, 52 byte packets
 1  10.0.20.1 (10.0.20.1)  1.658 ms  1.566 ms  1.249 ms
 2  172.22.122.2 (172.22.122.2)  2.085 ms  1.725 ms  1.769 ms
 3  172.22.121.1 (172.22.121.1)  373.895 ms  105.236 ms  5.557 ms
 4  10.0.10.2 (10.0.10.2)  4.125 ms  8.154 ms  13.172 ms
 5  20.20.1.1 (20.20.1.1)  9.111 ms  21.108 ms  12.488 ms
```

Question: What would be causing the sub-optimal path to the RIP network for the R3-2 router?

_____

_____

Answer: The R3-2 router must have a route in its routing table that is more preferred (lower preference number) than the RIP route learned from P3.

Step 4.34

Issue the **`run show route 20.20.1.1`** command to view the route that R3-2 is using to reach the P3 router.

```
[edit]
lab@mxD:R3-2# run show route 20.20.1.1

inet.0: 19 destinations, 20 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/21       *[OSPF/90] 00:34:24, metric 2, tag 0
                    > to 10.0.20.1 via ge-0/0/4.0
                    [RIP/100] 01:18:37, metric 2, tag 0
                    > to 172.22.126.2 via ge-0/0/3.0
```

Question: What type of route is being used by R3-2 to forward traffic to the P3 router?

_____

_____

Answer: The R3-2 router is using the OSPF route being advertised (originally) by the R3-1 router.

Question: Can you think of another problem that this behavior is causing?

_____

_____

Answer: In the previous steps you wrote and applied a policy to redistribute the 20.20.0.0/21 RIP route into OSPF, because the OSPF route is the active route in the routing table, R3-2 must not be advertising the 20.20.0.0/21 route into OSPF.

Step 4.35

Issue the **`run show ospf database nssa`** command to view the NSSA routes that the R3-2 router is advertising.

```
[edit]
lab@mxD:R3-2# run show ospf database nssa
```

```
    OSPF database, Area 0.0.0.10
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
NSSA     0.0.0.0          172.16.1.1       0x80000008  1282  0x20 0x362d  36
NSSA     0.0.0.0          172.16.2.1       0x80000006  2868  0x20 0x3331  36
NSSA     20.20.0.0        172.16.1.2       0x80000001  2372  0x28 0x35d2  36
```

Question: Is the R3-2 router advertising any NSSA routes?

_____

_____

Answer: No. The R3-2 router is only receiving NSSA routes.

Question: Without changing the route preference, how can you ensure that the R3-2 router will advertise the RIP-learned 20.20.0.0/21 route as an OSPF NSSA?

_____

_____

Answer: One way to fix the problem would be to write and apply an OSPF import policy that blocks the NSSA route being advertised by R3-1 from getting into the R3-2 router's routing table.

Step 4.36

Navigate to the `[edit policy-options policy-statement ospf-import]` hierarchy. Create the OSPF import policy to block the 20.20/21 network that is being advertised by R3-1 using an NSSA LSA from being installed in the routing table from OSPF.

```
[edit]
lab@mxD:R3-2# edit policy-options policy-statement ospf-import

[edit policy-options policy-statement ospf-import]
lab@mxD:R3-2# set term 1 from route-filter 20.20.0.0/21 orlonger

[edit policy-options policy-statement ospf-import]
lab@mxD:R3-2# set term 1 then reject

[edit policy-options policy-statement ospf-import]
lab@mxD:R3-2# show
term 1 {
    from {
        route-filter 20.20.0.0/21 orlonger;
    }
    then reject;
}

[edit policy-options policy-statement ospf-import]
lab@mxD:R3-2#
```

Step 4.37

Navigate to the top of the configuration hierarchy and apply the OSPF import policy. Commit the changes when completed and exit to operational mode.

```
[edit policy-options policy-statement ospf-import]
lab@mxD:R3-2# top

[edit]
lab@mxD:R3-2# set protocols ospf import ospf-import

[edit]
lab@mxD:R3-2# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R3-2>
```

Step 4.38

Verify that the OSPF import policy is working and that optimal routing is being used to the RIP network by using the **traceroute 20.20.1.1 routing-instance *R3-X*** command.

```
lab@mxD:R3-2> traceroute 20.20.1.1
traceroute to 20.20.1.1 (20.20.1.1), 30 hops max, 52 byte packets
 1  20.20.1.1 (20.20.1.1)  2.202 ms  1.697 ms  1.711 ms
```

Question: Is the OSPF import policy working?

_____

_____

Answer: Yes. The OSPF import policy is providing an optimal path to the RIP network for the R3-2 router.

Note

Because of the newly applied policy, you actually moved the suboptimal routing problem that the R3-2 router was facing over to the R3-1 router. Feel free to fix the suboptimal routing that is happening on the R3-1 router if you have the time.

Step 4.39

Log out of your assigned devices using the **exit** command.

```
lab@mxD:R3-2> exit
Connection to 172.25.11.4 closed.
```

Tell your instructor that you have completed this lab.

Network Management Diagram

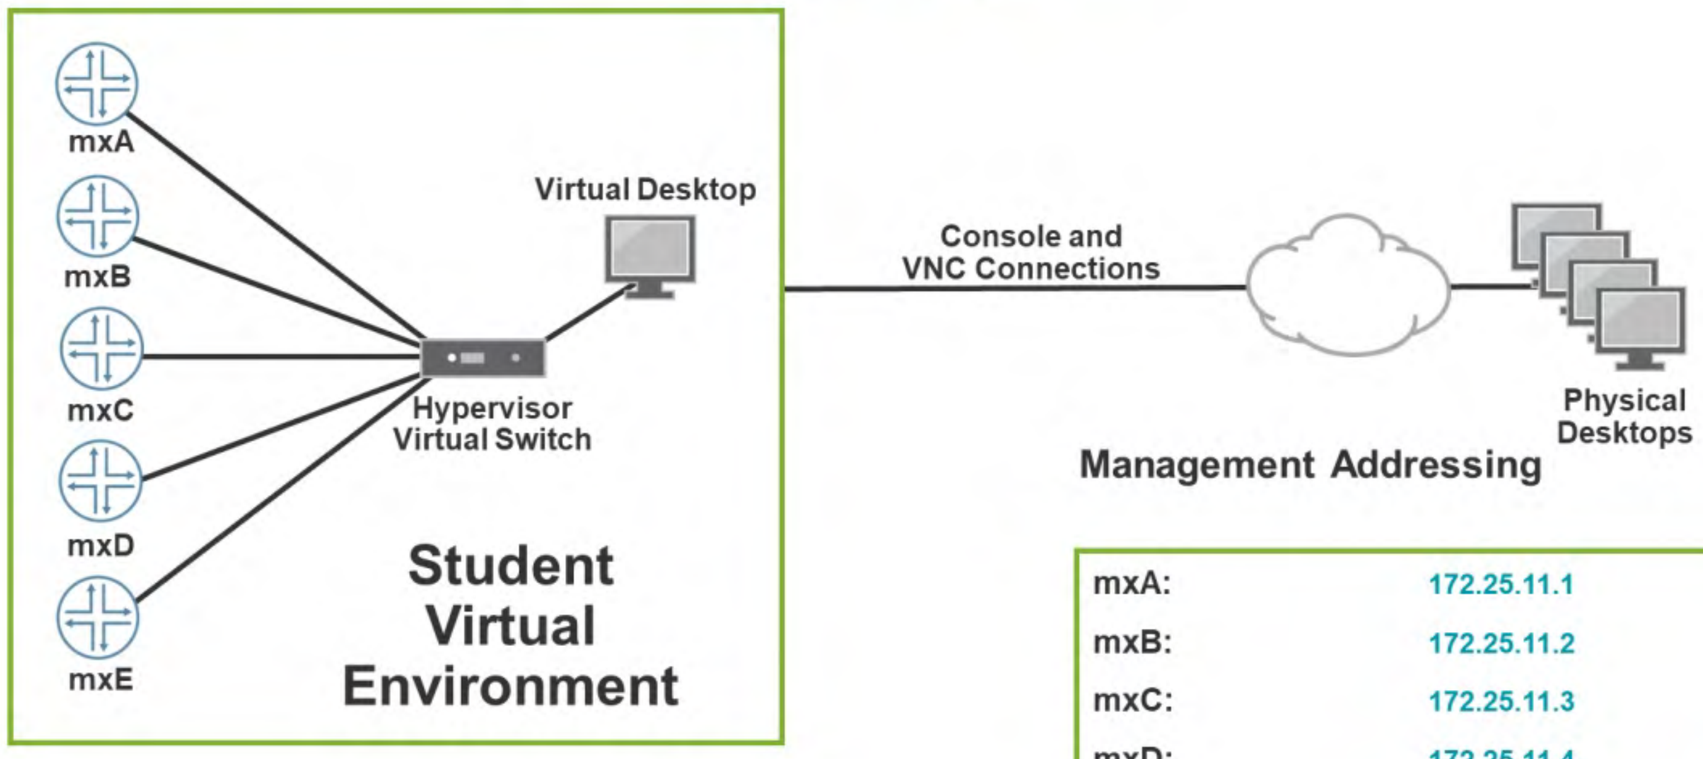

Management Addressing

mxA:	172.25.11.1
mxB:	172.25.11.2
mxC:	172.25.11.3
mxD:	172.25.11.4
mxE:	172.25.11.5
Virtual Desktop:	172.25.11.254

Note: Your instructor will provide address and access information.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab Network Diagram: Advanced OSPF Options – Parts 1-3

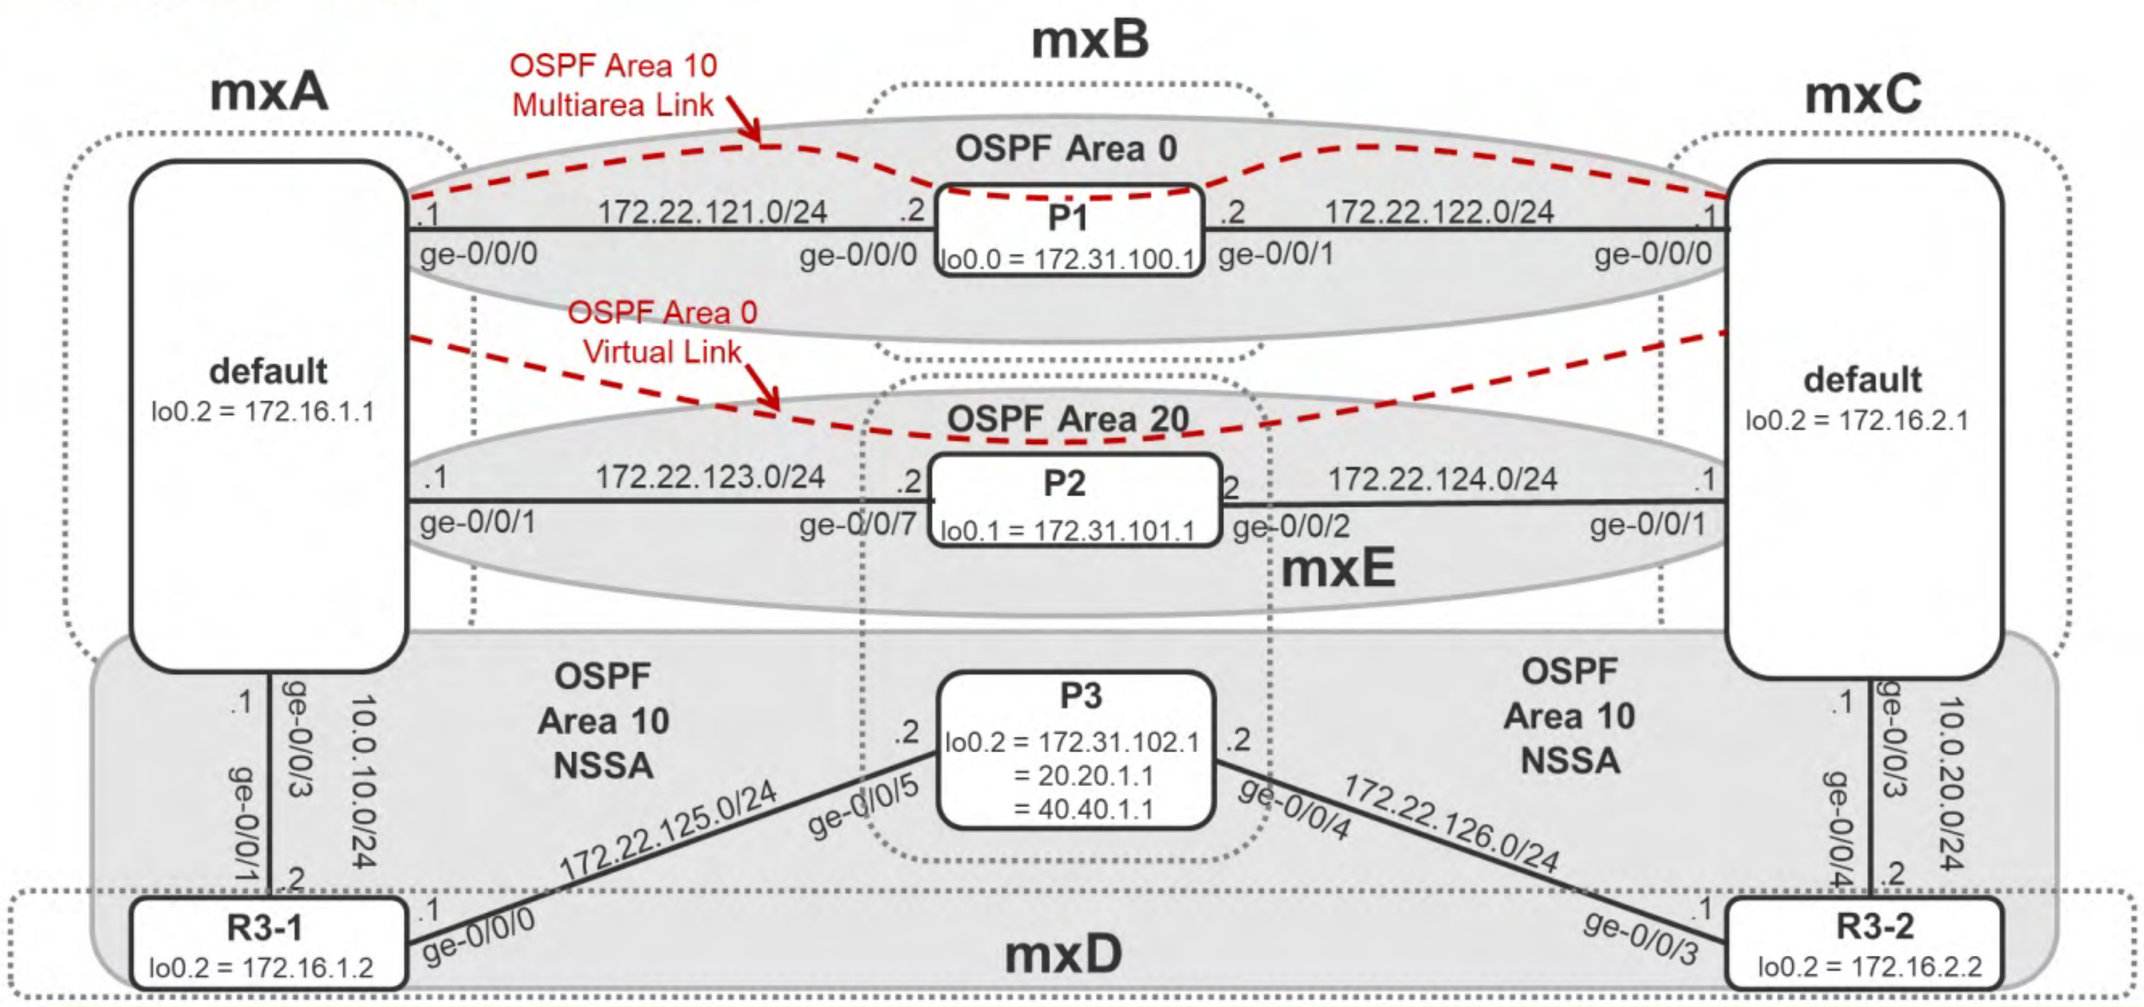

© 2020 Juniper Networks, Inc. All Rights Reserved.

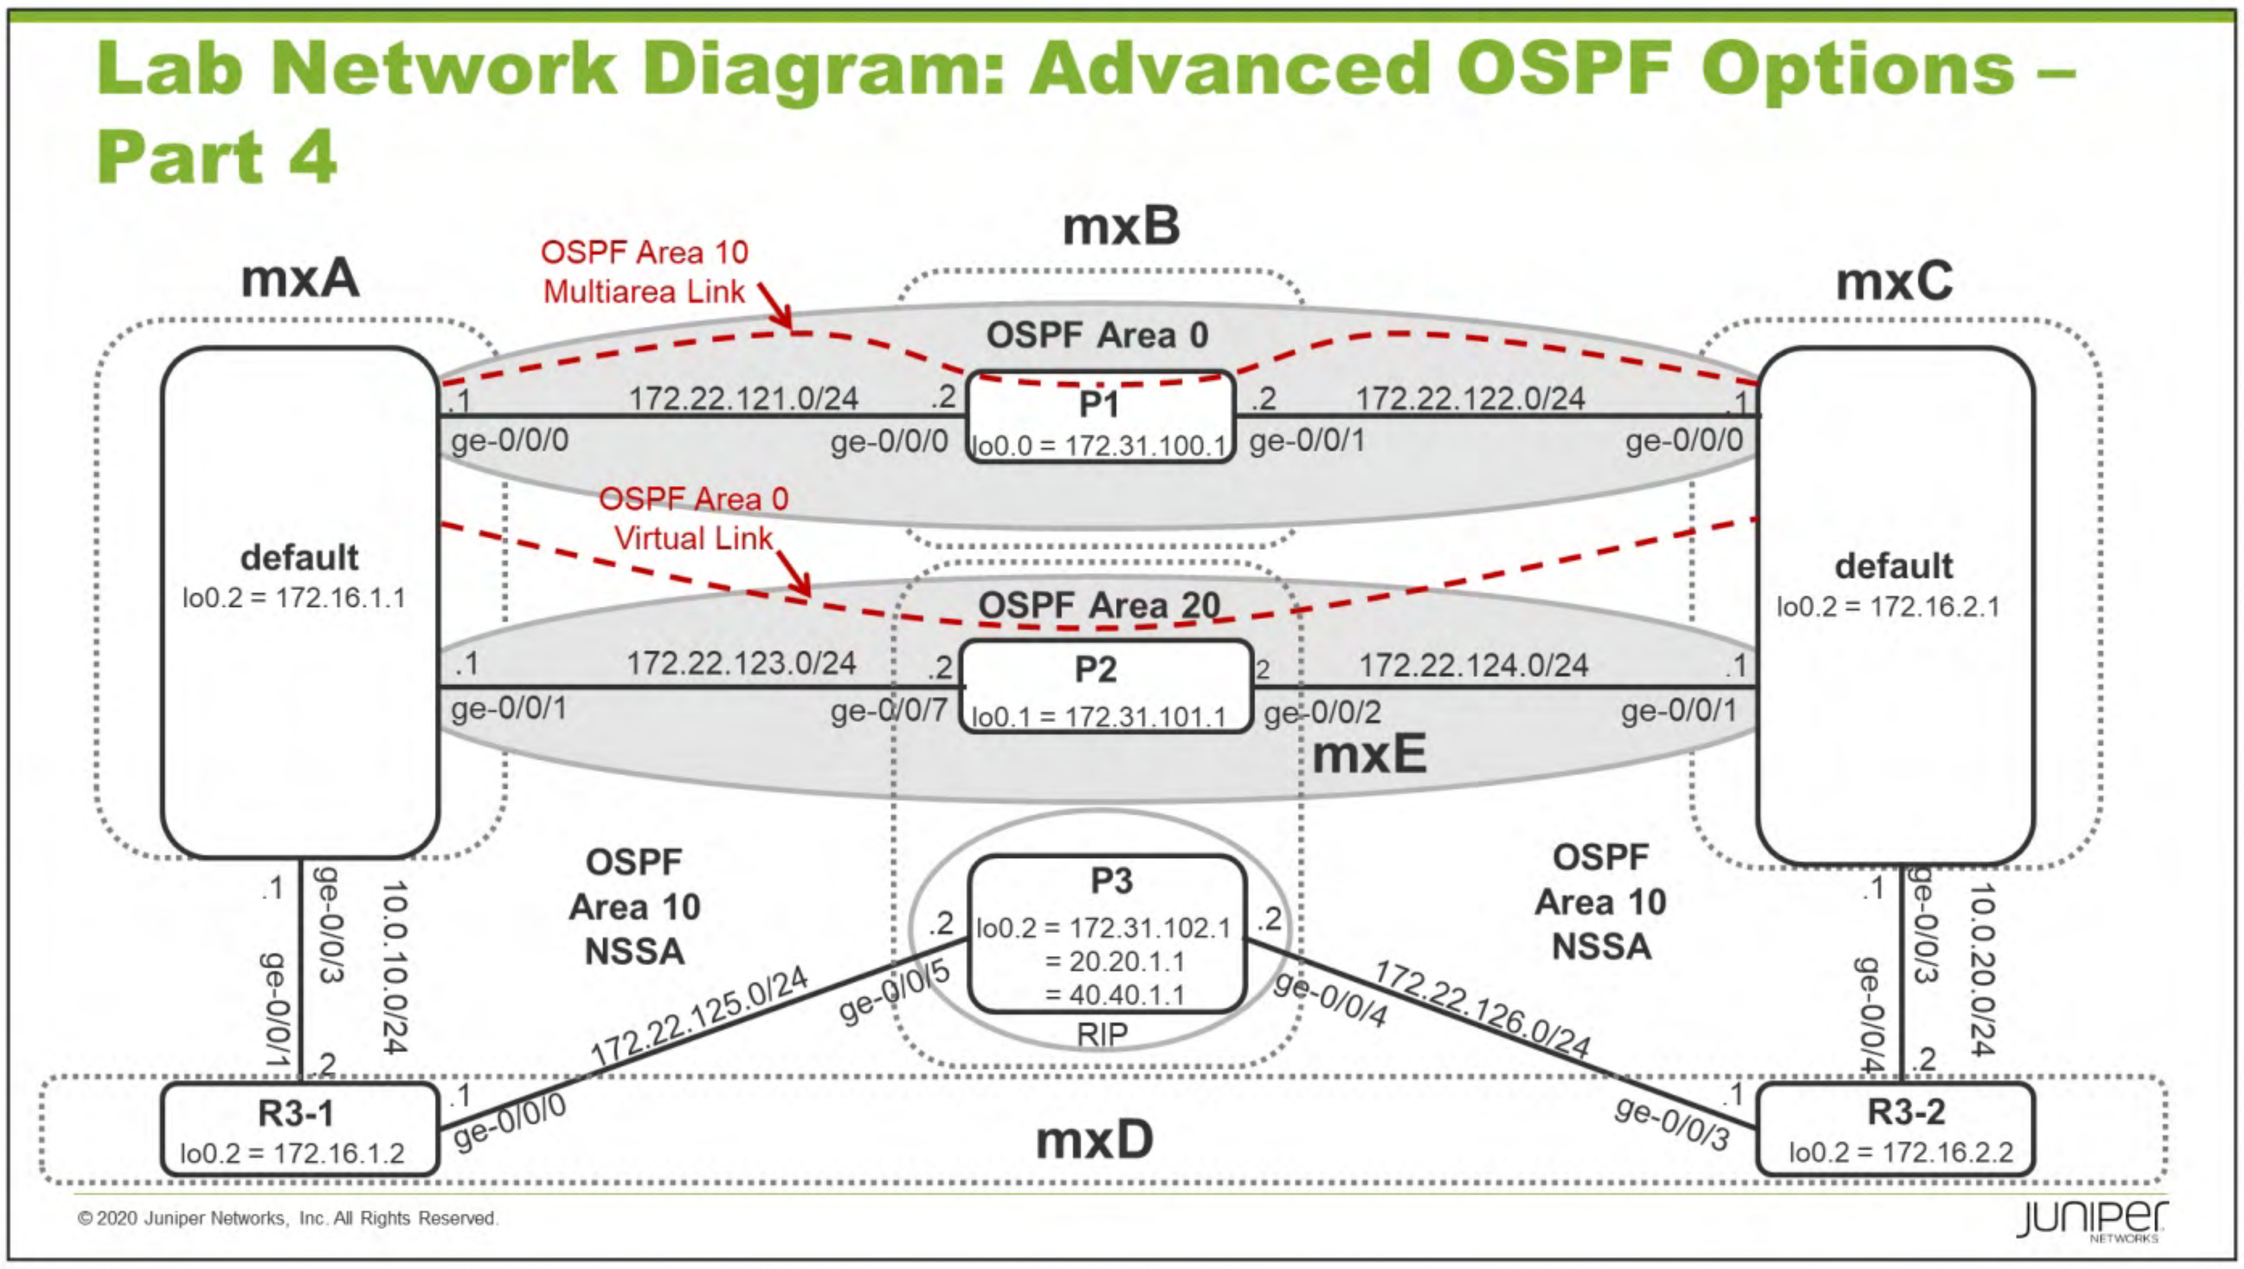
Lab Network Diagram: Advanced OSPF Options – Part 4
mxA
mxB
mxC
mxD
mxE
OSPF Area 10
Multiarea Link
OSPF Area 0
OSPF Area 0
Virtual Link
OSPF Area 20
default
lo0.2 = 172.16.1.1
default
lo0.2 = 172.16.2.1
P1
lo0.0 = 172.31.100.1
P2
lo0.1 = 172.31.101.1
P3
lo0.2 = 172.31.102.1
= 20.20.1.1
= 40.40.1.1
RIP
R3-1
lo0.2 = 172.16.1.2
R3-2
lo0.2 = 172.16.2.2
172.22.121.0/24
172.22.122.0/24
172.22.123.0/24
172.22.124.0/24
172.22.125.0/24
172.22.126.0/24
10.0.10.0/24
10.0.20.0/24
ge-0/0/0
ge-0/0/1
ge-0/0/2
ge-0/0/3
ge-0/0/4
ge-0/0/5
ge-0/0/7
OSPF
Area 10
NSSA
OSPF
Area 10
NSSA
© 2020 Juniper Networks, Inc. All Rights Reserved.
juniper NETWORKS

Lab

Troubleshooting OSPF

Overview

In this lab, you use the lab diagram titled “Lab Network Diagram: Troubleshooting OSPF” to troubleshoot different OSPF scenarios. The goal is to get all adjacencies up and make sure OSPF is distributing the interfaces of the endpoints, so that you will be able to send traffic through the backbone.

By completing this lab, you will perform the following tasks:

- Troubleshoot various OSPF peering issues.
- Make appropriate configuration changes to rectify peering problems.
- Verify and validate routes being shared and ensure that all routes are present.

Part 1: Load Starting Configuration

In this lab part, you load the initial configuration of the routers.

Note

The instructor will tell you the nature of your access and will provide you with the necessary details to access your assigned device.

Note

The lab topology requires you to display information in the different logical systems of mxB and mxD. You will be changing the perspective of the CLI by issuing the **`set cli logical-system`** ***`ls-name`*** command. Refer to the lab diagram for the correct logical system name.

By changing the perspective of the CLI, you will be able to perform network commands such as **`ping`** or **`traceroute`** from the perspective of the pertinent logical system.

Step 1.1

Ensure that you know to which student device you have been assigned. Consult the management network diagram to determine the management address of the student devices.

Access the CLI on your mxD device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab4-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.4
Password: lab123
Last login: Sun Nov 24 12:06:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab4-start.config
load complete

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

lab@mxD>
```

Step 1.2

Access the CLI on your mxC device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab4-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.3
Password: lab123
Last login: Sun Nov 24 11:45:27 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# load override ajspr/lab4-start.config
load complete

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

lab@mxC>
```

Step 1.3

Access the CLI on your mxB device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab4-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.2
Password: lab123
Last login: Sun Nov 24 11:45:09 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxB> configure
Entering configuration mode

[edit]
lab@mxB# load override ajspr/lab4-start.config
load complete

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

lab@mxB>
```

Step 1.4

Access the CLI on your mxA device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab4-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.1
Password: lab123
Last login: Sun Nov 24 11:45:09 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# load override ajspr/lab4-start.config
load complete

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

lab@mxA>
```

Part 2: Troubleshooting OSPF

In this lab part, you use the OSPF configuration to troubleshoot errors.

Step 2.1

Issue the **`show ospf interface detail`** command to view the current OSPF status for the mxA router. There should be an adjacency to the R2 router, and the interface connecting the mxA router to the Endpoint1 device should be set to passive.

```
lab@mxA> show ospf interface detail
Interface            State   Area            DR ID           BDR ID          Nbrs
ge-0/0/0.0           DR      0.0.0.2         192.168.71.1    0.0.0.0            0
  Type: LAN, Address: 172.22.131.1, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  DR addr: 172.22.131.1, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
lo0.0                DR      0.0.0.2         192.168.71.1    0.0.0.0            0
  Type: LAN, Address: 192.168.71.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  DR addr: 192.168.71.1, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 0
```

Question: Are you seeing the correct OSPF adjacencies? What seems to be the issue?

_____

_____

Answer: No, no OSPF adjacencies are currently active. The ge-0/0/0 interface has been put in area *2*, but should be in area *1*. Also the ge-0/0/3 interface is not listed.

Step 2.2

Use the **`show ospf statistics | find errors`** to view the current OSPF errors.

```
lab@mxA> show ospf statistics | find errors
Receive errors:
  4 area mismatches
```

Question: Are there any OSPF errors?

_____

Answer: Yes, as shown in the output, there should be some area mismatches. This confirms the previous findings of a mismatched area.

Step 2.3

Next, determine what the issue is with the ge-0/0/3 interface using the **`show configuration protocols ospf`** command.

```
lab@mxA> show configuration protocols ospf
area 0.0.0.2 {
    interface lo0.0;
    interface ge-0/0/0.0;
    interface ge-0/0/3.415;
}
```

Question: Is the OSPF configuration related to the ge-0/0/3 interface correct?

_____

Answer: No, the configuration is not correct for the ge-0/0/3 interface. It is currently using the logical unit of 415 and should be set to unit 0. Also the area is incorrect. The ge-0/0/3 interface as well as all other interfaces which should be in area 1, not area 2.

Question: What additional information is missing from the ge-0/0/3 interface?

_____

_____

Answer: The ge-0/0/3 interface connects the mxA router to the Endpoint1 device and should be operating in passive mode as per the guideline outlined in step 2.1.

Step 2.4

Enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Correct the issues identified in previous steps (wrong area ID, ge-0/0/3 unit value and ge-0/0/3 interface in passive mode). Commit your configuration and return to operational mode when complete.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit protocols ospf

[edit protocols ospf]
lab@mxA# show
area 0.0.0.2 {
    interface lo0.0;
    interface ge-0/0/0.0;
    interface ge-0/0/3.415;
}

[edit protocols ospf]
lab@mxA# rename area 2 to area 1

[edit protocols ospf]
lab@mxA# delete area 1 interface ge-0/0/3.415

[edit protocols ospf]
lab@mxA# set area 1 interface ge-0/0/3.0 passive

[edit protocols ospf]
lab@mxA# show
area 0.0.0.1 {
    interface lo0.0;
    interface ge-0/0/0.0;
    interface ge-0/0/3.0 {
        passive;
    }
}

[edit protocols ospf]
lab@mxA# commit and-quit
commit complete
Exiting configuration mode
```

```
lab@mxA>
```

Step 2.5

Issue the **`show ospf interface detail`** command again to see if the issue is resolved.

```
lab@mxA> show ospf interface detail
Interface            State   Area            DR ID            BDR ID           Nbrs
ge-0/0/0.0           BDR     0.0.0.1         192.168.71.2     192.168.71.1        1
  Type: LAN, Address: 172.22.131.1, Mask: 255.255.255.252, MTU: 1500, Cost: 1
  DR addr: 172.22.131.2, BDR addr: 172.22.131.1, Priority: 128
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 1
ge-0/0/3.0           DRother 0.0.0.1         0.0.0.0          0.0.0.0             0
  Type: LAN, Address: 172.22.131.38, Mask: 255.255.255.252, MTU: 1504, Cost: 1
  Adj count: 0, Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Passive, Cost: 1
lo0.0                DR      0.0.0.1         192.168.71.1     0.0.0.0             0
  Type: LAN, Address: 192.168.71.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  DR addr: 192.168.71.1, Priority: 128
  Adj count: 0
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Protection type: None
  Topology default (ID 0) -> Cost: 0
```

Question: Have the OSPF issues been resolved?

_____

_____

Answer: Yes, the OSPF issues should now be resolved.

Step 2.6

Issue the **`show ospf neighbor`** command to view the status of the OSPF neighbor relationships.

```
lab@mxA> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.131.2     ge-0/0/0.0             Full      192.168.71.2     128    32
```

Question: Has the mxA router established a full OSPF neighbor relationship with the R2 router?

_____

Answer: Yes, the OSPF neighbor relationship with the R2 router should be in the `Full` state.

Step 2.7

Return to the SSH session that you have open to the mxB router.

From the SSH session with the mxB router, change the perspective of the CLI to that of the R2 logical system by issuing the **`set cli logical-system R2`** command.

```
lab@mxB> set cli logical-system R2
Logical system: R2

lab@mxB:R2>
```

Step 2.8

Issue the **`show ospf interface`** command to view the interfaces that have been configured for OSPF.

```
lab@mxB:R2> show ospf interface
Interface          State   Area          DR ID          BDR ID         Nbrs
ge-0/0/2.0         DR      0.0.0.0       192.168.71.2   0.0.0.0           1
lo0.2              DR      0.0.0.0       192.168.71.2   0.0.0.0           0
lt-0/0/10.0        DR      0.0.0.0       192.168.71.2   0.0.0.0           0
ge-0/0/0.0         DR      0.0.0.1       192.168.71.2   192.168.71.1      1
```

Question: What issues can you see?

_____

_____

Answer: From the output, it looks like OSPF is running on the correct interfaces. But only two neighbor have been detected. Three neighbor should be detected.

Step 2.9

Issue the **`show ospf neighbor`** command to view the status of the OSPF neighbor relationships.

```
lab@mxB:R2> show ospf neighbor
Address          Interface              State     ID              Pri  Dead
172.22.131.10    ge-0/0/2.0             Init      192.168.71.3    128    34
172.22.131.1     ge-0/0/0.0             Full      192.168.71.1    128    36
```

Question: What issues can you see?

_____

_____

Answer: From the output, only one neighbor is in the `Full` state and the relationship with R3 is in `Init`, which usually means we also need to check the remote side for errors.

Step 2.10

Clear the OSPF statistics.

```
lab@mxB:R2> clear ospf statistics
```

Step 2.11

Wait one minute and check the error counters by issuing the **`show ospf statistics | find errors`**.

```
lab@mxB:R2> show ospf statistics | find errors
Receive errors:
  14 netmask mismatches
```

Question: Did the counter increase?

_____

Answer: Yes, the counter has increased and looks to be a network mismatch.

Step 2.12

Check the interfaces configuration with the **`show interfaces terse | match inet | except :`** command, and compare it with the diagram.

```
lab@mxB:R2> show interfaces terse | match inet | except :
ge-0/0/0.0              up    up   inet     172.22.131.2/30
ge-0/0/2.0              up    up   inet     172.22.131.9/30
lt-0/0/10.0             up    up   inet     172.22.131.5/24
lo0.2                   up    up   inet     192.168.71.2        --> 0/0
```

Question: Are all the interfaces configured correctly?

_____

Answer: No, not all the interfaces are configured correctly. The lt-0/0/0 interface has the wrong mask configured. It should be using a /30 mask, but instead has been configured with a /24.

Step 2.13

Enter configuration mode and navigate to the `[edit interfaces lt-0/0/0]`. Make the necessary configuration changes to alter the subnet mask applied to the lt-0/0/0 interface. Commit your configuration and return to operational mode.

```
lab@mxB:R2> configure
Entering configuration mode

[edit]
lab@mxB:R2# edit interfaces lt-0/0/10

[edit interfaces lt-0/0/10]
lab@mxB:R2# show
unit 0 {
    encapsulation ethernet;
    peer-unit 1;
    family inet {
        address 172.22.131.5/24;
    }
    family iso;
    family inet6;
```

```
    family mpls;
}

[edit interfaces lt-0/0/10]
lab@mxB:R2# delete unit 0 family inet address 172.22.131.5/24

[edit interfaces lt-0/0/10]
lab@mxB:R2# set unit 0 family inet address 172.22.131.5/30

[edit interfaces lt-0/0/10]
lab@mxB:R2# commit and-quit
commit complete
Exiting configuration mode

lab@mxB:R2>
```

Step 2.14

Use the **show ospf neighbor** command to determine if OSPF is now coming up on lt-0/0/0.

```
lab@mxB:R2> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.131.10    ge-0/0/2.0             Init      192.168.71.3     128    37
172.22.131.6     lt-0/0/10.0            Full      192.168.71.4     128    37
172.22.131.1     ge-0/0/0.0             Full      192.168.71.1     128    34
```

Question: Did the mask change fix the issue with the OSPF neighbor out lt-0/0/0?

_____

_____

Answer: Yes, the change fixed the issue and the neighbor adjacency is now established.

Step 2.15

Change the perspective of the CLI to that of the R3 logical system by issuing the **set cli logical-system *R3*** command.

```
lab@mxB:R2> set cli logical-system R3
Logical system: R3

lab@mxB:R3>
```

Step 2.16

Issue the **show ospf interface** command to view the interfaces that have been configured for OSPF.

```
lab@mxB:R3> show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/3.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            0
ge-0/0/4.0         BDR     0.0.0.0         192.168.71.4    192.168.71.3       1
lo0.3              DR      0.0.0.0         192.168.71.3    0.0.0.0            0
```

Question: Is OSPF running on the correct interfaces?

_____

_____

Answer: Yes, OSPF is running on the correct interfaces.

Step 2.17

Issue the **show ospf neighbor** command to view the status of the OSPF neighbor relationships.

```
lab@mxB:R3> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.131.14    ge-0/0/4.0             Full      192.168.71.4     128    39
```

Question: Are both OSPF neighbors up?

_____

_____

Answer: No. Only one of the two neighbors are up at this time.

Step 2.18

Check the OSPF errors using the **show ospf statistics | find error** command.

```
lab@mxB:R3> show ospf statistics | find errors
Receive errors:
  36 Hellos received on point-to-point LAN with DR/BDR elected
```

Question: What seems to be the issue?

_____

_____

Answer: Checking the configuration shows the problem. One of the interfaces has been configured as point-to-point.

Step 2.19

Issue the show configuration protocols ospf command to view the OSPF configuration of R3.

```
lab@mxB:R3> show configuration protocols ospf
area 0.0.0.0 {
    interface ge-0/0/3.0 {
        interface-type p2p;
    }
    interface ge-0/0/4.0;
    interface lo0.3;
}
```

Question: Have any interfaces been configured as point-to-point OSPF interfaces?

_____

_____

Answer: Yes, the ge-0/0/3.0 interfaces has been configured as a point-to-point link.

Step 2.20

Enter configuration mode and navigate to the `[edit protocols ospf]` hierarchy. Remove the `p2p` configuration from the R3 side. Commit the configuration and return to operational mode.

```
lab@mxB:R3> configure
Entering configuration mode

[edit]
lab@mxB:R3# edit protocols ospf

[edit protocols ospf]
lab@mxB:R3# delete area 0 interface ge-0/0/3.0 interface-type p2p

[edit protocols ospf]
lab@mxB:R3# commit and-quit
commit complete
Exiting configuration mode

lab@mxB:R3>
```

Step 2.21

Use the **`show ospf neighbor`** command to verify that both the neighbors are now established.

```
lab@mxB:R3> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
172.22.131.9     ge-0/0/3.0             Full      192.168.71.2     128    38
172.22.131.14    ge-0/0/4.0             Full      192.168.71.4     128    39
```

Question: Are both the neighbors established?

_____

_____

Answer: Yes, both the neighbors are now established.

Step 2.22

Change the perspective of the CLI to that of the R4 logical system by issuing the **`set cli logical-system R4`** command.

```
lab@mxB:R3> set cli logical-system R4
Logical system: R4

lab@mxB:R4>
```

Step 2.23

Issue the **show ospf interface** command to view the interfaces that have been configured for OSPF.

```
lab@mxB:R4> show ospf interface
Interface          State   Area            DR ID           BDR ID          Nbrs
ge-0/0/5.0         DR      0.0.0.0         192.168.71.4    192.168.71.3       1
lo0.4              DR      0.0.0.0         192.168.71.4    0.0.0.0            0
lt-0/0/10.1        DR      0.0.0.0         192.168.71.4    192.168.71.2       1
ge-0/0/1.0         BDR     0.0.0.2         192.168.71.5    192.168.71.4       1
```

Question: Is OSPF running on the correct interfaces?

_____

_____

Answer: Yes, OSPF is running on the correct interfaces.

Step 2.24

Issue the **show ospf neighbor** command to view the status of the OSPF neighbor relationships.

```
lab@mxB:R4> show ospf neighbor
Address          Interface              State     ID              Pri  Dead
172.22.131.13    ge-0/0/5.0             Full      192.168.71.3    128    37
172.22.131.5     lt-0/0/10.1            Full      192.168.71.2    128    39
172.22.131.18    ge-0/0/1.0             Full      192.168.71.5    128    37
```

Question: Are all three OSPF neighbors up?

_____

_____

Answer: Yes, all three of the neighbors are up.

Step 2.25

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, see if you can reach the Endpoint2 device using traceroute.

```
lab@mxA> traceroute 172.22.131.42
traceroute to 172.22.131.42 (172.22.131.42), 30 hops max, 52 byte packets
 1  172.22.131.2 (172.22.131.2)  1.662 ms  1.270 ms  1.298 ms
 2  172.22.131.6 (172.22.131.6)  2.975 ms  1.330 ms  1.493 ms
 3  172.22.131.18 (172.22.131.18)  40.815 ms  2.093 ms  1.876 ms
 4  172.22.131.42 (172.22.131.42)  3.825 ms  3.587 ms  4.308 ms
```

Question: Can you reach the Endpoint2 device? Does this indicate that OSPF is properly configured?

_____

_____

Answer: Yes, you should be able to reach the Endpoint2 device from the mxA router. No, there may still be some issues and you should check to see if you can see routes for each routers loopback address.

Step 2.26

Use the **show ospf route 192.168/16** to see if all the routers loopback addresses are listed.

```
lab@mxA> show ospf route 192.168/16
Topology default Route Table:

Prefix              Path  Route      NH       Metric NextHop       Nexthop
                    Type  Type       Type             Interface     Address/LSP
192.168.71.2        Intra Area BR    IP            1 ge-0/0/0.0    172.22.131.2
192.168.71.1/32     Intra Network    IP            0 lo0.0
192.168.71.2/32     Inter Network    IP            1 ge-0/0/0.0    172.22.131.2
192.168.71.3/32     Inter Network    IP            2 ge-0/0/0.0    172.22.131.2
192.168.71.4/32     Inter Network    IP            2 ge-0/0/0.0    172.22.131.2
```

Question: Are all the routers' loopback addresses listed?

_____

_____

Answer: No, the loopback address from the mxC router is missing.

Step 2.27

Check in the OSPF database to see if you can see a network summary LSA generated by the R2 router for the loopback of the mxC router

```
lab@mxA> show ospf database netsummary

OSPF database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Summary  172.22.131.4     192.168.71.2     0x80000002   150  0x22 0xd769  28
Summary  172.22.131.8     192.168.71.2     0x80000003    81  0x22 0xad8e  28
Summary  172.22.131.12    192.168.71.2     0x80000002    81  0x22 0x91a6  28
Summary  172.22.131.63    192.168.71.2     0x80000001   150  0x22 0x340d  28
Summary  192.168.71.2     192.168.71.2     0x80000002   676  0x22 0xa630  28
Summary  192.168.71.3     192.168.71.2     0x80000002    81  0x22 0xa62e  28
Summary  192.168.71.4     192.168.71.2     0x80000001   150  0x22 0x9e36  28
```

Question: Is there a summary for the mxC router's loopback in the `Area 0.0.0.1` database?

_____

_____

Answer: No, we seem to be missing a summary for the loopback of the mxC router.

Step 2.28

Return to the SSH session that you have open to the mxB router.

From the SSH session with the mxB router, change the perspective of the CLI to that of the R2 logical system by issuing the **`set cli logical-system R2`** command.

```
lab@mxB:R4> set cli logical-system R2
Logical system: R2

lab@mxB:R2>
```

Step 2.29

Check the area 0 database using the **`show ospf database area 0 netsummary`** command.

```
lab@mxB:R2> show ospf database area 0 netsummary

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Summary *172.22.131.0     192.168.71.2     0x80000024   669  0x22 0xbb67  28
Summary  172.22.131.0     192.168.71.4     0x80000023  1246  0x22 0x520b  28
Summary *172.22.131.36    192.168.71.2     0x80000002  1113  0x22 0xa07f  28
Summary *192.168.71.1     192.168.71.2     0x80000002   883  0x22 0xba1c  28
```

Question: Is there a summary LSA for the loopback of the mxC router in area 0?

_____

_____

Answer: No, it looks like the R4 router is not generating a route for the loopback of the mxC router.

Step 2.30

Change the perspective of the CLI to that of the R4 router by issuing the **`set cli logical-system R4`** command.

```
lab@mxB:R2> set cli logical-system R4
Logical system: R4

lab@mxB:R4>
```

Step 2.31

Issue the **`show ospf database area 2 router`** command to check if you have a router LSA for the mxC router.

```
lab@mxB:R4> show ospf database area 2 router

    OSPF database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router  *192.168.71.4     192.168.71.4     0x80000022   815  0x22 0x7a7d  36
Router   192.168.71.5     192.168.71.5     0x80000022  2165  0x22 0x575c  60
```

Question: Is there a router LSA for the mxC router?

_____

_____

Answer: Yes, there is a router LSA for the mxC router.

Step 2.32

Examine in detail the router LSA for the mxC router (be sure to use the correct LSA-id), and make sure that it includes a stub network for its loopback, with netmask 255.255.255.255.

```
lab@mxB:R4> show ospf database area 2 lsa-id 192.168.71.5 extensive

    OSPF database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   192.168.71.5     192.168.71.5     0x80000022  2241  0x22 0x575c  60
  bits 0x0, link count 3
  id 172.22.131.18, data 172.22.131.18, Type Transit (2)
    Topology count: 0, Default metric: 1
  id 172.22.131.40, data 255.255.255.252, Type Stub (3)
    Topology count: 0, Default metric: 1
  id 192.168.71.5, data 255.255.255.255, Type Stub (3)
    Topology count: 0, Default metric: 0
  Topology default (ID 0)
    Type: Transit, Node ID: 172.22.131.18
      Metric: 1, Bidirectional
  Aging timer 00:22:38
  Installed 00:37:18 ago, expires in 00:22:39
  Last changed 1d 00:47:21 ago, Change count: 2
```

Question: Does the router LSA of the mxC router contain its loopback address?

_____

_____

Answer: Yes, the router LSA contains the loopback address of the mxC router.

Step 2.33

Review the OSPF configuration on the R4 router by issuing the **`show configuration protocols ospf`** command.

```
lab@mxB:R4> show configuration protocols ospf
area 0.0.0.0 {
    interface ge-0/0/5.0;
    interface lt-0/0/10.1;
```

```
    interface lo0.4;
}
area 0.0.0.2 {
    area-range 172.22.131.0/26;
    area-range 0.0.0.0/0 restrict;
    interface ge-0/0/1.0;
}
```

Question: What is the issue with the configuration on the R4 router?

_____

_____

Answer: The problem is the area-range 0/0 restrict; this was probably meant to make sure that the router generates only one single /26 LSA for the whole area. But unfortunately the loopback does not belong to this aggregate, so it is filtered by the '0/0 restrict'.

Step 2.34

Enter configuration mode and allow a /32 for the loopback address of the mxC router. Commit your configuration and return to operational mode when complete.

```
lab@mxB:R4> configure
Entering configuration mode

[edit]
lab@mxB:R4# set protocols ospf area 2 area-range 192.168.71.5/32

[edit]
lab@mxB:R4# show protocols ospf
area 0.0.0.0 {
    interface ge-0/0/5.0;
    interface lt-0/0/10.1;
    interface lo0.4;
}
area 0.0.0.2 {
    area-range 172.22.131.0/26;
    area-range 0.0.0.0/0 restrict;
    area-range 192.168.71.5/32;
    interface ge-0/0/1.0;
}

[edit]
lab@mxB:R4# commit and-quit
commit complete
Exiting configuration mode

lab@mxB:R4>
```

Step 2.35

Finally, verify that a netsummary LSA for the loopback of the mxC router is now present in the area 0 database.

```
lab@mxB:R4> show ospf database area 0 netsummary

    OSPF database, Area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Summary  172.22.131.0     192.168.71.2     0x80000024  1210  0x22 0xbb67  28
Summary *172.22.131.0     192.168.71.4     0x80000024    30  0x22 0x500c  28
Summary  172.22.131.36    192.168.71.2     0x80000002  1654  0x22 0xa07f  28
Summary  192.168.71.1     192.168.71.2     0x80000002  1424  0x22 0xba1c  28
Summary *192.168.71.5     192.168.71.4     0x80000001    30  0x22 0x8849  28
```

Step 2.36

Log out of your assigned devices using the **exit** command.

```
lab@mxB:R4> exit

Connection to 172.25.11.2 closed
```

Tell your instructor that you have completed this lab.

Network Management Diagram

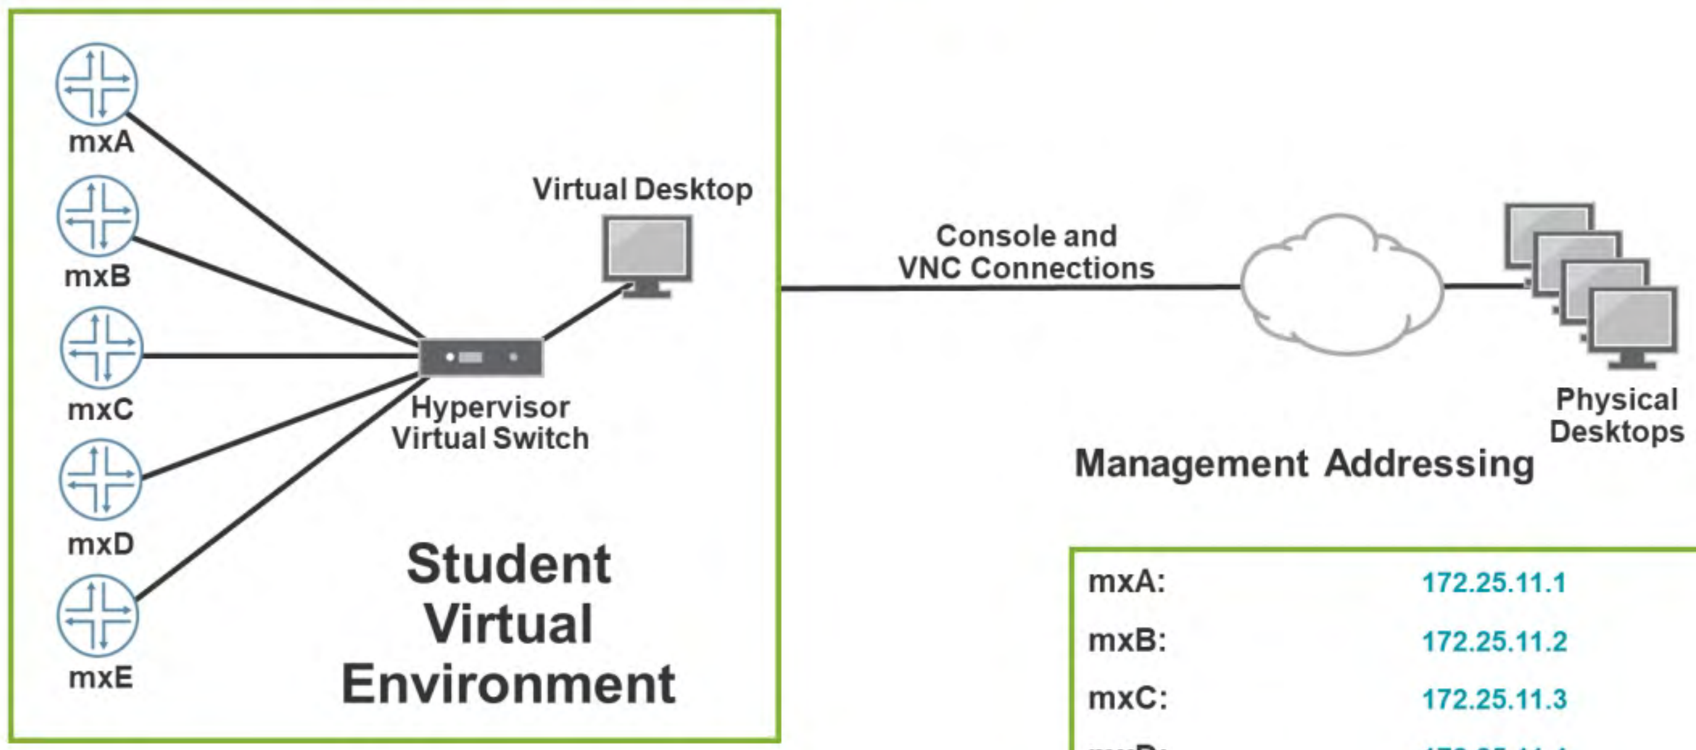

Management Addressing

mxA:	172.25.11.1
mxB:	172.25.11.2
mxC:	172.25.11.3
mxD:	172.25.11.4
mxE:	172.25.11.5
Virtual Desktop:	172.25.11.254

Note: Your instructor will provide address and access information.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab Network Diagram: Troubleshooting OSPF

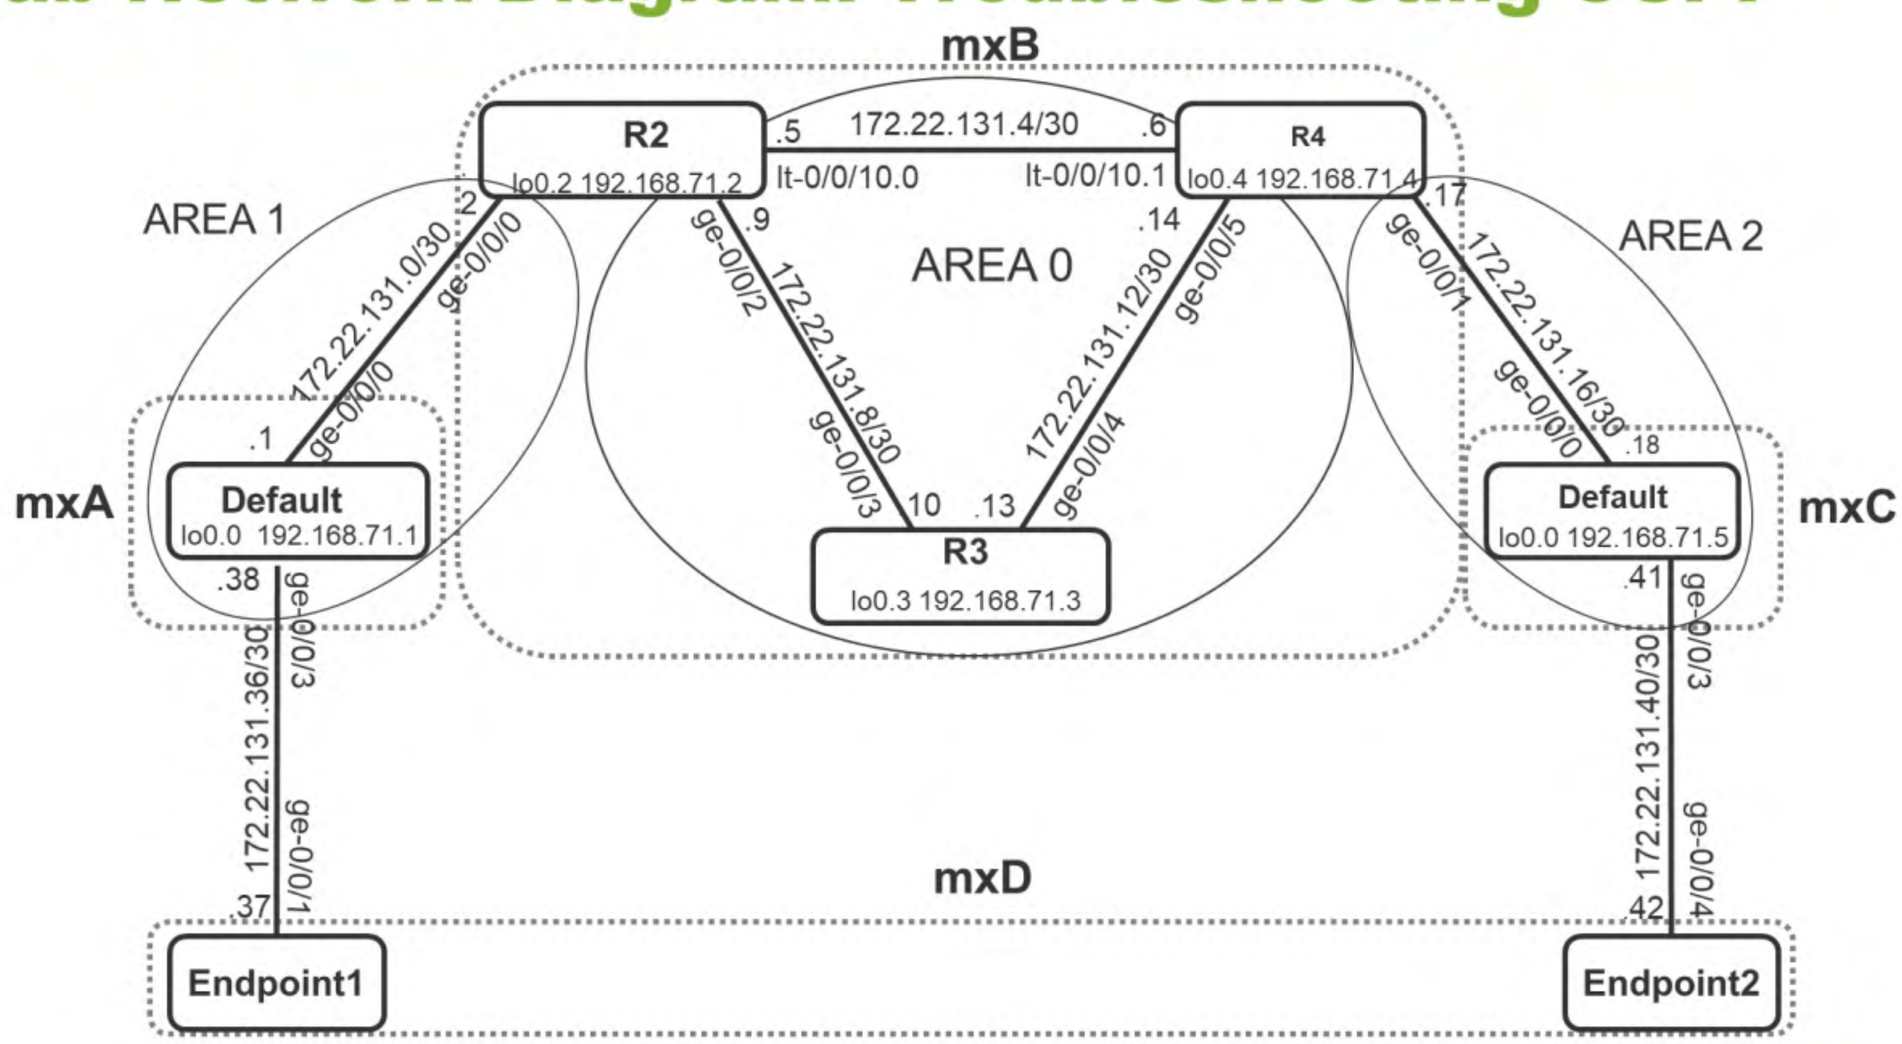

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab

IS-IS Configuration and Monitoring

Overview

In this lab, you use the lab diagram titled “Lab Network Diagram: IS-IS Configuration and Monitoring” to establish an IS-IS routing domain.

By completing this lab, you will perform the following tasks:

- Configure support of the ISO protocol data unit (PDU).
- Configure the IS-IS Network Entity Title (NET).
- Establish IS-IS adjacencies.
- Monitor IS-IS interfaces and adjacencies.

Part 1: Load Starting Configuration

In this lab part, you verify the initial configuration of the routers. You then verify that the interfaces are operational.

Note

The instructor will tell you the nature of your access and will provide you with the necessary details to access your assigned device.

Note

The lab topology requires you to display information in the different logical systems of mxD and mxE. You will be changing the perspective of the CLI by issuing the **`set cli logical-system`** ***`ls-name`*** command. Refer to the lab diagram for the correct logical system name.

By changing the perspective of the CLI, you will be able to perform network commands such as **`ping`** or **`traceroute`** from the perspective of the pertinent logical system.

Step 1.1

Ensure that you know to which student device you have been assigned. Check with your instructor if you are not certain. Consult the management network diagram to determine the management address of the student devices.

Step 1.2

Access the CLI on your mxD device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab5-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.4
Password: lab123
Last login: Tue Dec  3 14:41:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab5-start.config
load complete

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

lab@mxD>
```

Step 1.3

Issue the **show configuration** command. Use the lab diagram to verify that the mxD router has two logical systems, R3-1 and R3-2, that have the correct interface configuration.

```
lab@mxD> show configuration
system {
    host-name mxD;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    R3-1 {
        interfaces {
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 10.0.10.2/24;
                    }
                }
            }
            lo0 {
                unit 1 {
                    family inet {
                        address 172.16.1.2/32;
```

```
                    }
                }
            }
        }
    }
    R3-2 {
        interfaces {
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 10.0.14.2/24;
                    }
                }
            }
            lo0 {
                unit 2 {
                    family inet {
                        address 172.16.2.2/32;
                    }
                }
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.4/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxD router. If the configuration is incorrect, please notify your instructor.

Step 1.4

Access the CLI on your mxA device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab5-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.1
Password: lab123
Last login: Tue Dec  3 14:42:16 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
```

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# load override ajspr/lab5-start.config
load complete

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

lab@mxA>
```

Step 1.5

Issue the **show configuration** command. Use the lab diagram to verify that the mxA router has the correct interface configuration. Verify that there are four 20.20/24 static routes and a policy statement.

```
lab@mxA> show configuration
## Last commit: 2017-05-02 11:17:50 UTC by lab
version 17.1R1.8;
system {
    host-name mxA;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
```

```
        }
    }
}
interfaces {
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.22.121.1/24;
            }
        }
    }
    ge-0/0/2 {
        unit 0 {
            family inet {
                address 172.22.123.1/24;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.10.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.1/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.1.1/32;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxA router. If the configuration is incorrect, please notify your instructor.

Step 1.6

Use the **ping 10.0.10.2 count 5** command to verify that you can ping the physical interface on the R3-1 router.

```
lab@mxA> ping 10.0.10.2 count 5
PING 10.0.10.2 (10.0.10.2): 56 data bytes
64 bytes from 10.0.10.2: icmp_seq=0 ttl=64 time=37.896 ms
64 bytes from 10.0.10.2: icmp_seq=1 ttl=64 time=1.756 ms
64 bytes from 10.0.10.2: icmp_seq=2 ttl=64 time=163.903 ms
64 bytes from 10.0.10.2: icmp_seq=3 ttl=64 time=150.237 ms
64 bytes from 10.0.10.2: icmp_seq=4 ttl=64 time=4.449 ms

--- 10.0.10.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.756/71.648/163.903/71.031 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.7

Access the CLI on your mxE device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device’s starting configuration by issuing the **load override *ajspr/lab5-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.5
Password: lab123
Last login: Tue Dec  3 14:41:21 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxE> configure
Entering configuration mode

[edit]
lab@mxE# load override ajspr/lab5-start.config
load complete

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

lab@mxE>
```

Step 1.8

Issue the **show configuration** command. Use the lab diagram to verify that the mxE router has the correct interface configuration. Verify that OSPF has been enabled (all interfaces in area 0).

```
lab@mxE> show configuration
## Last commit: 2017-05-02 11:35:18 UTC by lab
version 17.1R1.8;
system {
    host-name mxE;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    P1 {
        interfaces {
            ge-0/0/2 {
                unit 0 {
                    family inet {
                        address 172.22.122.2/24;
                    }
                }
            }
            ge-0/0/7 {
                unit 0 {
                    family inet {
                        address 172.22.121.2/24;
                    }
                }
```

```
            }
            lt-0/0/10 {
                unit 0 {
                    encapsulation ethernet;
                    peer-unit 1;
                    family inet {
                        address 172.22.252.1/30;
                    }
                }
            }
            lo0 {
                unit 1 {
                    family inet {
                        address 172.31.100.1/32;
                    }
                }
            }
        }
        routing-options {
            autonomous-system 65412;
        }
    }
    P2 {
        interfaces {
            ge-0/0/3 {
                unit 0 {
                    family inet {
                        address 172.22.124.2/24;
                    }
                }
            }
            ge-0/0/6 {
                unit 0 {
                    family inet {
                        address 172.22.123.2/24;
                    }
                }
            }
            lt-0/0/10 {
                unit 1 {
                    encapsulation ethernet;
                    peer-unit 0;
                    family inet {
                        address 172.22.252.2/30;
                    }
                }
            }
            lo0 {
                unit 2 {
                    family inet {
                        address 172.31.101.1/32;
                    }
                }
            }
        }
```

```
        routing-options {
            autonomous-system 65412;
        }
    }
}
chassis {
    fpc 0 {
        pic 0 {
            tunnel-services {
                bandwidth 1g;
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.5/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxE router. If the configuration is incorrect, please notify your instructor.

Step 1.9

Use the **`ping 172.22.121.1 count 5 logical-system P1`** command to verify that you can ping the ge-0/0/1 interface on the mxA router.

```
lab@mxE> ping 172.22.121.1 count 5 logical-system P1
PING 172.22.121.1 (172.22.121.1): 56 data bytes
64 bytes from 172.22.121.1: icmp_seq=0 ttl=64 time=1.687 ms
64 bytes from 172.22.121.1: icmp_seq=1 ttl=64 time=2.577 ms
64 bytes from 172.22.121.1: icmp_seq=2 ttl=64 time=1.616 ms
64 bytes from 172.22.121.1: icmp_seq=3 ttl=64 time=1.408 ms
64 bytes from 172.22.121.1: icmp_seq=4 ttl=64 time=130.031 ms

--- 172.22.121.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.408/27.464/130.031/51.285 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.10

Use the **`ping 172.22.252.2 count 5 logical-system P1`** command to verify that you can ping the lt-0/0/10.1 interface on the P2 router.

```
lab@mxE> ping 172.22.252.2 count 5 logical-system P1
PING 172.22.252.2 (172.22.252.2): 56 data bytes
64 bytes from 172.22.252.2: icmp_seq=0 ttl=64 time=0.986 ms
64 bytes from 172.22.252.2: icmp_seq=1 ttl=64 time=0.906 ms
64 bytes from 172.22.252.2: icmp_seq=2 ttl=64 time=1.058 ms
64 bytes from 172.22.252.2: icmp_seq=3 ttl=64 time=69.544 ms
64 bytes from 172.22.252.2: icmp_seq=4 ttl=64 time=1.292 ms

--- 172.22.252.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.906/14.757/69.544/27.394 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.11

Use the **`ping 172.22.123.1 count 5 logical-system P2`** command to verify that you can ping the ge-0/0/2 interface on the mxA router.

```
lab@mxE> ping 172.22.123.1 count 5 logical-system P2
PING 172.22.123.1 (172.22.123.1): 56 data bytes
64 bytes from 172.22.123.1: icmp_seq=0 ttl=64 time=57.777 ms
64 bytes from 172.22.123.1: icmp_seq=1 ttl=64 time=1.492 ms
64 bytes from 172.22.123.1: icmp_seq=2 ttl=64 time=45.699 ms
64 bytes from 172.22.123.1: icmp_seq=3 ttl=64 time=102.426 ms
64 bytes from 172.22.123.1: icmp_seq=4 ttl=64 time=4.075 ms

--- 172.22.123.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.492/42.294/102.426/37.397 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.12

Access the CLI on your mxC device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123**.

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab5-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.3
Last login: Mon Apr 24 12:48:15 2017 from 172.25.11.254
--- JUNOS 17.1R1.8 Kernel 64-bit  JNPR-10.3-20170209.344539_build
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# load override ajspr/lab5-start.config
load complete

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

lab@mxC>
```

Step 1.13

Issue the **show configuration** command. Use the lab diagram to verify that the mxC router has the correct interface configuration.

```
lab@mxC> show configuration
## Last commit: 2017-05-02 11:18:57 UTC by lab
version 17.1R1.8;
system {
    host-name mxC;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
```

```
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.22.122.1/24;
            }
        }
    }
    ge-0/0/2 {
        unit 0 {
            family inet {
                address 172.22.124.1/24;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.14.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.3/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.2.1/32;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxC router. If the configuration is incorrect, please notify your instructor.

Step 1.14

Use the **ping 10.0.14.2 count 5** command to verify that you can ping the physical interface on the R3-2 router.

```
lab@mxC> ping 10.0.14.2 count 5
PING 10.0.14.2 (10.0.14.2): 56 data bytes
64 bytes from 10.0.14.2: icmp_seq=0 ttl=64 time=26.747 ms
64 bytes from 10.0.14.2: icmp_seq=1 ttl=64 time=1.515 ms
64 bytes from 10.0.14.2: icmp_seq=2 ttl=64 time=122.310 ms
64 bytes from 10.0.14.2: icmp_seq=3 ttl=64 time=1.777 ms
64 bytes from 10.0.14.2: icmp_seq=4 ttl=64 time=1.512 ms

--- 10.0.14.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.512/30.772/122.310/46.794 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.15

Use the **ping 172.22.122.2 count 5** command to verify that you can ping the physical interface on the P1 router.

```
lab@mxC> ping 172.22.122.2 count 5
PING 172.22.122.2 (172.22.122.2): 56 data bytes
64 bytes from 172.22.122.2: icmp_seq=0 ttl=64 time=1.745 ms
64 bytes from 172.22.122.2: icmp_seq=1 ttl=64 time=151.563 ms
64 bytes from 172.22.122.2: icmp_seq=2 ttl=64 time=1.497 ms
64 bytes from 172.22.122.2: icmp_seq=3 ttl=64 time=1.592 ms
64 bytes from 172.22.122.2: icmp_seq=4 ttl=64 time=1.845 ms

--- 172.22.122.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.497/31.648/151.563/59.957 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.16

Use the **`ping 172.22.124.2 count 5`** command to verify that you can ping the physical interface on the P2 router.

```
lab@mxC> ping 172.22.124.2 count 5
PING 172.22.124.2 (172.22.124.2): 56 data bytes
64 bytes from 172.22.124.2: icmp_seq=0 ttl=64 time=207.688 ms
64 bytes from 172.22.124.2: icmp_seq=1 ttl=64 time=1.793 ms
64 bytes from 172.22.124.2: icmp_seq=2 ttl=64 time=1.523 ms
64 bytes from 172.22.124.2: icmp_seq=3 ttl=64 time=1.643 ms
64 bytes from 172.22.124.2: icmp_seq=4 ttl=64 time=2.095 ms

--- 172.22.124.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.523/42.948/207.688/82.370 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Part 2: Configuring the Transit Interfaces to Support ISO Packets

In this lab part, you load the reset configuration and configure family iso on all transit interfaces that could support IS-IS PDUs. The ingress I/O manager verifies Layer 3 packet headers based on the configuration provided for each logical interface including IPv4 and IPv6 protocols. Because the IS-IS PDU is not an IPv4 or IPv6 packet, the hardware must be configured to process this ISO PDU.

Step 2.1

Enter configuration mode and navigate to the `[edit groups ISO]` hierarchy. Using wildcard commands, configure all transit interfaces (`interface <*-*>`) to support the IS-IS family iso on all logical units (`unit <*>`).

```
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# edit groups ISO

[edit groups ISO]
lab@mxC# set interfaces <*-*> unit <*> family iso

[edit groups ISO]
lab@mxC#
```

Step 2.2

Navigate to the `[edit interfaces]` hierarchy and apply the ***ISO*** group to all interfaces. Commit the configuration when completed.

```
[edit groups ISO]
lab@mxC# top edit interfaces

[edit interfaces]
lab@mxC# set apply-groups ISO

[edit interfaces]
lab@mxC# commit
commit complete

[edit interfaces]
lab@mxC#
```

Step 2.3

Use the **`show | display inheritance`** command to verify that `family iso` has been inherited by all transit logical interfaces.

```
[edit interfaces]
lab@mxC# show | display inheritance
ge-0/0/1 {
    unit 0 {
        family inet {
            address 172.22.122.1/24;
        }
        ##
```

```
        ## 'iso' was inherited from group 'ISO'
        ##
        family iso;
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 172.22.124.1/24;
        }
        ##
        ## 'iso' was inherited from group 'ISO'
        ##
        family iso;
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 10.0.14.1/24;
        }
        ##
        ## 'iso' was inherited from group 'ISO'
        ##
        family iso;
    }
}
fxp0 {
    unit 0 {
        family inet {
            address 172.25.11.3/24;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.16.2.1/32;
        }
    }
}

[edit interfaces]
lab@mxC#
```

Question: Did the permanent interfaces fxp0 and lo0 inherit the family iso command?

_____

Answer: No, the <*-*> wildcard parameter only matches on transit interfaces.

Step 2.4

Issue a **`run show interface terse`** command to verify the family iso configuration on the interfaces.

```
[edit interfaces]
lab@mxC# run show interfaces terse
Interface               Admin Link Proto    Local                 Remote
ge-0/0/0                up    up
lc-0/0/0                up    up
lc-0/0/0.32769          up    up   vpls
pfe-0/0/0               up    up
pfe-0/0/0.16383         up    up   inet
                                   inet6
pfh-0/0/0               up    up
pfh-0/0/0.16383         up    up   inet
pfh-0/0/0.16384         up    up   inet
ge-0/0/1                up    up
ge-0/0/1.0              up    up   inet     172.22.122.1/24
                                   iso
                                   multiservice
ge-0/0/2                up    up
ge-0/0/2.0              up    up   inet     172.22.124.1/24
                                   iso
                                   multiservice
ge-0/0/3                up    up
ge-0/0/3.0              up    up   inet     10.0.14.1/24
                                   iso
                                   multiservice
ge-0/0/4                up    up
ge-0/0/5                up    up
ge-0/0/6                up    up
ge-0/0/7                up    down
ge-0/0/8                up    down
ge-0/0/9                up    down
cbp0                    up    up
demux0                  up    up
dsc                     up    up
em1                     up    up
em1.0                   up    up   inet     10.0.0.4/8
                                            128.0.0.1/2
                                            128.0.0.4/2
                                   inet6    fe80::250:56ff:fea9:1817/64
                                            fec0::a:0:0:4/64
                                   tnp      0x4
esi                     up    up
fxp0                    up    up
fxp0.0                  up    up   inet     172.25.11.3/24
gre                     up    up
ipip                    up    up
irb                     up    up
jsrv                    up    up
jsrv.1                  up    up   inet     128.0.0.127/2
lo0                     up    up
lo0.0                   up    up   inet     172.16.2.1          --> 0/0
lo0.16384               up    up   inet     127.0.0.1           --> 0/0
```

```
lo0.16385                  up    up   inet
lsi                        up    up
mtun                       up    up
pimd                       up    up
pime                       up    up
pip0                       up    up
pp0                        up    up
rbeb                       up    up
tap                        up    up
vtep                       up    up
```

Question: Only the transit interfaces with configured logical units inherited the `family iso` parameter. Why?

_____

_____

Answer: The group interface command included the unit wildcard `<*>`, which only matches on interfaces with logical units configured.

Step 2.5

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, enter configuration mode and navigate to the `[edit groups ISO]` hierarchy. Using wildcard commands, configure all transit interfaces within the logical systems (`logical system <*> interface <*-*>`) to support the IS-IS family iso on all logical units (`unit <*>`).

```
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# edit groups ISO

[edit groups ISO]
lab@mxD# set logical-systems <*> interfaces <*-*> unit <*> family iso

[edit groups ISO]
lab@mxD#
```

Step 2.6

Navigate to the `[edit logical-system R3-1 interfaces]` hierarchy and apply the ***ISO*** group to all interfaces.

```
[edit groups ISO]
lab@mxD# top edit logical-systems R3-1 interfaces

[edit logical-systems R3-1 interfaces]
lab@mxD# set apply-groups ISO

[edit logical-systems R3-1 interfaces]
lab@mxD#
```

Step 2.7

Navigate to the `[edit logical-system R3-2 interfaces]` hierarchy and apply the ***ISO*** group to all interfaces. Commit the configuration when completed.

```
[edit logical-systems R3-1 interfaces]
lab@mxD# top edit logical-systems R3-2 interfaces

[edit logical-systems R3-2 interfaces]
lab@mxD# set apply-groups ISO

[edit logical-systems R3-2 interfaces]
lab@mxD# commit
commit complete

[edit logical-systems R3-2 interfaces]
lab@mxD#
```

Step 2.8

Return to the SSH session that you have open to the mxE router.

From the SSH session with the mxE router, enter configuration mode and navigate to the `[edit groups ISO]` hierarchy. Using wildcard commands, configure all transit interfaces within the logical systems (`logical system <*> interface <*-*>`) to support the IS-IS family iso on all logical units (`unit <*>`).

```
lab@mxE> configure
Entering configuration mode

[edit]
lab@mxE# edit groups ISO

[edit groups ISO]
lab@mxE# set logical-systems <*> interfaces <*-*> unit <*> family iso

[edit groups ISO]
lab@mxE#
```

Step 2.9

Navigate to the `[edit logical-system P1 interfaces]` hierarchy and apply the ***ISO*** group to all interfaces.

```
[edit groups ISO]
lab@mxE# top edit logical-systems P1 interfaces

[edit logical-systems P1 interfaces]
lab@mxE# set apply-groups ISO

[edit logical-systems P1 interfaces]
lab@mxE#
```

Step 2.10

Navigate to the `[edit logical-system P2 interfaces]` hierarchy and apply the ***ISO*** group to all interfaces. Commit the configuration when completed.

```
[edit logical-systems P1 interfaces]
lab@mxE# top edit logical-systems P2 interfaces

[edit logical-systems P2 interfaces]
lab@mxE# set apply-groups ISO

[edit logical-systems P2 interfaces]
lab@mxE# commit
commit complete

[edit logical-systems P2 interfaces]
lab@mxE#
```

Step 2.11

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, enter configuration mode and navigate to the `[edit groups ISO]` hierarchy. Using wildcard commands, configure all transit interfaces (`interface <*-*>`) to support the IS-IS family iso on all logical units (`unit <*>`).

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit groups ISO

[edit groups ISO]
lab@mxA# set interfaces <*-*> unit <*> family iso

[edit groups ISO]
lab@mxA#
```

Step 2.12

Navigate to the `[edit interfaces]` hierarchy and apply the ***ISO*** group to all interfaces. Commit the configuration when completed.

```
[edit groups ISO]
lab@mxA# top edit interfaces

[edit interfaces]
lab@mxA# set apply-groups ISO

[edit interfaces]
lab@mxA# commit
commit complete

[edit interfaces]
lab@mxA#
```

Part 3: Configuring the IS-IS Network Entity Title

In this lab part, you configure IS-IS to define the NET address on an active IS-IS interface. This address is typically configured on the loopback interface because it is always an active interface on the router. Use the following chart to determine the correct NET addresses for your routers.

Router	Interface	Network Entity Title
mxA	lo0.0	49.0001.1720.1600.1001.00
R3-1	lo0.1	49.0001.1720.1600.1002.00
R3-2	lo0.2	49.0002.1720.1600.2002.00
mxC	lo0.0	49.0002.1720.1600.2001.00
P1	lo0.1	49.1234.1720.3110.0001.00
P2	lo0.2	49.1234.1720.3110.1001.00

Step 3.1

Navigate to the `[edit interfaces lo0]` hierarchy and configure loopback interface with the appropriate IS-IS NET address. Commit the configuration when completed.

```
[edit interfaces]
lab@mxA# edit lo0

[edit interfaces lo0]
lab@mxA# set unit 0 family iso address 49.0001.1720.1600.1001.00

[edit interfaces lo0]
lab@mxA# commit
commit complete

[edit interfaces lo0]
lab@mxA#
```

Step 3.2

Use the **`run show interfaces lo0 terse`** command to verify that the loopback interfaces have the correct IS-IS NET address configured.

```
[edit interfaces lo0]
lab@mxA# run show interfaces lo0 terse
Interface               Admin Link Proto    Local                 Remote
lo0                     up    up
lo0.0                   up    up   inet     172.16.1.1          --> 0/0
                                   iso      49.0001.1720.1600.1001
lo0.16384               up    up   inet     127.0.0.1           --> 0/0
lo0.16385               up    up   inet
```

Question: What is the IS-IS area configured on the loopback interfaces?

_____

Answer: Reading the NET address from right to left, the Net-Selector is 00, Sys-ID is 1720.1600.1001, and the area in the above capture is 49.0001

Step 3.3

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, navigate to the `[edit interfaces lo0]` hierarchy and configure loopback interface with the appropriate IS-IS NET address. Commit the configuration when completed.

```
[edit interfaces]
lab@mxC# edit lo0

[edit interfaces lo0]
lab@mxC# set unit 0 family iso address 49.0002.1720.1600.2001.00

[edit interfaces lo0]
lab@mxC# commit
commit complete

[edit interfaces lo0]
lab@mxC#
```

Step 3.4

Use the **`run show interfaces lo0 terse`** command to verify that the loopback interfaces have the correct IS-IS NET address configured.

```
[edit interfaces lo0]
lab@mxC# run show interfaces lo0 terse
Interface               Admin Link Proto    Local                 Remote
lo0                     up    up
lo0.0                   up    up   inet     172.16.2.1          --> 0/0
                                   iso      49.0002.1720.1600.2001
lo0.16384               up    up   inet     127.0.0.1           --> 0/0
lo0.16385               up    up   inet
```

Question: Does the output show that the interface has been assigned the appropriate NET address?

_____

Answer: Yes, the loopback interface has been assigned the correct NET address.

Step 3.5

Return to the SSH session that you have open to the mxE router.

From the SSH session with the mxE router, navigate to the `[edit logical-systems P2 interfaces lo0]` hierarchy and configure loopback interface with the appropriate IS-IS NET address.

```
[edit logical-systems P2 interfaces]
lab@mxE# edit lo0

[edit logical-systems P2 interfaces lo0]
lab@mxE# set unit 2 family iso address 49.1234.1720.3110.1001.00

[edit logical-systems P2 interfaces lo0]
```

```
lab@mxE#
```

Step 3.6

Navigate to the `[edit logical-systems P1 interfaces lo0]` hierarchy and configure loopback interface with the appropriate IS-IS NET address. Commit the configuration when completed.

```
[edit logical-systems P2 interfaces lo0]
lab@mxE# top edit logical-systems P1 interfaces lo0

[edit logical-systems P1 interfaces lo0]
lab@mxE# set unit 1 family iso address 49.1234.1720.3110.0001.00

[edit logical-systems P1 interfaces lo0]
lab@mxE# commit and-quit
commit complete
Exiting configuration mode

lab@mxE>
```

Step 3.7

Use the **`run show interfaces lo0 terse`** command to verify that the loopback interfaces have the correct IS-IS NET address configured.

```
lab@mxE> show interfaces lo0 terse
Interface               Admin Link Proto    Local                 Remote
lo0                     up    up
lo0.1                   up    up   inet     172.31.100.1        --> 0/0
                                   iso      49.1234.1720.3110.0001
lo0.2                   up    up   inet     172.31.101.1        --> 0/0
                                   iso      49.1234.1720.3110.1001
lo0.16384               up    up   inet     127.0.0.1           --> 0/0
lo0.16385               up    up   inet
```

Question: Does the output show that the interfaces have been assigned the appropriate NET addresses?

_____

Answer: Yes, the loopback interfaces have been assigned the correct NET addresses.

Step 3.8

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, navigate to the `[edit logical-systems R3-2 interfaces lo0]` hierarchy and configure loopback interface with the appropriate IS-IS NET address.

```
[edit logical-systems R3-2 interfaces]
lab@mxD# edit lo0

[edit logical-systems R3-2 interfaces lo0]
lab@mxD# set unit 2 family iso address 49.0002.1720.1600.2002.00
```

```
[edit logical-systems R3-2 interfaces lo0]
lab@mxD#
```

Step 3.9

Navigate to the `[edit logical-systems R3-1 interfaces lo0]` hierarchy and configure loopback interface with the appropriate IS-IS NET address. Commit the configuration when completed.

```
[edit logical-systems R3-2 interfaces lo0]
lab@mxD# top edit logical-systems R3-1 interfaces lo0

[edit logical-systems R3-1 interfaces lo0]
lab@mxD# set unit 1 family iso address 49.0001.1720.1600.1002.00

[edit logical-systems R3-1 interfaces lo0]
lab@mxD# commit and-quit
commit complete
Exiting configuration mode

lab@mxD>
```

Step 3.10

Use the **`run show interfaces lo0 terse`** command to verify that the loopback interfaces have the correct IS-IS NET address configured.

```
lab@mxD> show interfaces lo0 terse
Interface               Admin Link Proto    Local                 Remote
lo0                     up    up
lo0.1                   up    up   inet     172.16.1.2          --> 0/0
                                   iso      49.0001.1720.1600.1002
lo0.2                   up    up   inet     172.16.2.2          --> 0/0
                                   iso      49.0002.1720.1600.2002
lo0.16384               up    up   inet     127.0.0.1           --> 0/0
lo0.16385               up    up   inet
```

Question: Does the output show that the interfaces have been assigned the appropriate NET addresses?

_____

Answer: Yes, the loopback interfaces have been assigned the correct NET addresses.

Part 4: Configuring Interfaces as Part of the IS-IS Protocol

In this lab part, you configure IS-IS on the participating interfaces under `[edit protocols isis]`. The IS-IS default operation for all configured interfaces is to form an adjacency in both Level 1 and Level 2 when possible.

Step 4.1

Change the perspective of the mxD router''s CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxD> set cli logical-system R3-1
```

```
Logical system: R3-1

lab@mxD:R3-1>
```

Step 4.2

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Configure each of the R3-1 router's interfaces to run the ISIS protocol. Commit the configuration when completed.

```
lab@mxD:R3-1> configure
Entering configuration mode

[edit]
lab@mxD:R3-1# edit protocols isis

[edit protocols isis]
lab@mxD:R3-1# set interface lo0.1

[edit protocols isis]
lab@mxD:R3-1# set interface ge-0/0/1.0

[edit protocols isis]
lab@mxD:R3-1# commit
commit complete

[edit protocols isis]
lab@mxD:R3-1#
```

Step 4.3

Use the **`run show isis interface`** command to verify IS-IS is active on the interfaces.

```
[edit protocols isis]
lab@mxD:R3-1# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/1.0            3   0x1 mxD-R3-1.00       mxD-R3-1.00            10/10
lo0.1                 3   0x1 Passive           Passive                 0/0
```

Question: Which IS-IS levels are configured on the two interfaces?

_____

_____

Answer: Both Level 1 and Level 2 are configured on these interfaces, which is shown by a 3 (1+2) in the "L" column.

Step 4.4

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, navigate to the `[edit protocols isis]` hierarchy and configure each of the mxA router's interfaces to run the ISIS protocol. Commit the configuration when completed.

```
[edit interfaces lo0]
lab@mxA# top edit protocols isis

[edit protocols isis]
lab@mxA# set interface lo0.0

[edit protocols isis]
lab@mxA# set interface ge-0/0/1.0

[edit protocols isis]
lab@mxA# set interface ge-0/0/2.0

[edit protocols isis]
lab@mxA# set interface ge-0/0/3.0

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

[edit protocols isis]
lab@mxA#
```

Step 4.5

Use the **`run show isis interface`** command to verify IS-IS is active on the interfaces. It may take a minute for the ge-0/0/3.0 interface to show up as mxA.02 rather than mxA.00

```
[edit protocols isis]
lab@mxA# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/1.0            3   0x1 mxA.00            mxA.00                 10/10
ge-0/0/2.0            3   0x1 mxA.00            mxA.00                 10/10
ge-0/0/3.0            3   0x2 mxA.02            mxA.02                 10/10
lo0.0                 3   0x1 Passive           Passive                 0/0
```

Question: Which IS-IS levels are configured on the four interfaces?

_____

_____

Answer: Both Level 1 and Level 2 are configured on these interfaces, which is shown by a 3 (1+2) in the “L” column.

Step 4.6

Use the **`run show isis adjacency`** command to verify IS-IS adjacencies in the default routing instance.

```
[edit protocols isis]
lab@mxA# run show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/3.0            mxD-R3-1       1  Up                   20  0:50:56:a9:3d:65
ge-0/0/3.0            mxD-R3-1       2  Up                   21  0:50:56:a9:3d:65
```

Question: How many adjacencies have been formed? Why?

_____

_____

_____

Answer: Two adjacencies have formed because the transit interfaces are configured for both level 1 and level 2.

Step 4.7

Disable IS-IS Level 1 on the ge-0/0/3.0 interface of the mxA router. Commit the configuration and exit to operational mode when completed.

```
[edit protocols isis]
lab@mxA# set interface ge-0/0/3.0 level 1 disable

[edit protocols isis]
lab@mxA# commit and-quit
commit complete
Exiting configuration mode

lab@mxA>
```

Step 4.8

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, Disable IS-IS Level 1 on the ge-0/0/1.0 interface of the R3-1 router. Commit the configuration and exit to operational mode when completed.

```
[edit protocols isis]
lab@mxD:R3-1# set interface ge-0/0/1.0 level 1 disable

[edit protocols isis]
lab@mxD:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R3-1>
```

Step 4.9

Use the **show isis interface** command to verify IS-IS is active on the interfaces.

```
lab@mxD:R3-1> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/1.0            2   0x1 Disabled          mxA.02                 10/10
lo0.1                 3   0x1 Passive           Passive                 0/0
```

Question: Which IS-IS level is configured on the ge-0/0/1.0 interface?

_____

_____

Answer: Only Level 2 is configured on the interfaces.

Step 4.10

Issue the **show isis adjacency** commands to verify that only one adjacency has been formed between the R3-1 and mxA routers.

```
lab@mxD:R3-1> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/1.0            mxA            2  Up                    7  0:50:56:a9:61:f9
```

Question: How many ISIS adjacencies have been formed between the R3-1 and mxA routers?

_____

_____

Answer: A single Level 2 adjacency has formed between the R3-1 and mxA router. If the previous Level 1 adjacency is shown in the down state, it can be ignored. It can take the system a few moments for the changes to take affect.

Step 4.11

Return to the SSH session that you have open to the mxE router.

From the SSH session with the mxE router, change the perspective of the CLI to that of the P1 logical system by issuing the **set cli logical-system *P1*** command.

```
lab@mxE> set cli logical-system P1
Logical system: P1

lab@mxE:P1>
```

Step 4.12

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Configure the P1 router to run ISIS at level 2 only.

```
lab@mxE:P1> configure
Entering configuration mode

[edit]
lab@mxE:P1# edit protocols isis

[edit protocols isis]
lab@mxE:P1# set level 1 disable

[edit protocols isis]
lab@mxE:P1#
```

Step 4.13

Configure each of the P1 router's interfaces to run the ISIS protocol. Commit the configuration and exit to operational mode when completed.

```
[edit protocols isis]
lab@mxE:P1# set interface lo0.1

[edit protocols isis]
lab@mxE:P1# set interface ge-0/0/7.0

[edit protocols isis]
lab@mxE:P1# set interface ge-0/0/2.0

[edit protocols isis]
lab@mxE:P1# set interface lt-0/0/10.0

[edit protocols isis]
lab@mxE:P1# commit and-quit
commit complete
Exiting configuration mode

lab@mxE:P1>
```

Step 4.14

Use the **show isis interface** command to verify IS-IS is active on the interfaces.

```
lab@mxE:P1> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/2.0            2   0x1 Disabled          mxE-P1.00              10/10
ge-0/0/7.0            2   0x2 Disabled          mxE-P1.02              10/10
lo0.1                 2   0x1 Passive           Passive                 0/0
lt-0/0/10.0           2   0x1 Disabled          mxE-P1.00              10/10
```

Question: Which IS-IS levels are configured on the transit interfaces?

_____

_____

Answer: Only level 2 has been enabled on the transit interfaces.

Step 4.15

Issue the **show isis adjacency** commands to verify that an adjacency has been formed between the P1 and mxA routers.

```
lab@mxE:P1> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/7.0            mxA            2  Up                    7  0:50:56:a9:1b:d0
```

Question: Has an ISIS adjacency formed between the P1 and mxA router?

_____

_____

Answer: Yes. A level 2 adjacency has formed between the two routers.

Step 4.16

Change the perspective of the CLI to that of the P2 logical system by issuing the **`set cli logical-system P2`** command.

```
lab@mxE:P1> set cli logical-system P2
Logical system: P2

lab@mxE:P2>
```

Step 4.17

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Configure the P2 router to run ISIS at level 2 only.

```
lab@mxE:P2> configure
Entering configuration mode

[edit]
lab@mxE:P2# edit protocols isis

[edit protocols isis]
lab@mxE:P2# set level 1 disable

[edit protocols isis]
lab@mxE:P2#
```

Step 4.18

Configure each of the P2 router's interfaces to run the ISIS protocol. Commit the configuration and exit to operational mode when completed.

```
[edit protocols isis]
lab@mxE:P2# set interface lo0.2

[edit protocols isis]
lab@mxE:P2# set interface ge-0/0/6.0

[edit protocols isis]
lab@mxE:P2# set interface ge-0/0/3.0

[edit protocols isis]
lab@mxE:P2# set interface lt-0/0/10.1

[edit protocols isis]
lab@mxE:P2# commit and-quit
commit complete
Exiting configuration mode
```

```
lab@mxE:P2>
```

Step 4.19

Use the **`show isis interface`** command to verify IS-IS is active on the interfaces.

```
lab@mxE:P2> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/3.0            2   0x1 Disabled          mxE-P2.00              10/10
ge-0/0/6.0            2   0x2 Disabled          mxE-P2.02              10/10
lo0.2                 2   0x1 Passive           Passive                 0/0
lt-0/0/10.1           2   0x1 Disabled          mxE-P2.00              10/10
```

Question: Which IS-IS levels are configured on the transit interfaces?

_____

_____

Answer: Only level 2 has been enabled on the transit interfaces.

Step 4.20

Issue the **`show isis adjacency`** commands to verify that an adjacency has been formed between the P2, P1, and mxA routers.

```
lab@mxE:P2> show isis adjacency
Interface             System          L State        Hold (secs) SNPA
ge-0/0/6.0            mxA             2  Up                   23  0:50:56:a9:12:91
lt-0/0/10.1           mxE-P1          2  Up                   22  0:50:56:a9:0:92
```

Question: Has an ISIS adjacency formed between the P2 and P1 routers as well as the P2 and the mxA routers?

_____

_____

Answer: Yes. A level 2 adjacency has formed between the P2 router and the other routers.

Step 4.21

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, navigate to the `[edit protocols isis]` hierarchy. Configure the mxC router to run ISIS at level 2 only.

```
[edit interfaces lo0]
lab@mxC# top edit protocols isis

[edit protocols isis]
lab@mxC# set level 1 disable

[edit protocols isis]
lab@mxC#
```

Step 4.22

Configure each of the mxC router's interfaces to run the ISIS protocol. Commit the configuration when completed.

```
[edit protocols isis]
lab@mxC# set interface lo0.0

[edit protocols isis]
lab@mxC# set interface ge-0/0/1.0

[edit protocols isis]
lab@mxC# set interface ge-0/0/2.0

[edit protocols isis]
lab@mxC# set interface ge-0/0/3.0

[edit protocols isis]
lab@mxC# commit and-quit
commit complete
Exiting configuration mode

lab@mxC>
```

Step 4.23

Use the **show isis interface** command to verify IS-IS is active on the interfaces.

```
lab@mxC> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/1.0            2   0x1 Disabled          mxE-P1.03              10/10
ge-0/0/2.0            2   0x2 Disabled          mxC.02                 10/10
ge-0/0/3.0            2   0x1 Disabled          mxC.00                 10/10
lo0.0                 2   0x1 Passive           Passive                 0/0
```

Question: Which IS-IS levels are configured on the transit interfaces?

_____

_____

Answer: Only level 2 has been enabled on the transit interfaces. The DR selected may be different for you.

Step 4.24

Issue the **show isis adjacency** commands to verify that an adjacency has been formed between the P2, P1, and mxC routers.

```
lab@mxC> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/1.0            mxE-P1         2  Up                    8  0:50:56:a9:78:33
ge-0/0/2.0            mxE-P2         2  Up                    8  0:50:56:a9:7c:c1
```

Question: Has an ISIS adjacency formed between the mxC and P1 routers as well as the mxC and P2 routers?

_____

_____

Answer: Yes. A level 2 adjacency has formed between the mxC router and the other routers.

Step 4.25

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, change the perspective of the CLI to that of the R3-2 logical system by issuing the **`set cli logical-system R3-2`** command.

```
lab@mxD:R3-1> set cli logical-system R3-2
Logical system: R3-2

lab@mxD:R3-2>
```

Step 4.26

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Configure the R3-2 router to run ISIS at level 2 only.

```
lab@mxD:R3-2> configure
Entering configuration mode

[edit]
lab@mxD:R3-2# edit protocols isis

[edit protocols isis]
lab@mxD:R3-2# set level 1 disable

[edit protocols isis]
lab@mxD:R3-2#
```

Step 4.27

Configure each of the R3-2 router's interfaces to run the ISIS protocol. Commit the configuration and exit to operational mode when completed.

```
[edit protocols isis]
lab@mxD:R3-2# set interface lo0.2

[edit protocols isis]
lab@mxD:R3-2# set interface ge-0/0/4.0

[edit protocols isis]
lab@mxD:R3-2# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R3-2>
```

Step 4.28

Use the **show isis interface** command to verify IS-IS is active on the interfaces.

```
lab@mxD:R3-2> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/4.0            2   0x1 Disabled          mxC.02                 10/10
lo0.2                 2   0x1 Passive           Passive                 0/0
```

Question: Which IS-IS levels are configured on the transit interface?

_____

_____

Answer: Only level 2 has been enabled on the transit interface.

Step 4.29

Issue the **show isis adjacency** commands to verify that an adjacency has been formed between the R3-2 and mxC routers.

```
lab@mxD:R3-2> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/4.0            mxC            2  Up                    6  0:50:56:a9:68:d8
```

Question: Has an ISIS adjacency formed between the R3-2 and mxC routers?

_____

_____

Answer: Yes. A level 2 adjacency has formed between the two routers.

Part 5: Examining the IS-IS Database

In this lab part, you examine both the IS-IS link state database (LSDB) and the shortest-path-first (SPF) tree database. The IS-IS LSDB data is input into the SPF algorithm and the algorithm creates two more databases, a candidate database and tree database. The candidate database is temporary and is deleted after the SPF calculation is complete. The SPF tree database is used to populate the local routing table.

Step 5.1

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, use the **show isis database** command to display the IS-IS LSDB.

```
lab@mxA> show isis database
IS-IS level 1 link-state database:
LSP ID                    Sequence Checksum Lifetime Attributes
mxA.00-00                      0x4   0x3d8a      543 L1 L2 Attached
```

```
mxD-R3-1.00-00                    0x3   0xe33f       77 L1 L2
  3 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
mxA.00-00                        0x7   0x4b99      658 L1 L2
mxA.02-00                        0x2   0x12c4      914 L1 L2
mxD-R3-1.00-00                   0x7   0x8c92     1189 L1 L2
mxC.00-00                        0x4   0xcdf4     1165 L1 L2
mxC.02-00                        0x1    0xb48      772 L1 L2
mxC.03-00                        0x1   0xa3e2     1165 L1 L2
mxD-R3-2.00-00                   0x2   0x7e5d     1163 L1 L2
mxE-P1.00-00                     0x4   0xe46c      744 L1 L2
mxE-P1.02-00                     0x1   0xe5a2      541 L1 L2
mxE-P1.03-00                     0x1   0x2344      744 L1 L2
mxE-P2.00-00                     0x4   0x7790      774 L1 L2
mxE-P2.02-00                     0x1   0x3820      656 L1 L2
mxE-P2.03-00                     0x1   0xd749      683 L1 L2
  13 LSPs
```

Question: The mxA.00-00 LSP in the Level 1 database has the attach bit set. Why?

_____

_____

Answer: Any IS-IS router connected to both Level 1 and Level 2 will enable the attach bit in the Level 1 LSP. All routers in the Level 1 database will create a default route to the closest L1/L2 attached router.

Question: Why are there more than one mxA LSPs?

_____

_____

Answer: The LSP represented by mxA.00-00 is the default LSP advertising the locally attached topology, while the mxA.02-00 is the Designated Intermediate System (DIS) pseudo LSP. Your output might be different based on DIS node elected.

Step 5.2

Use the **`show isis database level 2 mxA.00-00 extensive`** command to display the IS-IS header and the type/length/value (TLV) entries for the mxA router's LSP.

```
lab@mxA> show isis database level 2 mxA.00-00 extensive
IS-IS level 2 link-state database:

mxA.00-00 Sequence: 0xa, Checksum: 0x7990, Lifetime: 985 secs
   IS neighbor: mxA.02                          Metric:       10
     Two-way fragment: mxA.02-00, Two-way first fragment: mxA.02-00
```

```
 IS neighbor: mxA.03                          Metric:       10
   Two-way fragment: mxA.03-00, Two-way first fragment: mxA.03-00
 IS neighbor: mxE-P2.02                       Metric:       10
   Two-way fragment: mxE-P2.02-00, Two-way first fragment: mxE-P2.02-00
 IP prefix: 10.0.10.0/24                      Metric:       10 Internal Up
 IP prefix: 172.16.1.1/32                     Metric:        0 Internal Up
 IP prefix: 172.22.121.0/24                   Metric:       10 Internal Up
 IP prefix: 172.22.123.0/24                   Metric:       10 Internal Up

Header: LSP ID: mxA.00-00, Length: 269 bytes
  Allocated length: 1492 bytes, Router ID: 172.16.1.1
  Remaining lifetime: 985 secs, Level: 2, Interface: 0
  Estimated free bytes: 1171, Actual free bytes: 1223
  Aging timer expires in: 985 secs
  Protocols: IP, IPv6

Packet: LSP ID: mxA.00-00, Length: 269 bytes, Lifetime : 1198 secs
  Checksum: 0x7990, Sequence: 0xa, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

TLVs:
  Area address: 49.0001 (3)
  LSP Buffer Size: 1492
  Speaks: IP
  Speaks: IPV6
  IP router id: 172.16.1.1
  IP address: 172.16.1.1
  Hostname: mxA
  IS neighbor: mxA.02, Internal, Metric: default 10
  IS neighbor: mxA.03, Internal, Metric: default 10
  IS neighbor: mxE-P2.02, Internal, Metric: default 10
  IS extended neighbor: mxA.02, Metric: default 10
    IP address: 10.0.10.1
    Local interface index: 334, Remote interface index: 0
  IS extended neighbor: mxA.03, Metric: default 10
    IP address: 172.22.121.1
    Local interface index: 332, Remote interface index: 0
  IS extended neighbor: mxE-P2.02, Metric: default 10
    IP address: 172.22.123.1
    Local interface index: 333, Remote interface index: 0
  IP prefix: 172.22.121.0/24, Internal, Metric: default 10, Up
  IP prefix: 172.22.123.0/24, Internal, Metric: default 10, Up
  IP prefix: 10.0.10.0/24, Internal, Metric: default 10, Up
  IP prefix: 172.16.1.1/32, Internal, Metric: default 0, Up
  IP extended prefix: 172.22.121.0/24 metric 10 up
  IP extended prefix: 172.22.123.0/24 metric 10 up
  IP extended prefix: 10.0.10.0/24 metric 10 up
  IP extended prefix: 172.16.1.1/32 metric 0 up
  Router Capability:  Router ID 172.16.1.1, Flags: 0x00
No queued transmissions
```

Question: Which protocols are supported in this LSP?

_____

_____

Answer: Both IP and IPV6 are supported in this LSP.

Question: Under the TLV section, the mxA router's loopback address is being advertised in which two TLVs?

_____

_____

Answer: The IP prefix, TLV 128, and the IP extended prefix, TLV 135.

Step 5.3

Use the **show isis route** command to display the SPF tree database.

```
lab@mxA> show isis route
 IS-IS routing table              Current version: L1: 9 L2: 16
IPv4/IPv6 Routes
----------------
Prefix              L Version   Metric Type Interface       NH   Via
   Backup Score
10.0.14.0/24        2      16       30 int  ge-0/0/1.0      IPV4 mxE-P1
                                           ge-0/0/2.0      IPV4 mxE-P2
172.16.1.2/32       2      16       10 int  ge-0/0/3.0      IPV4 mxD-R3-1
172.16.2.1/32       2      16       20 int  ge-0/0/1.0      IPV4 mxE-P1
                                           ge-0/0/2.0      IPV4 mxE-P2
172.16.2.2/32       2      16       30 int  ge-0/0/1.0      IPV4 mxE-P1
                                           ge-0/0/2.0      IPV4 mxE-P2
172.22.122.0/24     2      16       20 int  ge-0/0/1.0      IPV4 mxE-P1
172.22.124.0/24     2      16       20 int  ge-0/0/2.0      IPV4 mxE-P2
172.22.252.0/30     2      16       20 int  ge-0/0/2.0      IPV4 mxE-P2
                                           ge-0/0/1.0      IPV4 mxE-P1
172.31.100.1/32     2      16       10 int  ge-0/0/1.0      IPV4 mxE-P1
172.31.101.1/32     2      16       10 int  ge-0/0/2.0      IPV4 mxE-P2
```

Question: The IS-IS tree table shows the output of the SPF algorithm. Examine the output for a prefix that matches R3-1's loopback address. From which database was this route learned?

_____

_____

Answer: The 172.16.1.2/32 entry was learned from the Level 2 database.

Step 5.4

Log out of your assigned devices using the **exit** command.

```
lab@mxA> exit

mxA (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

Network Management Diagram

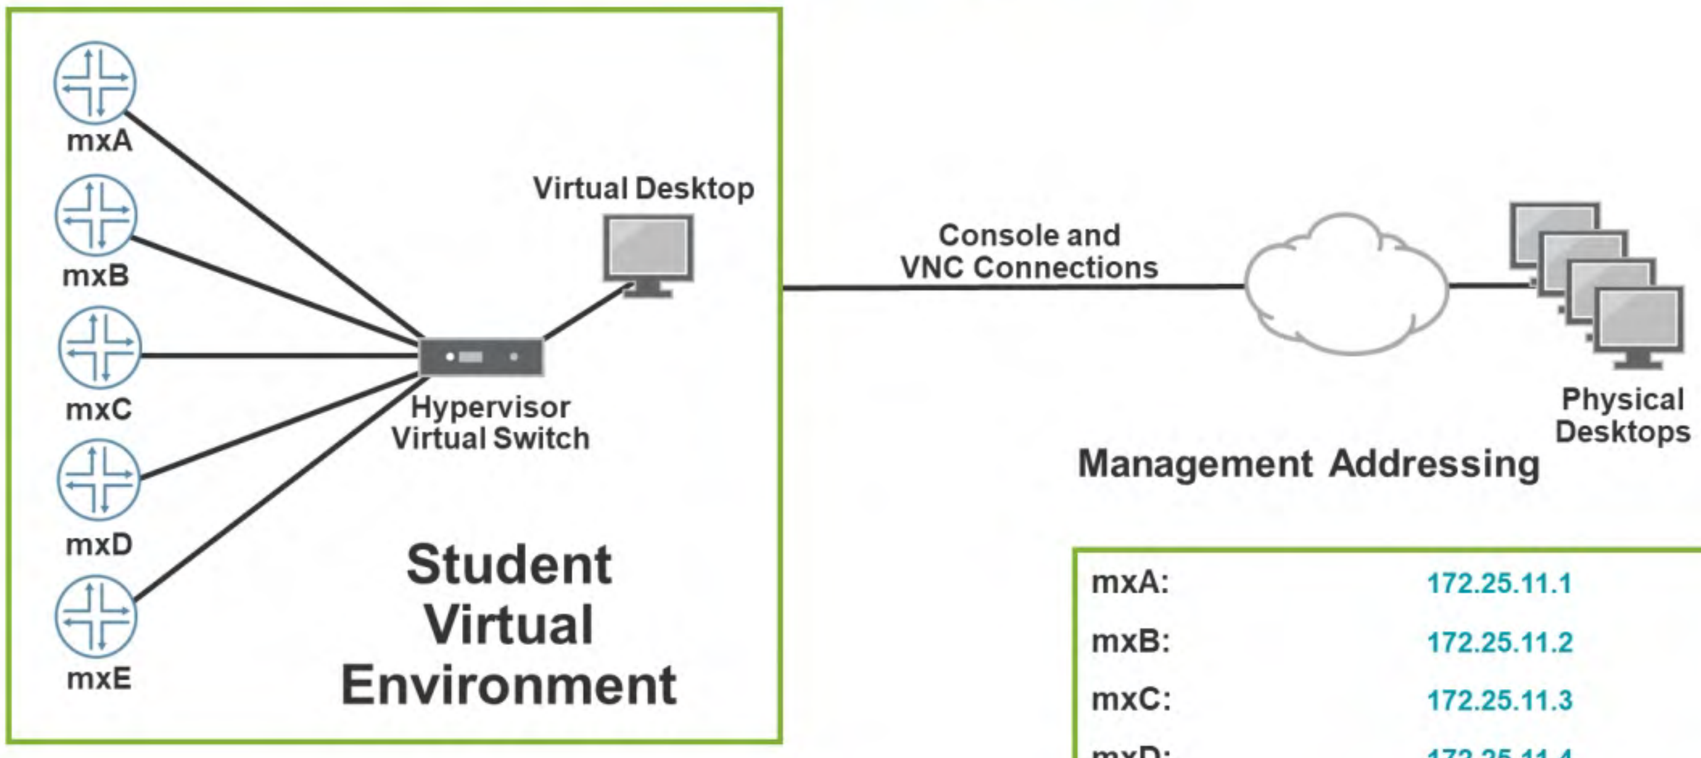

Management Addressing

mxA:	172.25.11.1
mxB:	172.25.11.2
mxC:	172.25.11.3
mxD:	172.25.11.4
mxE:	172.25.11.5
Virtual Desktop:	172.25.11.254

Note: Your instructor will provide address and access information.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab Network Diagram: IS-IS Configuration and Monitoring

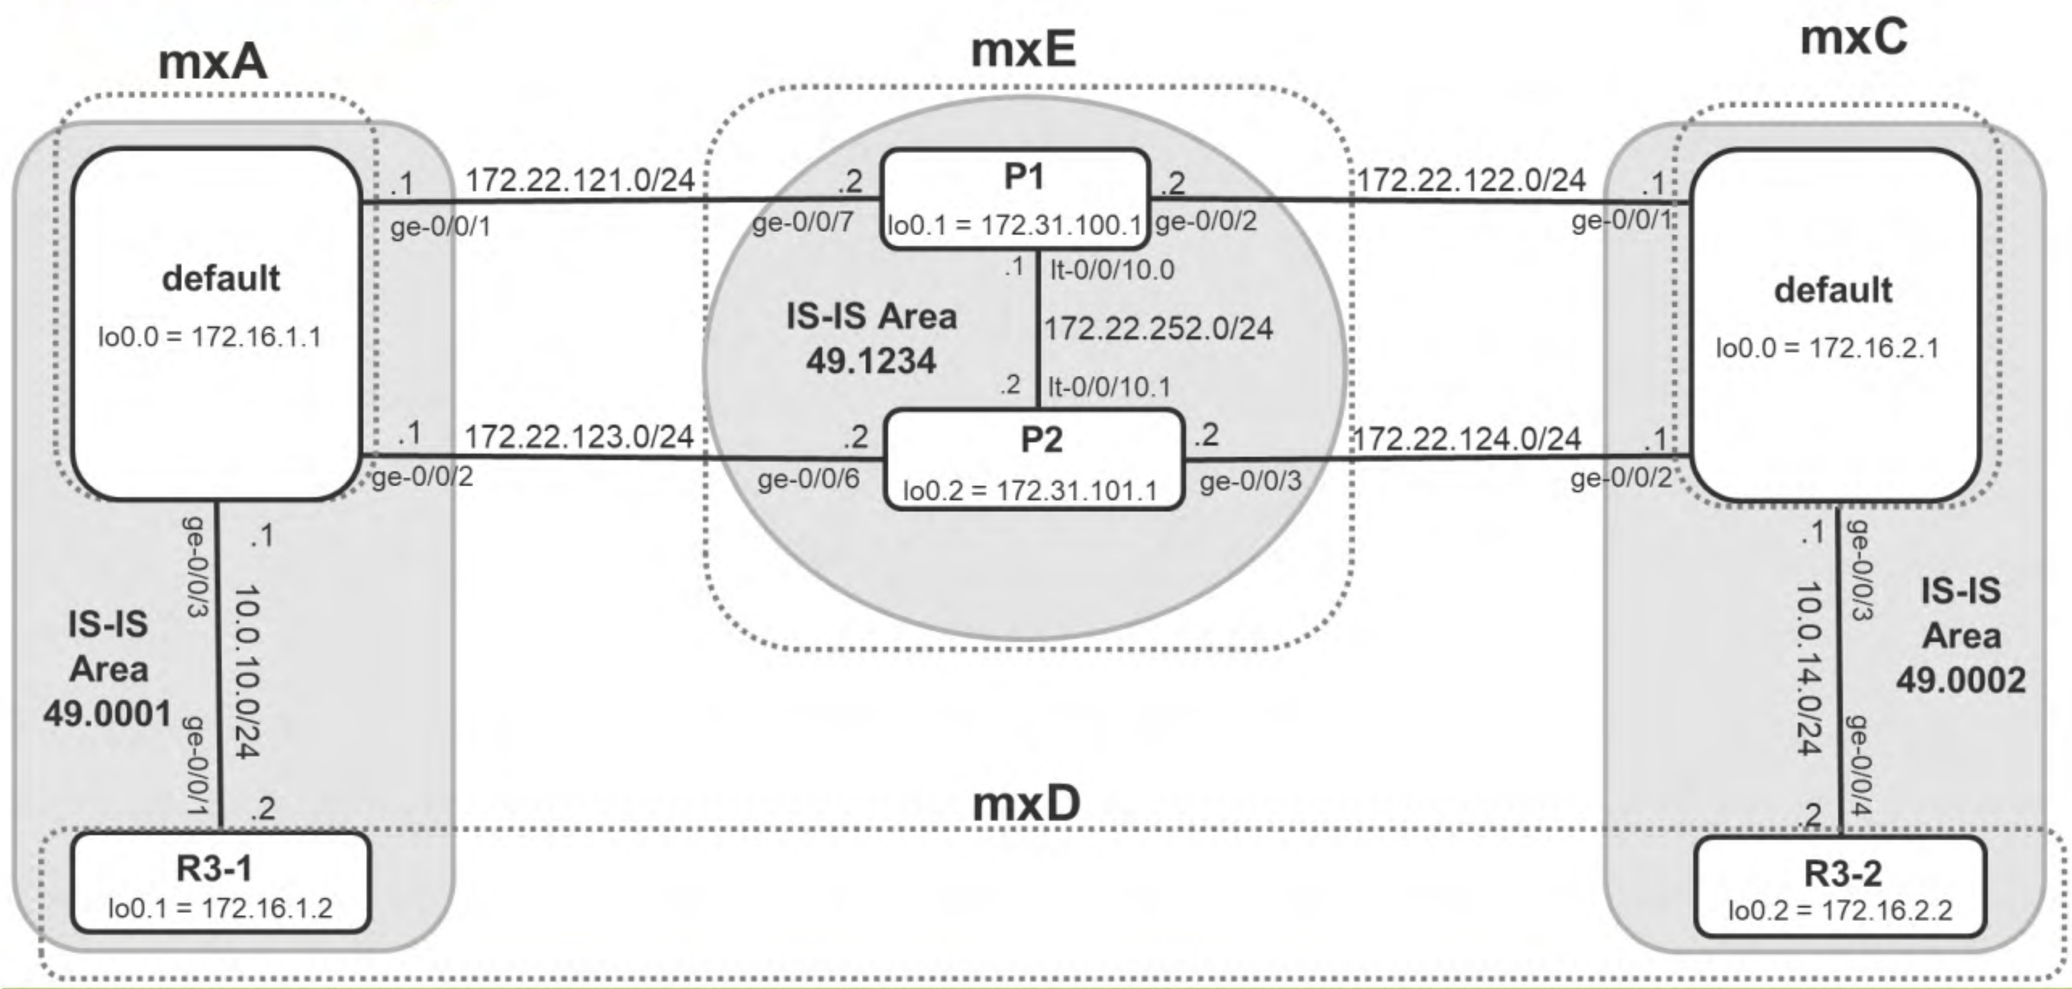

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab

Advanced IS-IS Configuration and Routing Policy

Overview

In this lab, you use the lab diagram titled "Lab Network Diagram: Advanced IS-IS Configuration and Routing Policy Lab" to establish an extended IS-IS Level 2 routing domain. This extended network allows the implementation and monitoring of some of the IS-IS advanced configuration options. In addition, this lab provides further insight into the IS-IS database with the redistribution of external routes.

By completing this lab, you perform the following tasks:

- Manipulate routing using various metrics.
- Authenticate IS-IS hello packets.
- Explore the use of the overload bit.
- Redistribute routes between IS-IS and RIP.

Part 1: Load Starting Configuration

In this lab part, you verify the initial configuration of the routers. You then verify that the interfaces are operational and that ISIS neighbor relationships have formed.

Note

The instructor will tell you the nature of your access and will provide you with the necessary details to access your assigned device.

Note

The lab topology requires you to display information in the different logical systems of mxA, mxB, mxC, mxD, and mxE. You will be changing the perspective of the CLI by issuing the **`set cli logical-system`** ***`ls-name`*** command. Refer to the lab diagram for the correct logical system name.

By changing the perspective of the CLI, you will be able to perform network commands such as **`ping`** or **`traceroute`** from the perspective of the pertinent logical system.

Step 1.1

Ensure that you know to which student device you have been assigned. Check with your instructor if you are not certain. Consult the management network diagram to determine the management address of the student devices.

Step 1.2

Access the CLI on your mxB device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab6-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.2
Password: lab123
Last login: Fri Dec  6 12:53:16 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxB> configure
Entering configuration mode

[edit]
lab@mxB# load override ajspr/lab6-start.config
load complete

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

lab@mxB>
```

Step 1.3

Issue the **show configuration** command. Use the lab diagram to verify that the mxB router has the correct logical system and interface configuration. Verify that RIP has been enabled.

```
lab@mxB> show configuration
## Last commit: 2019-12-06 12:57:34 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxB;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    P3 {
        interfaces {
            ge-0/0/6 {
                unit 0 {
                    family inet {
                        address 172.22.125.2/24;
                    }
                }
            }
            lo0 {
                unit 0 {
```

```
                family inet {
                    address 172.31.102.1/32;
                    address 20.20.1.1/32;
                    address 40.40.1.1/32;
                }
            }
        }
    }
    protocols {
        rip {
            group rip {
                export export-rip-policy;
                neighbor ge-0/0/6.0;
            }
        }
    }
    policy-options {
        policy-statement export-rip-policy {
            term 1 {
                from {
                    protocol [ static rip aggregate ];
                    route-filter 20.20.0.0/21 orlonger;
                }
                then accept;
            }
            term 2 {
                from {
                    protocol rip;
                    route-filter 0.0.0.0/0 exact;
                }
                then accept;
            }
        }
    }
    routing-options {
        static {
            route 20.20.0.0/24 reject;
            route 20.20.1.0/24 reject;
            route 20.20.2.0/24 reject;
            route 20.20.3.0/24 reject;
            route 20.20.4.0/25 reject;
            route 20.20.4.128/25 reject;
            route 20.20.5.0/26 reject;
            route 20.20.5.64/26 reject;
            route 20.20.5.128/26 reject;
            route 20.20.5.192/26 reject;
            route 40.40.0.0/24 reject;
            route 40.40.1.0/24 reject;
            route 40.40.2.0/24 reject;
            route 40.40.3.0/24 reject;
            route 172.16.1.2/32 next-hop 172.22.125.1;
            route 172.16.2.2/32 next-hop 172.22.126.1;
            route 172.16.1.4/32 next-hop 172.22.125.1;
            route 172.16.2.4/32 next-hop 172.22.126.1;
            route 10.0.20.0/22 next-hop 172.22.125.1;
```

```
                route 10.0.24.0/22 next-hop 172.22.126.1;
            }
            aggregate {
                route 20.20.0.0/21;
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.2/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxB router. If the configuration is incorrect, please notify your instructor.

Step 1.4

Access the CLI on your mxE device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab6-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.5
Password: lab123
Last login: Fri Dec  6 12:53:16 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil.
lab@mxE> configure
Entering configuration mode

[edit]
lab@mxE# load override ajspr/lab6-start.config
load complete

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

lab@mxE>
```

Step 1.5

Issue the **show configuration** command. Use the lab diagram to verify that the mxE router has the correct logical system and interface configuration. Verify that ISIS has been enabled (all interfaces at level 2).

```
lab@mxE> show configuration
## Last commit: 2019-12-26 12:07:00 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxE;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    P1 {
        interfaces {
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 172.22.122.2/24;
                    }
                    family iso;
                }
            }
```

```
        ge-0/0/5 {
            unit 0 {
                family inet {
                    address 172.22.121.2/24;
                }
                family iso;
            }
        }
        lo0 {
            unit 0 {
                family inet {
                    address 172.31.100.1/32;
                }
                family iso {
                    address 49.1234.1720.3110.0001.00;
                }
            }
        }
    }
    protocols {
        isis {
            interface ge-0/0/4.0 {
                level 1 disable;
            }
            interface ge-0/0/5.0 {
                level 1 disable;
            }
            interface lo0.0;
        }
    }
    routing-options {
        autonomous-system 65412;
    }
}
R5-1 {
    interfaces {
        ge-0/0/0 {
            unit 0 {
                family inet {
                    address 172.22.125.1/24;
                }
            }
        }
        ge-0/0/6 {
            unit 0 {
                family inet {
                    address 10.0.12.2/24;
                }
                family iso;
            }
        }
        ge-0/0/7 {
            unit 0 {
                family inet {
                    address 10.0.13.2/24;
```

```
                    }
                    family iso;
                }
            }
            lo0 {
                unit 1 {
                    family inet {
                        address 172.16.1.4/32;
                    }
                    family iso {
                        address 49.0001.1720.1600.1004.00;
                    }
                }
            }
        }
        protocols {
            isis {
                interface ge-0/0/6.0 {
                    level 1 disable;
                }
                interface ge-0/0/7.0 {
                    level 1 disable;
                }
                interface lo0.1;
            }
        }
    }
    R5-2 {
        interfaces {
            ge-0/0/2 {
                unit 0 {
                    family inet {
                        address 10.0.16.2/24;
                    }
                    family iso;
                }
            }
            ge-0/0/3 {
                unit 0 {
                    family inet {
                        address 10.0.17.2/24;
                    }
                    family iso;
                }
            }
            lo0 {
                unit 2 {
                    family inet {
                        address 172.16.2.4/32;
                    }
                    family iso {
                        address 49.0002.1720.1600.2004.00;
                    }
                }
            }
```

```
        }
        protocols {
            isis {
                interface ge-0/0/2.0 {
                    level 1 disable;
                }
                interface ge-0/0/3.0 {
                    level 1 disable;
                }
                interface lo0.2;
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.5/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxE router. If the configuration is incorrect, please notify your instructor.

Question: Based on the configuration, which routers (logical systems) will mxE represent in this lab?

_____

Answer: The mxE device will, using logical systems, represent the R5-1, R5-2, and P1 routers.

Step 1.6

Use the **ping 172.22.125.2 count 5 logical-system *R5-1*** command to verify that you can ping the physical interface on the P3 router.

```
lab@mxE> ping 172.22.125.2 count 5 logical-system R5-1
PING 172.22.125.2 (172.22.125.2): 56 data bytes
64 bytes from 172.22.125.2: icmp_seq=0 ttl=64 time=1.664 ms
64 bytes from 172.22.125.2: icmp_seq=1 ttl=64 time=1.387 ms
64 bytes from 172.22.125.2: icmp_seq=2 ttl=64 time=1.299 ms
64 bytes from 172.22.125.2: icmp_seq=3 ttl=64 time=1.486 ms
64 bytes from 172.22.125.2: icmp_seq=4 ttl=64 time=118.364 ms

--- 172.22.125.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.299/24.840/118.364/46.762 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.7

Access the CLI on your mxD device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device’s starting configuration by issuing the **load override *ajspr/lab6-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.4
Password: lab123
Last login: Fri Dec  6 12:53:16 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab6-start.config
load complete

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

lab@mxD>
```

Step 1.8

Issue the **show configuration** command. Use the lab diagram to verify that the mxD router has two logical systems, R2-1 and R2-2, that have the correct interface and ISIS configuration.

```
lab@mxD> show configuration
## Last commit: 2019-12-06 12:57:59 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxD;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    R2-1 {
        interfaces {
            ge-0/0/0 {
                unit 0 {
                    family inet {
                        address 172.22.121.1/24;
                    }
                    family iso;
                }
            }
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 10.0.10.1/24;
                    }
```

```
                    family iso;
                }
            }
            ge-0/0/2 {
                unit 0 {
                    family inet {
                        address 10.0.11.1/24;
                    }
                    family iso;
                }
            }
            lo0 {
                unit 1 {
                    family inet {
                        address 172.16.1.1/32;
                    }
                    family iso {
                        address 49.0001.1720.1600.1001.00;
                    }
                }
            }
        }
        protocols {
            isis {
                interface ge-0/0/0.0 {
                    level 1 disable;
                }
                interface ge-0/0/1.0 {
                    level 1 disable;
                }
                interface ge-0/0/2.0 {
                    level 1 disable;
                }
                interface lo0.1;
            }
        }
    }
    R2-2 {
        interfaces {
            ge-0/0/3 {
                unit 0 {
                    family inet {
                        address 172.22.122.1/24;
                    }
                    family iso;
                }
            }
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 10.0.14.1/24;
                    }
                    family iso;
                }
            }
```

```
        ge-0/0/5 {
            unit 0 {
                family inet {
                    address 10.0.15.1/24;
                }
                family iso;
            }
        }
        lo0 {
            unit 2 {
                family inet {
                    address 172.16.2.1/32;
                }
                family iso {
                    address 49.0002.1720.1600.2001.00;
                }
            }
        }
        protocols {
            isis {
                interface ge-0/0/3.0 {
                    level 1 disable;
                }
                interface ge-0/0/4.0 {
                    level 1 disable;
                }
                interface ge-0/0/5.0 {
                    level 1 disable;
                }
                interface lo0.2;
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.4/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications. If the configuration is incorrect, please notify your instructor.

Step 1.9

Issue the **show isis adjacency logical-system *R2-1*** command to ensure the R2-1 router has established an ISIS neighbor relationship with the P1 router.

```
lab@mxD> show isis adjacency logical-system R2-1
Interface             System         L State        Hold (secs) SNPA
ge-0/0/0.0            mxE-P1         2  Up                    8  0:50:56:a9:53:31
```

Question: Have the two routers formed an ISIS neighbor relationship with one another?

_____

Answer: Yes, the two routers should have an ISIS neighbor relationship that is in the `Up` state. You may see that R2-1 has also formed an ISIS neighbor relationship with R3-1 and/or R4-1 depending on the state of the lab.

Step 1.10

Issue the **show isis adjacency logical-system *R2-2*** command to ensure the R2-2 router has established an OSPF neighbor relationship with the P1 router.

```
lab@mxD> show isis adjacency logical-system R2-2
Interface             System         L State        Hold (secs) SNPA
ge-0/0/3.0            mxE-P1         2  Up                    7  0:50:56:a9:53:86
```

Question: Have the two routers formed an ISIS neighbor relationship with one another?

_____

Answer: Yes, the two routers should have an ISIS neighbor relationship that is in the `Up` state. You may see that R2-2 has also formed an ISIS neighbor relationship with R3-2 and/or R4-2 depending on the state of the lab.

Step 1.11

Access the CLI on your mxA device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab6-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.1
Password: lab123
Last login: Fri Dec  6 12:53:16 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# load override ajspr/lab6-start.config
load complete
```

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

lab@mxA>
```

Step 1.12

Issue the **show configuration** command. Use the lab diagram to verify that the mxA router has two logical systems, R3-1 and R4-1, that have the correct interface and ISIS configuration.

```
lab@mxA> show configuration
## Last commit: 2019-12-06 12:57:39 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxA;
    root-authentication {
        encrypted-password
"$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
.3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    R3-1 {
        interfaces {
            ge-0/0/2 {
```

```
                unit 0 {
                    family inet {
                        address 10.0.12.1/24;
                    }
                    family iso;
                }
            }
            ge-0/0/3 {
                unit 0 {
                    family inet {
                        address 10.0.10.2/24;
                    }
                    family iso;
                }
            }
            lo0 {
                unit 1 {
                    family inet {
                        address 172.16.1.2/32;
                    }
                    family iso {
                        address 49.0001.1720.1600.1002.00;
                    }
                }
            }
        }
        protocols {
            isis {
                interface ge-0/0/2.0 {
                    level 1 disable;
                }
                interface ge-0/0/3.0 {
                    level 1 disable;
                }
                interface lo0.1;
            }
        }
    }
    R4-1 {
        interfaces {
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 10.0.13.1/24;
                    }
                    family iso;
                }
            }
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 10.0.11.2/24;
                    }
                    family iso;
                }
```

```
        }
        lo0 {
            unit 2 {
                family inet {
                    address 172.16.1.3/32;
                }
                family iso {
                    address 49.0001.1720.1600.1003.00;
                }
            }
        }
    }
    protocols {
        isis {
            interface ge-0/0/1.0 {
                level 1 disable;
            }
            interface ge-0/0/4.0 {
                level 1 disable;
            }
            interface lo0.2;
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.1/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxA router. If the configuration is incorrect, please notify your instructor.

Step 1.13

Issue the **`show isis adjacency logical-system`** ***`R3-1`*** command to ensure the R3-1 router has established an ISIS neighbor relationship with R2-1 and R5-1.

```
lab@mxA> show isis adjacency logical-system R3-1
Interface             System         L State        Hold (secs) SNPA
ge-0/0/2.0            mxE-R5-1       2  Up                    7  0:50:56:a9:4b:32
ge-0/0/3.0            mxD-R2-1       2  Up                   24  0:50:56:a9:3d:65
```

Question: Has the R3-1 router formed an ISIS neighbor relationship with R2-1 and R5-1?

_____

Answer: Yes, an ISIS neighbor relationship exists between R3-1 and the other neighboring routers.

Step 1.14

Issue the **`show isis adjacency logical-system R4-1`** command to ensure the R4-1 router has established an ISIS neighbor relationship with R2-1 and R5-1.

```
lab@mxA> show isis adjacency logical-system R4-1
Interface             System         L State        Hold (secs) SNPA
ge-0/0/1.0            mxE-R5-1       2  Up                   25  0:50:56:a9:c:9e
ge-0/0/4.0            mxD-R2-1       2  Up                   21  0:50:56:a9:57:d1
```

Question: Has the R4-1 router formed an ISIS neighbor relationship with R2-1 and R5-1?

_____

Answer: Yes, an ISIS neighbor relationship exists between R4-1 and the other neighboring routers.

Step 1.15

Access the CLI on your mxC device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override ajspr/lab6-start.config`** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.3
Password: lab123
Last login: Fri Dec  6 12:53:16 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# load override ajspr/lab6-start.config
load complete

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

lab@mxC>
```

Step 1.16

Issue the **`show configuration`** command. Use the lab diagram to verify that the mxC router has two logical systems, R3-2 and R4-2, that have the correct interface and ISIS configuration.

```
lab@mxC> show configuration
## Last commit: 2017-05-08 17:12:47 UTC by lab
version 17.1R1.8;
system {
    host-name mxC;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    R3-2 {
        interfaces {
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 10.0.16.1/24;
                    }
                    family iso;
                }
            }
            ge-0/0/3 {
                unit 0 {
                    family inet {
                        address 10.0.14.2/24;
                    }
```

```
                    family iso;
                }
            }
            lo0 {
                unit 1 {
                    family inet {
                        address 172.16.2.2/32;
                    }
                    family iso {
                        address 49.0002.1720.1600.2002.00;
                    }
                }
            }
        }
        protocols {
            isis {
                interface ge-0/0/1.0 {
                    level 1 disable;
                }
                interface ge-0/0/3.0 {
                    level 1 disable;
                }
                interface lo0.1;
            }
        }
    }
    R4-2 {
        interfaces {
            ge-0/0/2 {
                unit 0 {
                    family inet {
                        address 10.0.17.1/24;
                    }
                    family iso;
                }
            }
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 10.0.15.2/24;
                    }
                    family iso;
                }
            }
            lo0 {
                unit 2 {
                    family inet {
                        address 172.16.2.3/32;
                    }
                    family iso {
                        address 49.0002.1720.1600.2003.00;
                    }
                }
            }
        }
```

```
        protocols {
            isis {
                interface ge-0/0/2.0 {
                    level 1 disable;
                }
                interface ge-0/0/4.0 {
                    level 1 disable;
                }
                interface lo0.2;
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.3/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxC router. If the configuration is incorrect, please notify your instructor.

Step 1.17

Issue the **`show isis adjacency logical-system R3-2`** command to ensure the R3-2 router has established an ISIS neighbor relationship with R2-2 and R5-2.

```
lab@mxC> show isis adjacency logical-system R3-2
Interface             System         L State        Hold (secs) SNPA
ge-0/0/1.0            mxE-R5-2       2  Up                    8  0:50:56:a9:78:33
ge-0/0/3.0            mxD-R2-2       2  Up                   24  0:50:56:a9:56:53
```

Question: Has the R3-2 router formed an ISIS neighbor relationship with R2-2 and R5-2?

_____

Answer: Yes, an ISIS neighbor relationship exists between R3-2 and the other neighboring routers.

Step 1.18

Issue the **`show isis adjacency logical-system R4-2`** command to ensure the R4-2 router has established an ISIS neighbor relationship with R2-2 and R5-2.

```
lab@mxC> show isis adjacency logical-system R4-2
Interface             System         L State        Hold (secs) SNPA
ge-0/0/2.0            mxE-R5-2       2  Up                   6  0:50:56:a9:7c:c1
ge-0/0/4.0            mxD-R2-2       2  Up                  18  0:50:56:a9:5c:39
```

Question: Has the R4-2 router formed an ISIS neighbor relationship with R2-2 and R5-2?

_____

Answer: Yes, an ISIS neighbor relationship exists between R4-2 and the other neighboring routers.

Step 1.19

Use the **show isis database logical-system *R4-2*** command to examine the current IS-IS LSDB.

```
lab@mxC> show isis database logical-system R4-2
IS-IS level 1 link-state database:
LSP ID                    Sequence Checksum Lifetime Attributes
mxC-R4-2.00-00                 0x2   0x6d98      808 L1 L2 Attached
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                    Sequence Checksum Lifetime Attributes
mxD-R2-1.00-00                 0x6   0x2c49     1050 L1 L2
mxA-R3-1.00-00                 0x5   0xf32e     1186 L1 L2
mxA-R3-1.02-00                 0x2   0xf3e1     1166 L1 L2
mxA-R4-1.00-00                 0x5   0x8f79     1171 L1 L2
mxA-R4-1.02-00                 0x2   0x1cb6     1170 L1 L2
mxA-R4-1.03-00                 0x2   0xe7e3     1171 L1 L2
mxE-R5-1.00-00                 0x5   0x15ca     1012 L1 L2
mxE-R5-1.02-00                 0x2   0x8a43     1136 L1 L2
mxD-R2-2.00-00                 0x6   0xcb2c      790 L1 L2
mxC-R3-2.00-00                 0x4   0xecb8      790 L1 L2
mxC-R3-2.02-00                 0x1   0x8cf9      767 L1 L2
mxC-R4-2.00-00                 0x4   0xa7fa      792 L1 L2
mxC-R4-2.02-00                 0x1   0xb4ce      792 L1 L2
mxE-R5-2.00-00                 0x5   0x5a1c      792 L1 L2
mxE-R5-2.02-00                 0x1   0x6913      766 L1 L2
mxE-R5-2.03-00                 0x1   0x1c61      792 L1 L2
mxE-P1.00-00                   0x5   0xadca      663 L1 L2
mxE-P1.02-00                   0x2   0xe3a3      613 L1 L2
mxE-P1.03-00                   0x2   0x2145      663 L1 L2
  19 LSPs
```

Question: All of the IS-IS adjacencies are IS-IS Level 2. Why is there an LSP in the IS-IS Level 1 database?

_____

_____

Answer: Each IS-IS router maintains a complete LSDB for each level configured. Because Level 1 has not been globally disabled on R4-2, an LSP has been created by R4-2 representing itself within the Level 1 database. However, since R4-2 has no level 1 adjacencies, the level 1 LSP has not been flooded to any neighbors.

Question: Why is the attach bit set on the Level 1 LSP?

_____

_____

Answer: The R4-2 router is aware that there is another Level 2 router (P1 for example) that belongs to a different IS-IS area, IS-IS Area 49.1234. Once R4-2 (a L1/L2 router) detects a Level 2 router in a different IS-IS area, it turns on the attach bit in all Level 1 LSPs. This setting allows all Level 1 routers to create a default route to the closest L1/L2 attached router.

Part 2: Configuring IS-IS Authentication

IS-IS has three methods of authentication: none (default), simple authentication, and MD5 authentication. In addition, IS-IS authentication can be performed at the global or interface level hierarchy. The global level authenticates all IS-IS packets, hello, link-state, and sequence number PDUs generated by the router. The interface level authentication only authenticates the hello PDU generated by the router. In this lab part, you configure MD5 hello-authentication in the Level 2 interface hierarchy between the R2-1 router and the R3-1 and R4-1 routers.

Step 2.1

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, change the perspective of the CLI to that of the R2-1 logical system by issuing the **`set cli logical-system R2-1`** command.

```
lab@mxD> set cli logical-system R2-1
Logical system: R2-1

lab@mxD:R2-1>
```

Step 2.2

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Using `md5` as the type and *`juniper`* as the key, configure Level 2 hello-authentication on the interface that faces the R3-1 router.

```
lab@mxD:R2-1> configure
Entering configuration mode
```

```
[edit]
lab@mxD:R2-1# edit protocols isis

[edit protocols isis]
lab@mxD:R2-1# set interface ge-0/0/1.0 level 2 hello-authentication-type md5

[edit protocols isis]
lab@mxD:R2-1# set interface ge-0/0/1.0 level 2 hello-authentication-key juniper

[edit protocols isis]
lab@mxD:R2-1#
```

Step 2.3

Using `md5` as the type and *`juniper`* as the key, configure Level 2 hello-authentication on the interface that is connected to the R4-1 router. Commit your configuration and exit to operational mode when you are finished.

```
[edit protocols isis]
lab@mxD:R2-1# set interface ge-0/0/2.0 level 2 hello-authentication-type md5

[edit protocols isis]
lab@mxD:R2-1# set interface ge-0/0/2.0 level 2 hello-authentication-key juniper

[edit protocols isis]
lab@mxD:R2-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R2-1>
```

Step 2.4

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxA> set cli logical-system R3-1
Logical system: R3-1

lab@mxA:R3-1>
```

Step 2.5

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Using `md5` as the type and *`juniper`* as the key, configure Level 2 hello-authentication on the interface that faces the R2-1 router. Commit your configuration and exit to operational mode when you are finished.

```
lab@mxA:R3-1> configure
Entering configuration mode

[edit]
lab@mxA:R3-1# edit protocols isis

[edit protocols isis]
lab@mxA:R3-1# set interface ge-0/0/3.0 level 2 hello-authentication-type md5
```

```
[edit protocols isis]
lab@mxA:R3-1# set interface ge-0/0/3.0 level 2 hello-authentication-key juniper

[edit protocols isis]
lab@mxA:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R3-1>
```

Step 2.6

Change the perspective of the CLI to that of the R4-1 logical system by issuing the **`set cli logical-system R4-1`** command.

```
lab@mxA:R3-1> set cli logical-system R4-1
Logical system: R4-1

lab@mxA:R4-1>
```

Step 2.7

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Using `md5` as the type and *`juniper`* as the key, configure level 2 hello-authentication on the interface that is connected to the R2-1 router. Commit the changes and return to operational mode.

```
lab@mxA:R4-1> configure
Entering configuration mode

[edit]
lab@mxA:R4-1# edit protocols isis

[edit protocols isis]
lab@mxA:R4-1# set interface ge-0/0/4.0 level 2 hello-authentication-type md5

[edit protocols isis]
lab@mxA:R4-1# set interface ge-0/0/4.0 level 2 hello-authentication-key juniper

[edit protocols isis]
lab@mxA:R4-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R4-1>
```

Step 2.8

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, use the **`show isis adjacency`** command to verify that the IS-IS adjacencies are established using the hello-authentication configuration.

```
lab@mxD:R2-1> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/0.0            mxE-P1         2  Up                    8  0:50:56:a9:53:31
ge-0/0/1.0            mxA-R3-1       2  Up                    6  0:50:56:a9:61:f9
ge-0/0/2.0            mxA-R4-1       2  Up                    8  0:50:56:a9:71:86
```

Question: What is the state of the IS-IS adjacencies with the MD5 authentication configured?

_____

_____

Answer: The two adjacencies are in the `Up` state with authentication configured.

Step 2.9

Change the perspective of the CLI to that of the default logical system by issuing the **`clear cli logical-system`** command.

```
lab@mxD:R2-1> clear cli logical-system
Cleared default logical system

lab@mxD>
```

Step 2.10

Use the **`monitor traffic interface ge-0/0/1.0 detail no-resolve`** command to verify that the IS-IS hello packets are using the MD5 authentication. Use the Ctrl+c key sequence to stop the monitor output after a couple of seconds.

```
lab@mxD> monitor traffic interface ge-0/0/1.0 detail no-resolve
Address resolution is OFF.
Listening on ge-0/0/1.0, capture size 1514 bytes

19:00:45.726595  In IS-IS, length 75
        L2 Lan IIH, hlen: 27, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
          source-id: 1720.1600.1002,  holding time: 9s, Flags: [Level 2 only]
          lan-id:    1720.1600.1002.02, Priority: 64, PDU length: 75
            IS Neighbor(s) TLV #6, length: 6
              SNPA: 0050.56a9.3d65
            Protocols supported TLV #129, length: 2
              NLPID(s): IPv4 (0xcc), IPv6 (0x8e)
            IPv4 Interface address(es) TLV #132, length: 4
              IPv4 interface address: 10.0.10.2
            Area address(es) TLV #1, length: 4
              Area address (length: 3): 49.0001
            Restart Signaling TLV #211, length: 3
              Flags [none], Remaining holding time 0s
            Authentication TLV #10, length: 17
              HMAC-MD5 password: 430b9152315ee084e476eec5cab88329
19:00:48.128914  In IS-IS, length 75
        L2 Lan IIH, hlen: 27, v: 1, pdu-v: 1, sys-id-len: 6 (0), max-area: 3 (0)
          source-id: 1720.1600.1002,  holding time: 9s, Flags: [Level 2 only]
          lan-id:    1720.1600.1002.02, Priority: 64, PDU length: 75
            IS Neighbor(s) TLV #6, length: 6
              SNPA: 0050.56a9.3d65
            Protocols supported TLV #129, length: 2
              NLPID(s): IPv4 (0xcc), IPv6 (0x8e)
            IPv4 Interface address(es) TLV #132, length: 4
              IPv4 interface address: 10.0.10.2
            Area address(es) TLV #1, length: 4
```

```
            Area address (length: 3): 49.0001
          Restart Signaling TLV #211, length: 3
            Flags [none], Remaining holding time 0s
          Authentication TLV #10, length: 17
            HMAC-MD5 password: 430b9152315ee084e476eec5cab88329
^C
2 packets received by filter
0 packets dropped by kernel
```

Question: Which TLV carries the authentication key in the hello packet?

_____

_____

Answer: The authentication key is carried in TLV #10 in the IS-IS hello packet.

Part 3: Manipulating IS-IS Metrics

Several methods of manipulating routes exist within IS-IS. By changing the metrics on the IS-IS interfaces, default routing behavior can be affected. In this lab part, you explore some of these methods.

Step 3.1

Change the perspective of the CLI to that of the R2-1 logical system by issuing the **`set cli logical-system`** ***`R2-1`*** command.

```
lab@mxD> set cli logical-system R2-1
Logical system: R2-1

lab@mxD:R2-1>
```

Step 3.2

Use the **`show isis interface`** command to examine the default metrics assigned to the IS-IS interfaces.

```
lab@mxD:R2-1> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/0.0            2   0x1 Disabled          mxE-P1.02              10/10
ge-0/0/1.0            2   0x1 Disabled          mxA-R3-1.02            10/10
ge-0/0/2.0            2   0x1 Disabled          mxA-R4-1.02            10/10
lo0.1                 3   0x1 Passive           Passive                 0/0
```

Question: What is the default Level 1 and Level 2 metric for all Gigabit Ethernet interfaces?

_____

_____

Answer: The default Level 1 and Level 2 metric is 10 for all Gigabit Ethernet interfaces. System ID names may vary due to IS-IS DIS elections.

Step 3.3

Use the **show route 172.16.1.4/32** command to display the route to the loopback interface of the R5-1 router. Note that there are two equal cost paths and the Junos OS has chosen one of them.

```
lab@mxD:R2-1> show route 172.16.1.4/32

inet.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.4/32       *[IS-IS/18] 00:17:49, metric 20
                       to 10.0.10.2 via ge-0/0/1.0
                     > to 10.0.11.2 via ge-0/0/2.0
```

Question: What is the cost to reach the R5-1 router?

_____

_____

Answer: The route has an equal cost of 20 between the two intermediate nodes.

Step 3.4

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Use the **reference-bandwidth** command to change the default metric. Use 1 gigabit as the calculating bandwidth. Commit the change when completed.

```
lab@mxD:R2-1> configure
Entering configuration mode

[edit]
lab@mxD:R2-1# edit protocols isis

[edit protocols isis]
lab@mxD:R2-1# set reference-bandwidth 1g

[edit protocols isis]
lab@mxD:R2-1# commit
commit complete

[edit protocols isis]
lab@mxD:R2-1#
```

Step 3.5

Use the **`run show isis interface`** command to verify that the change in metric of the IS-IS Level 2 interfaces.

```
[edit protocols isis]
lab@mxD:R2-1# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/0.0            2   0x1 Disabled          mxE-P1.02               1/1
ge-0/0/1.0            2   0x1 Disabled          mxA-R3-1.02             1/1
ge-0/0/2.0            2   0x1 Disabled          mxA-R4-1.02             1/1
lo0.1                 3   0x1 Passive           Passive                 0/0
```

Question: What has happened to the Level 1 and Level 2 metrics on the IS-IS interfaces?

_____

_____

Answer: The Level 1 and Level 2 metrics have changed from the default of 10 to a cost of 1. The reference-bandwidth command takes the calculating bandwidth (1 Gbps) and divides it by the static interface bandwidth (1 Gbps). The result is a cost of 1. System ID names may vary due to IS-IS DIS elections.

Step 3.6

Use the **`run show route 172.16.1.4/32`** command to display the route to the loopback interface of the R5-1 router.

```
[edit protocols isis]
lab@mxD:R2-1# run show route 172.16.1.4/32

inet.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.4/32      *[IS-IS/18] 00:02:09, metric 11
                    to 10.0.10.2 via ge-0/0/1.0
                  > to 10.0.11.2 via ge-0/0/2.0
```

Question: What is the cost to reach the R5-1 router?

_____

_____

Answer: The cost to reach the remote network has changed from 20 to 11.

Step 3.7

On the ge-0/0/2.0 interface, change the Level 2 metric to *`1000`*. Commit the change when completed.

```
[edit protocols isis]
lab@mxD:R2-1# set interface ge-0/0/2.0 level 2 metric 1000

[edit protocols isis]
lab@mxD:R2-1# commit
commit complete

[edit protocols isis]
lab@mxD:R2-1#
```

Step 3.8

Use the **`run show isis interface`** command to examine the metrics now assigned to the IS-IS interfaces.

```
[edit protocols isis]
lab@mxD:R2-1# run show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/0.0            2   0x1 Disabled          mxE-P1.02           1/1
ge-0/0/1.0            2   0x1 Disabled          mxA-R3-1.02         1/1
ge-0/0/2.0            2   0x1 Disabled          mxA-R4-1.02         1/1000
lo0.1                 3   0x1 Passive           Passive             0/0
```

Step 3.9

Use the **`run show route 172.16.1.4/32`** command to display the route to the loopback interface of the R5-1 router.

```
[edit protocols isis]
lab@mxD:R2-1# run show route 172.16.1.4/32

inet.0: 22 destinations, 23 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.4/32       *[IS-IS/18] 00:01:02, metric 11
                     > to 10.0.10.2 via ge-0/0/1.0
```

Question: What has happened to the route?

_____

_____

_____

Answer: Because of the metric change, there is now only one next hop listed to reach the R5-1 node. The configured Level 2 metric is more specific than the `reference-bandwidth` command and therefore takes precedence in the configuration. The higher metric also causes all routing to the remote network to go through the R3-1 router.

Step 3.10

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxA:R4-1> set cli logical-system R3-1
Logical system: R3-1

lab@mxA:R3-1>
```

Step 3.11

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Enable IS-IS `overload` on the R3-1 router. Commit the change and exit to operational mode when completed.

```
lab@mxA:R3-1> configure
Entering configuration mode

[edit]
lab@mxA:R3-1# edit protocols isis

[edit protocols isis]
lab@mxA:R3-1# set overload

[edit protocols isis]
lab@mxA:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R3-1>
```

Step 3.12

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, use the **`run show isis database level 2`** command to ensure that the overload bit has been enabled in the LSP.

```
[edit protocols isis]
lab@mxD:R2-1# run show isis database level 2
IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
mxD-R2-1.00-00                  0x14   0xf3e6      699 L1 L2
mxA-R3-1.00-00                  0x10   0xdd35     1078 L1 L2 Overload
mxA-R3-1.02-00                   0xc   0xdfeb      522 L1 L2
mxA-R4-1.00-00                   0xf   0xe21c      833 L1 L2
mxA-R4-1.02-00                   0xd    0x6c1      833 L1 L2
mxA-R4-1.03-00                   0xb   0xd5ec      496 L1 L2
mxE-R5-1.00-00                   0xe    0x3d3      603 L1 L2
mxE-R5-1.02-00                   0xb   0x784c      729 L1 L2
mxD-R2-2.00-00                  0x10   0xb736      741 L1 L2
mxC-R3-2.00-00                   0xd   0xdac1      429 L1 L2
mxC-R3-2.02-00                   0xb   0x7804     1107 L1 L2
mxC-R4-2.00-00                   0xe   0x9305      993 L1 L2
mxC-R4-2.02-00                   0xb   0xa0d8     1184 L1 L2
mxE-R5-2.00-00                   0xe   0x4825      666 L1 L2
```

```
mxE-R5-2.02-00                    0xa    0x571c      436 L1 L2
mxE-R5-2.03-00                    0xa     0xa6a      666 L1 L2
mxE-P1.00-00                      0xe    0x9bd3      416 L1 L2
mxE-P1.02-00                      0xc    0xcfad     1038 L1 L2
mxE-P1.03-00                      0xb     0xf4e      587 L1 L2
  19 LSPs
```

Question: How can you tell if the overload bit is enabled?

_____

_____

Answer: The word "Overload" is present in the attributes column.

Step 3.13

Verify that the next hop has changed for the route to R5-1's loopback interface using the **`run show route 172.16.1.4/32`** command.

```
[edit protocols isis]
lab@mxD:R2-1# run show route 172.16.1.4/32

inet.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.4/32       *[IS-IS/18] 00:04:06, metric 73
                     > to 10.0.11.2 via ge-0/0/2.0
```

Question: Has the next hop changed for the route to the R5-1 router's loopback address?

_____

_____

_____

Answer: Yes, the R4-1 router is now the next hop used by the R2-1 router to reach the loopback of the R5-1 router.

Question: The IS-IS Level 2 interface (ge-0/0/2.0) has a metric of 1000 configured. Why is the metric to the remote network 73?

_____

_____

_____

Answer: TLV 128 only has 6 bits it can use for the metric value of the prefix. Therefore, the highest metric value that can be advertised in the TLV is 63. The remote network cost is 63 plus 10, or 73.

Step 3.14

Issue the **`run show isis database`** command and identify the LSP ID for the R2-1 router.

```
[edit protocols isis]
lab@mxD:R2-1# run show isis database
IS-IS level 1 link-state database:
LSP ID                    Sequence Checksum Lifetime Attributes
mxD-R2-1.00-00                 0xd    0x51a      704 L1 L2 Attached
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                    Sequence Checksum Lifetime Attributes
mxD-R2-1.00-00                0x15   0xf1e7     1029 L1 L2
mxA-R3-1.00-00                0x10   0xdd35      661 L1 L2 Overload
mxA-R3-1.02-00                 0xd   0xddec      984 L1 L2
mxA-R4-1.00-00                0x10   0xe01d     1080 L1 L2
mxA-R4-1.02-00                 0xe    0x4c2     1130 L1 L2
mxA-R4-1.03-00                 0xc   0xd3ed      924 L1 L2
mxE-R5-1.00-00                 0xf    0x1d4      996 L1 L2
mxE-R5-1.02-00                 0xc   0x764d     1181 L1 L2
mxD-R2-2.00-00                0x11   0xb537     1089 L1 L2
mxC-R3-2.00-00                 0xe   0xd8c2      810 L1 L2
mxC-R3-2.02-00                 0xb   0x7804      690 L1 L2
mxC-R4-2.00-00                 0xe   0x9305      577 L1 L2
mxC-R4-2.02-00                 0xb   0xa0d8      767 L1 L2
mxE-R5-2.00-00                 0xf   0x4626      966 L1 L2
mxE-R5-2.02-00                 0xb   0x551d      780 L1 L2
mxE-R5-2.03-00                 0xb    0x86b     1120 L1 L2
mxE-P1.00-00                   0xf   0x99d4      838 L1 L2
mxE-P1.02-00                   0xc   0xcfad      621 L1 L2
mxE-P1.03-00                   0xc    0xd4f      905 L1 L2
  19 LSPs
```

Question: What is the LSP ID of the R2-1 router?

_____

_____

_____

Answer: The LSP ID of the R2-1 router is `mxD-R2-1.00-00`.

Step 3.15

Use the **`run show isis database level 2 mxD-R2-1.00-00 extensive | find tlv`** command to display the TLVs that are being advertised by the R2-1 router.

```
[edit protocols isis]
lab@mxD:R2-1# run show isis database level 2 mxD-R2-1.00-00 extensive | find
   tlv
  TLVs:
    Area address: 49.0001 (3)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
```

```
    IP router id: 172.16.1.1
    IP address: 172.16.1.1
    Hostname: mxD-R2-1
    IS neighbor: mxE-P1.02, Internal, Metric: default 1
    IS neighbor: mxA-R3-1.02, Internal, Metric: default 1
    IS neighbor: mxA-R4-1.02, Internal, Metric: default 63
    IS extended neighbor: mxE-P1.02, Metric: default 1
      IP address: 172.22.121.1
      Local interface index: 340, Remote interface index: 0
    IS extended neighbor: mxA-R3-1.02, Metric: default 1
      IP address: 10.0.10.1
      Local interface index: 341, Remote interface index: 0
    IS extended neighbor: mxA-R4-1.02, Metric: default 63
      IP address: 10.0.11.1
      Local interface index: 342, Remote interface index: 0
    IP prefix: 172.16.1.1/32, Internal, Metric: default 0, Up
    IP prefix: 172.22.121.0/24, Internal, Metric: default 1, Up
    IP prefix: 10.0.10.0/24, Internal, Metric: default 1, Up
    IP prefix: 10.0.11.0/24, Internal, Metric: default 63, Up
    IP extended prefix: 172.16.1.1/32 metric 0 up
    IP extended prefix: 172.22.121.0/24 metric 1 up
    IP extended prefix: 10.0.10.0/24 metric 1 up
    IP extended prefix: 10.0.11.0/24 metric 63 up
    Router Capability:  Router ID 172.16.1.1, Flags: 0x00
  No queued transmissions
```

Question: The IP extended prefix TLV (#135) has 4 octets for a metric value. Why is the metric value on the adjacent interface set to 63 when the metric value is configured at 1000?

_____

_____

_____

Answer: When both the IP prefix (TLV #128) and the IP extended prefix (TLV #135) are advertised, the narrow metric is always used in the SPF calculation. Therefore, both metrics are set to the maximum value of 63.

Step 3.16

Enable `wide-metrics-only` for Level 2 interfaces on the default routing instance. Commit the configuration when completed.

```
[edit protocols isis]
lab@mxD:R2-1# set level 2 wide-metrics-only

[edit protocols isis]
lab@mxD:R2-1# commit
commit complete

[edit protocols isis]
lab@mxD:R2-1#
```

Step 3.17

Use the **`run show isis database level 2 mxD-R2-1.00-00 extensive | find tlv`** command to display the TLVs that are now being advertised by the R2-1 router.

```
[edit protocols isis]
lab@mxD:R2-1# run show isis database level 2 mxD-R2-1.00-00 extensive | find
   tlv
  TLVs:
    Area address: 49.0001 (3)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 172.16.1.1
    IP address: 172.16.1.1
    Hostname: mxD-R2-1
    IS extended neighbor: mxE-P1.02, Metric: default 1
      IP address: 172.22.121.1
      Local interface index: 340, Remote interface index: 0
  IS extended neighbor: mxA-R3-1.02, Metric: default 1
      IP address: 10.0.10.1
      Local interface index: 341, Remote interface index: 0
  IS extended neighbor: mxA-R4-1.02, Metric: default 1000
      IP address: 10.0.11.1
      Local interface index: 342, Remote interface index: 0
  IP extended prefix: 172.16.1.1/32 metric 0 up
  IP extended prefix: 172.22.121.0/24 metric 1 up
  IP extended prefix: 10.0.10.0/24 metric 1 up
  IP extended prefix: 10.0.11.0/24 metric 1000 up
  Router Capability:  Router ID 172.16.1.1, Flags: 0x00
No queued transmissions
```

Question: What is the metric value in the IP extended prefix (TLV #135) on the ge-0/0/2 interface?

_____

_____

Answer: It is now 1000. The narrow TLVs (TLV #2 and TLV #128) are no longer being advertised in the LSP. Therefore, the configured metric value can be advertised within the 4-byte metric field in the extended TLV #135.

Step 3.18

Use the **`run show route 172.16.1.4/32`** command to display the route to the R5-1 router.

```
[edit protocols isis]
lab@mxD:R2-1# run show route 172.16.1.4/32

inet.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
172.16.1.4/32        *[IS-IS/18] 00:03:33, metric 1010
                      > to 10.0.11.2 via ge-0/0/2.0
```

Question: What is the cost to the remote network now?

_____

_____

Answer: Using the wide-metric TLVs in the SPF calculation results in a cost of 1010.

Part 4: Configuring IS-IS External Reachability

In this lab part, you configure external routes to be redistributed into IS-IS using routing policy. An external connection to a RIP network must be established. Once established, you create a RIP import policy to only accept routes from the RIP router that have a prefix-length of /24. These routes are redistributed into IS-IS and a default route is advertised into RIP.

Step 4.1

Return to the SSH session that you have open to the mxE router.

From the SSH session with the mxE router, change the perspective of the CLI to that of the R5-1 logical system by issuing the **`set cli logical-system R5-1`** command.

```
lab@mxE> set cli logical-system R5-1
Logical system: R5-1

lab@mxE:R5-1>
```

Step 4.2

Enter configuration mode and navigate to the `[edit protocols rip group P3]` hierarchy. Add the interface connected to the P3 router as a neighbor in the *P3* group. Commit your changes when complete.

```
lab@mxE:R5-1> configure
Entering configuration mode

[edit]
lab@mxE:R5-1# edit protocols rip group P3

[edit protocols rip group P3]
lab@mxE:R5-1# set neighbor ge-0/0/0.0

[edit protocols rip group P3]
lab@mxE:R5-1# commit
commit complete

[edit protocols rip group P3]
lab@mxE:R5-1#
```

Step 4.3

Verify that a RIP neighbor has been established with the P3 router using the **`run show rip neighbor`** command.

```
[edit protocols rip group P3]
lab@mxE:R5-1# run show rip neighbor
                  Local  Source          Destination     Send   Receive   In
Neighbor          State  Address         Address         Mode   Mode     Met
--------          -----  -------         -----------     ----   -------  ---
ge-0/0/0.0           Up 172.22.125.1    224.0.0.9       mcast  both       1
```

Question: Is RIP running on the ge-0/0/0.0 interface? What is the RIP version 2 destination multicast address?

_____

_____

Answer: Yes, RIP is running on the ge-0/0/0.0 interface. The RIP version 2 multicast destination address is 224.0.0.9.

Step 4.4

Use the **`run show route receive-protocol rip 172.22.125.2`** command to view all routes being received from the RIP router (P3).

```
[edit protocols rip group P3]
lab@mxE:R5-1# run show route receive-protocol rip 172.22.125.2

inet.0: 35 destinations, 35 routes (35 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/21       *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.0.0/24       *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.1.0/24       *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.2.0/24       *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.3.0/24       *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.4.0/25       *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.4.128/25     *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.5.0/26       *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.5.64/26      *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.5.128/26     *[RIP/100] 00:03:53, metric 2, tag 0
```

```
                    > to 172.22.125.2 via ge-0/0/0
20.20.5.192/26     *[RIP/100] 00:03:53, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: What routes have been received from the RIP router?

_____

_____

Answer: The RIP router is advertising routes in the 20.20/21 range including some routes with a /21, /24, /25, and /26 subnet range.

Step 4.5

Navigate to the `[edit policy-options policy-statement import-rip-routes]` hierarchy. Create a policy to accept only RIP routes with a prefix-length of /24. No other RIP routes should be accepted.

```
[edit protocols rip group P3]
lab@mxE:R5-1# top edit policy-options policy-statement import-rip-routes

[edit policy-options policy-statement import-rip-routes]
lab@mxE:R5-1# set term 1 from protocol rip


[edit policy-options policy-statement import-rip-routes]
lab@mxE:R5-1# set term 1 from route-filter 0/0 prefix-length-range /24-/24

[edit policy-options policy-statement import-rip-routes]
lab@mxE:R5-1# set term 1 then accept

[edit policy-options policy-statement import-rip-routes]
lab@mxE:R5-1# set term 2 then reject

[edit policy-options policy-statement import-rip-routes]
lab@mxE:R5-1# show
term 1 {
    from {
        protocol rip;
        route-filter 0.0.0.0/0 prefix-length-range /24-/24;
    }
    then accept;
}
term 2 {
    then reject;
}

[edit policy-options policy-statement import-rip-routes]
lab@mxE:R5-1#
```

Step 4.6

Navigate to the `[edit protocols rip group P3]` hierarchy. Apply the *import-rip-routes* policy as an import policy to the RIP group *P3* and commit the changes.

```
[edit policy-options policy-statement import-rip-routes]
lab@mxE:R5-1# top edit protocols rip group P3

[edit protocols rip group P3]
lab@mxE:R5-1# set import import-rip-routes

[edit protocols rip group P3]
lab@mxE:R5-1# commit
commit complete

[edit protocols rip group P3]
lab@mxE:R5-1#
```

Step 4.7

Verify that the policy is working by using the **`run show route receive-protocol rip 172.22.125.2`** command to view all routes being received from the RIP router.

```
[edit protocols rip group P3]
lab@mxE:R5-1# run show route receive-protocol rip 172.22.125.2

inet.0: 28 destinations, 28 routes (28 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/24       *[RIP/100] 00:09:46, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.1.0/24       *[RIP/100] 00:09:46, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.2.0/24       *[RIP/100] 00:09:46, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.3.0/24       *[RIP/100] 00:09:46, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Is the RIP import routing policy accepting only routes with a prefix-length of /24?

_____

_____

Answer: Yes. The import policy is accepting only the four routes with a prefix-length of /24?

Step 4.8

Navigate to the `[edit policy-options policy-statement export-rip-default]` hierarchy. Configure the policy to advertise a default route.

```
[edit protocols rip group P3]
lab@mxE:R5-1# top edit policy-options policy-statement export-rip-default

[edit policy-options policy-statement export-rip-default]
lab@mxE:R5-1# set term 1 from route-filter 0/0 exact

[edit policy-options policy-statement export-rip-default]
lab@mxE:R5-1# set term 1 then accept


[edit policy-options policy-statement export-rip-default]
lab@mxE:R5-1#
```

Step 4.9

Navigate to the `[edit protocols rip group P3]` hierarchy. Apply the *`export-rip-default`* policy as a RIP export policy under group *`P3`* and commit the changes.

```
[edit policy-options policy-statement export-rip-default]
lab@mxE:R5-1# top edit protocols rip group P3

[edit protocols rip group P3]
lab@mxE:R5-1# set export export-rip-default

[edit protocols rip group P3]
lab@mxE:R5-1# commit
commit complete

[edit protocols rip group P3]
lab@mxE:R5-1#
```

Step 4.10

Use the **`run show route advertising-protocol rip 172.22.125.1`** command to display the routes that are being advertised to the RIP router.

```
[edit protocols rip group P3]
lab@mxE:R5-1# run show route advertising-protocol rip 172.22.125.1

[edit protocols rip group P3]
lab@mxE:R5-1#
```

Question: Is the export routing policy advertising the default route?

_____

_____

Answer: According to the above capture, no routes are being advertised to the RIP network.

Step 4.11

Use the **run show route 0/0 exact** command to provide you with the reason why the default route is not being advertised to the RIP network.

```
[edit protocols rip group P3]
lab@mxE:R5-1# run show route 0/0 exact

[edit protocols rip group P3]
lab@mxE:R5-1#
```

Question: Why is no default route being redistributed into RIP?

_____

_____

Answer: Only active routes can be acted upon by export policy. A default route is not being advertised because there is no active default route in the routing table.

Step 4.12

Navigate to the top of the hierarchy and configure an aggregate default route and commit the change.

```
[edit protocols rip group P3]
lab@mxE:R5-1# top

[edit]
lab@mxE:R5-1# set routing-options aggregate route 0/0

[edit]
lab@mxE:R5-1# commit
commit complete

[edit]
lab@mxE:R5-1#
```

Step 4.13

Use the **run show route advertising-protocol rip 172.22.125.1** command to verify that the default route is now being advertised to the RIP network.

```
[edit]
lab@mxE:R5-1# run show route advertising-protocol rip 172.22.125.1
```

```
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Aggregate/130] 00:00:46
                      Reject
```

Question: Is the default route being advertised to the RIP network?

_____

_____

Answer: Yes.The default aggregate route is being advertised to the RIP network.

Step 4.14

Navigate to the `[edit policy-options policy-statement export-rip-to-isis]` hierarchy. Configure the *`export-rip-to-isis`* policy to advertise the RIP routes.

```
[edit]
lab@mxE:R5-1# edit policy-options policy-statement export-rip-to-isis

[edit policy-options policy-statement export-rip-to-isis]
lab@mxE:R5-1# set term 1 from protocol rip

[edit policy-options policy-statement export-rip-to-isis]
lab@mxE:R5-1# set term 1 then accept

[edit policy-options policy-statement export-rip-to-isis]
lab@mxE:R5-1#
```

Step 4.15

Navigate to the `[edit protocols isis]` hierarchy. Apply the *`export-rip-to-isis`* policy as an export policy to IS-IS and commit the changes and return to operational mode.

```
[edit policy-options policy-statement export-rip-to-isis]
lab@mxE:R5-1# top edit protocols isis

[edit protocols isis]
lab@mxE:R5-1# set export export-rip-to-isis

[edit protocols isis]
lab@mxE:R5-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxE:R5-1>
```

Step 4.16

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, use the **run show route 20.20/22** command to verify that the RIP routes are active in the inet.0 routing table as IS-IS external routes.

```
[edit protocols isis]
lab@mxD:R2-1# run show route 20.20/22

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

20.20.0.0/24       *[IS-IS/165] 00:02:10, metric 1012
                    > to 10.0.11.2 via ge-0/0/2.0
20.20.1.0/24       *[IS-IS/165] 00:02:10, metric 1012
                    > to 10.0.11.2 via ge-0/0/2.0
20.20.2.0/24       *[IS-IS/165] 00:02:10, metric 1012
                    > to 10.0.11.2 via ge-0/0/2.0
20.20.3.0/24       *[IS-IS/165] 00:02:10, metric 1012
                    > to 10.0.11.2 via ge-0/0/2.0
```

Question: Are the RIP routes installed in the inet.0 routing instance as active IS-IS external routes?

_____

_____

Answer: Yes, the four RIP routes are active IS-IS external routes in the inet.0 table.

Step 4.17

Log out of your assigned device using the **exit** command.

```
[edit protocols isis]
lab@mxD:R2-1# exit

[edit]
lab@mxD:R2-1# exit
Exiting configuration mode

lab@mxD:R2-1> exit


mxD (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

Network Management Diagram

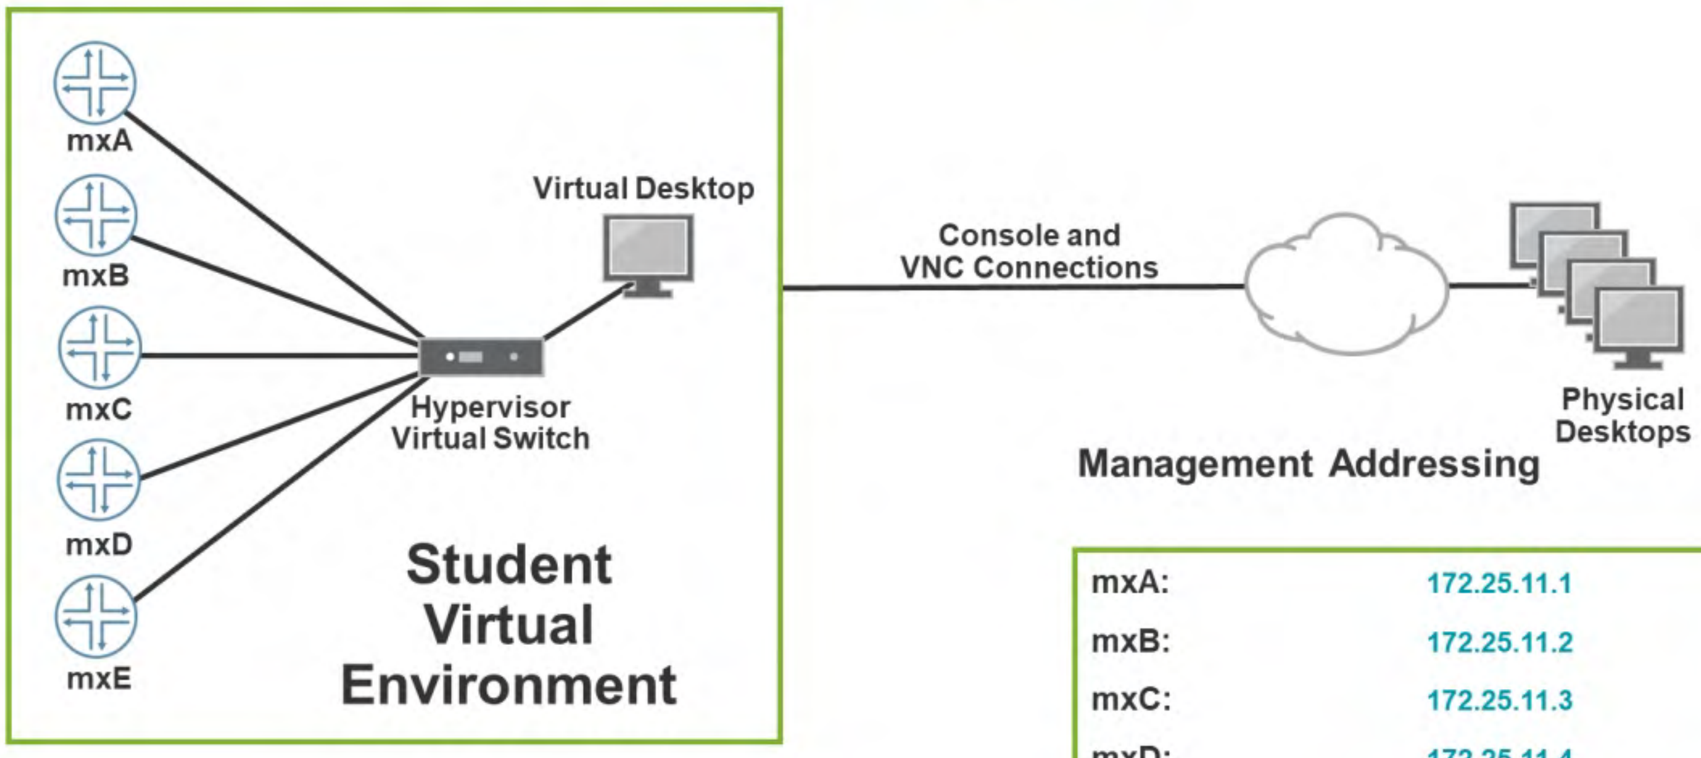

Management Addressing

mxA:	172.25.11.1
mxB:	172.25.11.2
mxC:	172.25.11.3
mxD:	172.25.11.4
mxE:	172.25.11.5
Virtual Desktop:	172.25.11.254

Note: Your instructor will provide address and access information.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab Network Diagram: Advanced IS-IS Configuration Options and Routing Policy

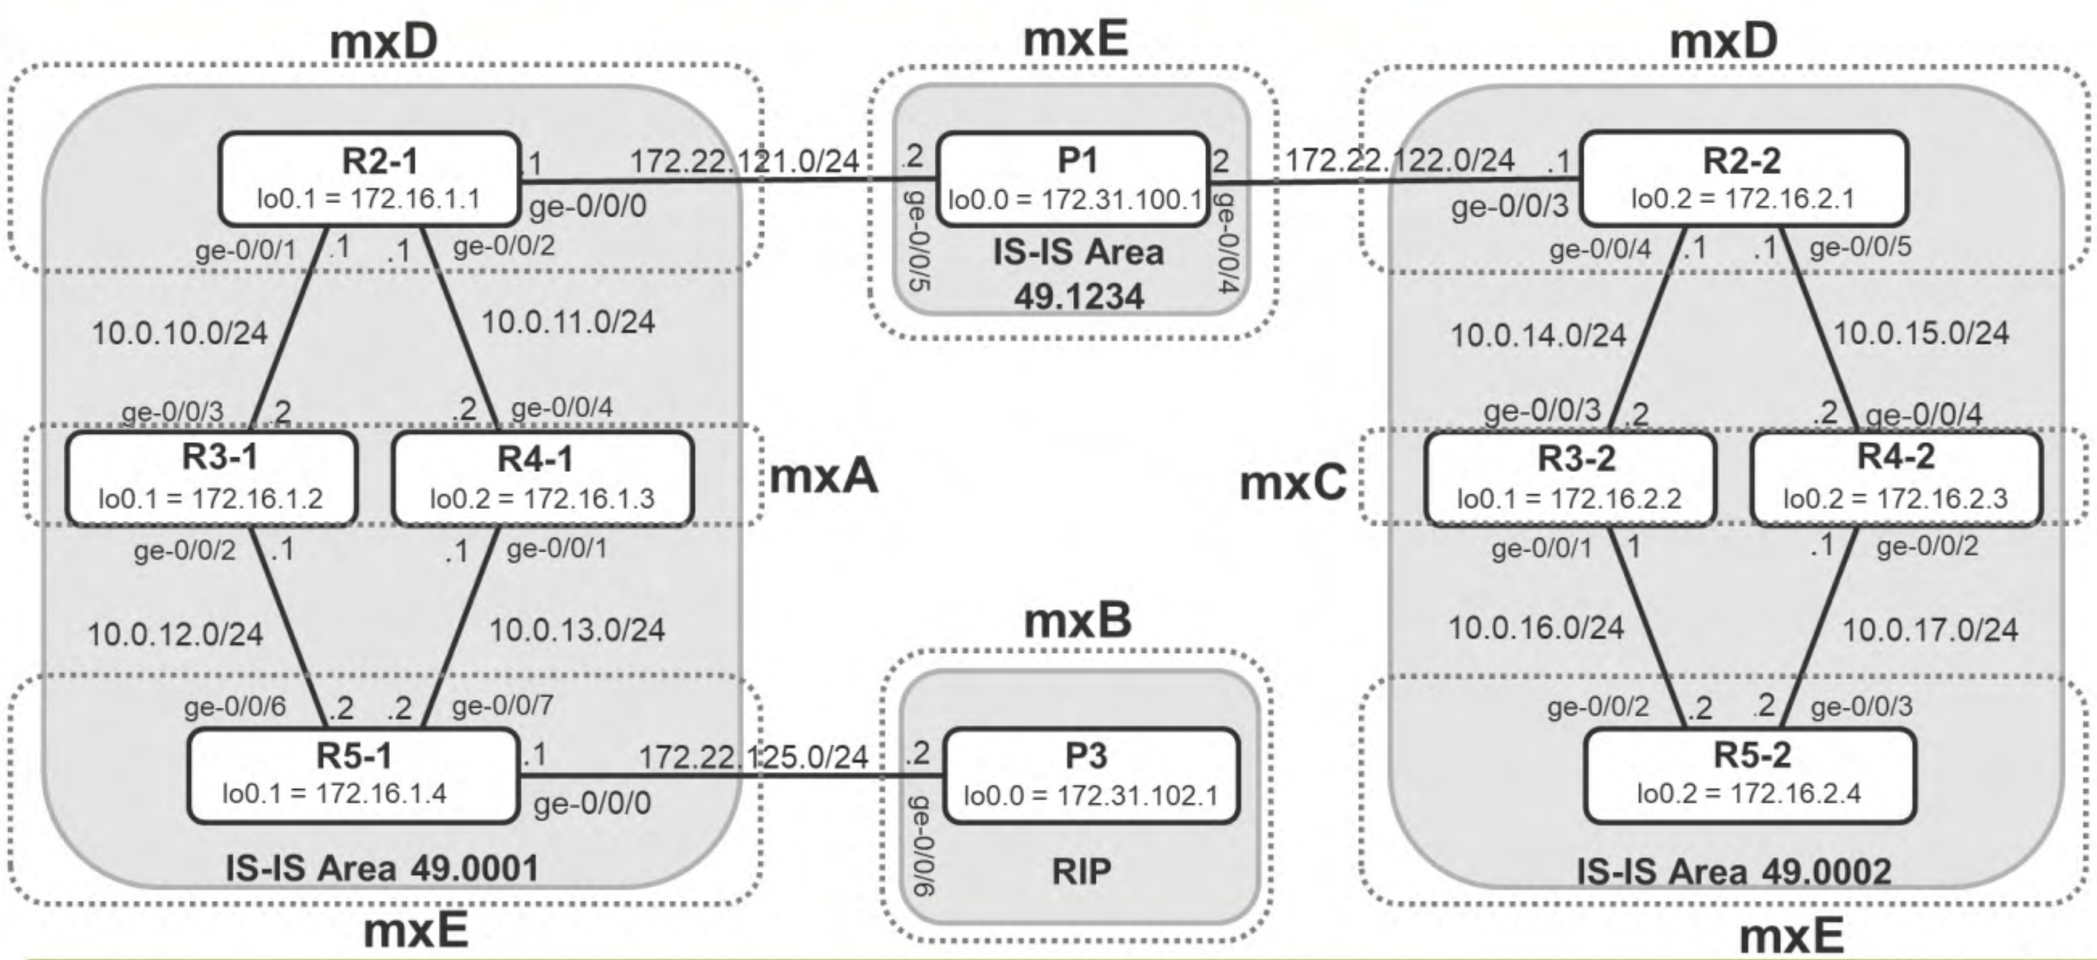

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab

Configuring a Multilevel IS-IS Network

Overview

In this lab, you use the lab diagram titled “Lab Network Diagram: Configuring a Multilevel IS-IS Network” to establish a multilevel IS-IS network. This diagram provides you with the topology to explore the default operation of a multilevel IS-IS environment, including the flooding scope through the L1/L2 attached router. It also provides you with the means to change these default operations using routing policies.

By completing this lab, you perform the following tasks:

- Establish a multilevel IS-IS network.
- Explore the default flooding scope between L1 and L2.
- Explore the use of the attach bit.
- Use routing policy to summarize routes from L1 to L2.
- Redistribute external routes from L1 to L2.
- Use routing policies to leak routes from L2 to L1.

Part 1: Load Starting Configuration

In this lab part, you verify the initial configuration of the routers. You then verify that the interfaces are operational and that IS-IS neighbor relationships have formed.

Note

The instructor will tell you the nature of your access and will provide you with the necessary details to access your assigned device.

Note

The lab topology requires you to display information in the different logical systems of mxA, mxB, mxC, mxD, and mxE. You will be changing the perspective of the CLI by issuing the **`set cli logical-system`** ***`ls-name`*** command. Refer to the lab diagram for the correct logical system name.

By changing the perspective of the CLI, you will be able to perform network commands such as **`ping`** or **`traceroute`** from the perspective of the pertinent logical system.

Step 1.1

Ensure that you know to which student device you have been assigned. Check with your instructor if you are not certain. Consult the management network diagram to determine the management address of the student devices.

Access the CLI on your mxB device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab7-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.2
Password: lab123
Last login: Fri Dec  6 12:59:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxB> configure
Entering configuration mode

[edit]
lab@mxB# load override ajspr/lab7-start.config
load complete

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

lab@mxB>
```

Step 1.2

Issue the **show configuration** command. Use the lab diagram to verify that the mxB router has the correct logical system and interface configuration. Verify that IS-IS has been enabled (all interfaces are at level 2).

```
lab@mxB> show configuration
## Last commit: 2019-12-06 13:50:43 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxB;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    P3 {
        interfaces {
            ge-0/0/6 {
                unit 0 {
                    family inet {
                        address 172.22.125.2/24;
                    }
                }
            }
            lo0 {
```

```
                unit 0 {
                    family inet {
                        address 172.31.102.1/32;
                        address 20.20.1.1/32;
                        address 40.40.1.1/32;
                    }
                }
            }
        }
        protocols {
            rip {
                group rip {
                    export export-rip-policy;
                    neighbor ge-0/0/6.0;
                }
            }
        }
        policy-options {
            policy-statement export-rip-policy {
                term 1 {
                    from {
                        protocol [ static rip aggregate ];
                        route-filter 20.20.0.0/21 orlonger;
                    }
                    then accept;
                }
                term 2 {
                    from {
                        protocol rip;
                        route-filter 0.0.0.0/0 exact;
                    }
                    then accept;
                }
            }
        }
        routing-options {
            static {
                route 20.20.0.0/24 reject;
                route 20.20.1.0/24 reject;
                route 20.20.2.0/24 reject;
                route 20.20.3.0/24 reject;
                route 20.20.4.0/25 reject;
                route 20.20.4.128/25 reject;
                route 20.20.5.0/26 reject;
                route 20.20.5.64/26 reject;
                route 20.20.5.128/26 reject;
                route 20.20.5.192/26 reject;
                route 40.40.0.0/24 reject;
                route 40.40.1.0/24 reject;
                route 40.40.2.0/24 reject;
                route 40.40.3.0/24 reject;
                route 172.16.1.2/32 next-hop 172.22.125.1;
                route 172.16.2.2/32 next-hop 172.22.126.1;
                route 172.16.1.4/32 next-hop 172.22.125.1;
                route 172.16.2.4/32 next-hop 172.22.126.1;
```

```
                route 10.0.20.0/22 next-hop 172.22.125.1;
                route 10.0.24.0/22 next-hop 172.22.126.1;
            }
            aggregate {
                route 20.20.0.0/21;
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.2/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxB router. If the configuration is incorrect, please notify your instructor.

Step 1.3

Access the CLI on your mxE device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab7-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.5
Password: lab123
Last login: Fri Dec  6 12:59:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxE> configure
Entering configuration mode

[edit]
lab@mxE# load override ajspr/lab7-start.config
load complete

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

lab@mxE>
```

Step 1.4

Issue the **show configuration** command. Use the lab diagram to verify that the mxE router has the correct logical system and interface configuration. Verify that IS-IS has been enabled (all interfaces at level 2).

```
lab@mxE> show configuration
## Last commit: 2019-12-27 10:11:30 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxE;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    P1 {
        interfaces {
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 172.22.122.2/24;
                    }
                    family iso;
                }
            }
```

```
            ge-0/0/5 {
                unit 0 {
                    family inet {
                        address 172.22.121.2/24;
                    }
                    family iso;
                }
            }
            lo0 {
                unit 0 {
                    family inet {
                        address 172.31.100.1/32;
                    }
                    family iso {
                        address 49.1234.1720.3110.0001.00;
                    }
                }
            }
        }
        protocols {
            isis {
                interface ge-0/0/4.0 {
                    level 1 disable;
                }
                interface ge-0/0/5.0 {
                    level 1 disable;
                }
                interface lo0.0;
            }
        }
        routing-options {
            autonomous-system 65412;
        }
    }
    R5-1 {
        interfaces {
            ge-0/0/0 {
                unit 0 {
                    family inet {
                        address 172.22.125.1/24;
                    }
                }
            }
            ge-0/0/6 {
                unit 0 {
                    family inet {
                        address 10.0.12.2/24;
                    }
                    family iso;
                }
            }
            ge-0/0/7 {
                unit 0 {
                    family inet {
                        address 10.0.13.2/24;
```

```
                    }
                    family iso;
                }
            }
            lo0 {
                unit 1 {
                    family inet {
                        address 172.16.1.4/32;
                    }
                    family iso {
                        address 49.0001.1720.1600.1004.00;
                    }
                }
            }
        }
        protocols {
            isis {
                interface ge-0/0/6.0 {
                    level 2 disable;
                }
                interface ge-0/0/7.0 {
                    level 2 disable;
                }
                interface lo0.1;
                export export-rip-to-isis;
            }
            rip {
                group P3 {
                    export export-rip-default;
                    import import-rip-routes;
                    neighbor ge-0/0/0.0;
                }
            }
        }
        policy-options {
            policy-statement export-rip-default {
                term 1 {
                    from {
                        route-filter 0.0.0.0/0 exact;
                    }
                    then accept;
                }
            }
            policy-statement export-rip-to-isis {
                term 1 {
                    from protocol rip;
                    then {
                        external {
                            type 1;
                        }
                        accept;
                    }
                }
            }
            policy-statement import-rip-routes {
```

```
            term 1 {
                from {
                    protocol rip;
                    route-filter 0.0.0.0/0 prefix-length-range /24-/24;
                }
                then accept;
            }
            term 2 {
                then reject;
            }
        }
    }
}
R5-2 {
    interfaces {
        ge-0/0/2 {
            unit 0 {
                family inet {
                    address 10.0.16.2/24;
                }
                family iso;
            }
        }
        ge-0/0/3 {
            unit 0 {
                family inet {
                    address 10.0.17.2/24;
                }
                family iso;
            }
        }
        lo0 {
            unit 2 {
                family inet {
                    address 172.16.2.4/32;
                }
                family iso {
                    address 49.0002.1720.1600.2004.00;
                }
            }
        }
    }
    protocols {
        isis {
            interface ge-0/0/2.0 {
                level 2 disable;
            }
            interface ge-0/0/3.0 {
                level 2 disable;
            }
            interface lo0.2;
        }
    }
}
```

```
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.5/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxE router. If the configuration is incorrect, please notify your instructor.

Question: Based on the configuration, which routers (logical systems) will mxE represent in this lab?

_____

Answer: The mxE device will, using logical systems, represent the R5-1, R5-2, and P1 routers.

Question: There have been three policies configured and applied to the R5-1 logical system. What do the policies accomplish?

_____

_____

_____

_____

_____

_____

_____

Answer: The R5-1 logical system has an import policy applied to its RIP neighbor. The policy only allows /24 routes to be accepted into the routing table. The R5-1 logical system has an export policy applied to its RIP neighbor. This policy advertises a default route to its RIP neighbor. The R5-1 logical system has an export policy applied to the IS-IS protocol. This policy redistributes RIP routes into IS-IS.

Step 1.5

Use the **ping 172.22.125.2 count 5 logical-system *R5-1*** command to verify that you can ping the physical interface on the P3 router.

```
lab@mxE> ping 172.22.125.2 count 5 logical-system R5-1
PING 172.22.125.2 (172.22.125.2): 56 data bytes
64 bytes from 172.22.125.2: icmp_seq=0 ttl=64 time=1.664 ms
64 bytes from 172.22.125.2: icmp_seq=1 ttl=64 time=1.387 ms
64 bytes from 172.22.125.2: icmp_seq=2 ttl=64 time=1.299 ms
64 bytes from 172.22.125.2: icmp_seq=3 ttl=64 time=1.486 ms
64 bytes from 172.22.125.2: icmp_seq=4 ttl=64 time=118.364 ms

--- 172.22.125.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.299/24.840/118.364/46.762 ms
```

Question: Were the pings successful?

_____

Answer: Yes, the pings should be successful. If not, verify that the interfaces have the correct IP addresses and are “up” and “up”.

Step 1.6

Access the CLI on your mxD device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device’s starting configuration by issuing the **load override *ajspr/lab7-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.4
Password: lab123
Last login: Fri Dec  6 12:59:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab7-start.config
load complete

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

lab@mxD>
```

Step 1.7

Issue the **show configuration** command. Use the lab diagram to verify that the mxD router has two logical systems, R2-1 and R2-2, that have the correct interface and IS-IS configuration.

```
lab@mxD> show configuration
## Last commit: 2019-12-06 14:04:21 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxD;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    R2-1 {
        interfaces {
            ge-0/0/0 {
                unit 0 {
                    family inet {
                        address 172.22.121.1/24;
                    }
                    family iso;
                }
            }
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 10.0.10.1/24;
                    }
```

```
                family iso;
            }
        }
        ge-0/0/2 {
            unit 0 {
                family inet {
                    address 10.0.11.1/24;
                }
                family iso;
            }
        }
        lo0 {
            unit 1 {
                family inet {
                    address 172.16.1.1/32;
                }
                family iso {
                    address 49.0001.1720.1600.1001.00;
                }
            }
        }
    }
    protocols {
        isis {
            interface ge-0/0/0.0 {
                level 1 disable;
            }
            interface ge-0/0/1.0 {
                level 1 disable;
            }
            interface ge-0/0/2.0 {
                level 1 disable;
            }
            interface lo0.1;
        }
    }
}
R2-2 {
    interfaces {
        ge-0/0/3 {
            unit 0 {
                family inet {
                    address 172.22.122.1/24;
                }
                family iso;
            }
        }
        ge-0/0/4 {
            unit 0 {
                family inet {
                    address 10.0.14.1/24;
                }
                family iso;
            }
        }
```

```
            ge-0/0/5 {
                unit 0 {
                    family inet {
                        address 10.0.15.1/24;
                    }
                    family iso;
                }
            }
            lo0 {
                unit 2 {
                    family inet {
                        address 172.16.2.1/32;
                    }
                    family iso {
                        address 49.0002.1720.1600.2001.00;
                    }
                }
            }
        }
        protocols {
            isis {
                interface ge-0/0/3.0 {
                    level 1 disable;
                }
                interface ge-0/0/4.0 {
                    level 1 disable;
                }
                interface ge-0/0/5.0 {
                    level 1 disable;
                }
                interface lo0.2;
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.4/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications. If the configuration is incorrect, please notify your instructor.

Step 1.8

Issue the **show isis adjacency logical-system *R2-1*** command to ensure the R2-1 router has established an IS-IS neighbor relationship with the P1 router.

```
lab@mxD> show isis adjacency logical-system R2-1
Interface             System         L State        Hold (secs) SNPA
ge-0/0/0.0            mxE-P1         2  Up                    8  0:50:56:a9:53:31
```

Question: Have the two routers formed an IS-IS neighbor relationship with one another? At what level is the adjacency?

_____

Answer: Yes, the two routers should have a level 2 IS-IS neighbor relationship that is in the `Up` state.

Note

You might see that R2-1 has also formed an IS-IS neighbor relationship with R3-1 and/or R4-1 depending on the state of the lab.

Step 1.9

Issue the **show isis adjacency logical-system *R2-2*** command to ensure the R2-2 router has established an OSPF neighbor relationship with the P1 router.

```
lab@mxD> show isis adjacency logical-system R2-2
Interface             System         L State        Hold (secs) SNPA
ge-0/0/3.0            mxE-P1         2  Up                    7  0:50:56:a9:53:86
```

Question: Have the two routers formed an IS-IS neighbor relationship with one another? At what level is the adjacency?

_____

_____

Answer: Yes, the two routers should have a level 2 IS-IS neighbor relationship that is in the `Up` state.

Question:

Note

You might see that R2-2 has also formed an IS-IS neighbor relationship with R3-2 and/or R4-2 depending on the state of the lab.

Step 1.10

Access the CLI on your mxA device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab7-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode..

```
[lab@desktop ~]$ ssh lab@172.25.11.1
Password: lab123
```

```
Last login: Fri Dec  6 12:59:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# load override ajspr/lab7-start.config
load complete

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

lab@mxA>
```

Step 1.11

Issue the **show configuration** command. Use the lab diagram to verify that the mxA router has two logical systems, R3-1 and R4-1, that have the correct interface and IS-IS configuration.

```
lab@mxA> show configuration
## Last commit: 2019-12-06 13:50:43 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxA;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
```

```
                interactive-commands any;
            }
        }
    }
    logical-systems {
        R3-1 {
            interfaces {
                ge-0/0/2 {
                    unit 0 {
                        family inet {
                            address 10.0.12.1/24;
                        }
                        family iso;
                    }
                }
                ge-0/0/3 {
                    unit 0 {
                        family inet {
                            address 10.0.10.2/24;
                        }
                        family iso;
                    }
                }
                lo0 {
                    unit 1 {
                        family inet {
                            address 172.16.1.2/32;
                        }
                        family iso {
                            address 49.0001.1720.1600.1002.00;
                        }
                    }
                }
            }
            protocols {
                isis {
                    interface ge-0/0/2.0 {
                        level 2 disable;
                    }
                    interface ge-0/0/3.0 {
                        level 1 disable;
                    }
                    interface lo0.1;
                }
            }
        }
        R4-1 {
            interfaces {
                ge-0/0/1 {
                    unit 0 {
                        family inet {
                            address 10.0.13.1/24;
                        }
                        family iso;
                    }
```

```
                }
                ge-0/0/4 {
                    unit 0 {
                        family inet {
                            address 10.0.11.2/24;
                        }
                        family iso;
                    }
                }
                lo0 {
                    unit 2 {
                        family inet {
                            address 172.16.1.3/32;
                        }
                        family iso {
                            address 49.0001.1720.1600.1003.00;
                        }
                    }
                }
            }
            protocols {
                isis {
                    interface ge-0/0/1.0 {
                        level 2 disable;
                    }
                    interface ge-0/0/4.0 {
                        level 1 disable;
                    }
                    interface lo0.2;
                }
            }
        }
    }
    interfaces {
        fxp0 {
            unit 0 {
                family inet {
                    address 172.25.11.1/24;
                }
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxA router. If the configuration is incorrect, please notify your instructor.

Step 1.12

Issue the **show isis adjacency logical-system *R3-1*** command to ensure the R3-1 router has established an IS-IS neighbor relationship with R2-1 and R5-1.

```
lab@mxA> show isis adjacency logical-system R3-1
Interface             System         L State        Hold (secs) SNPA
ge-0/0/2.0            mxE-R5-1       1  Up                    8  0:50:56:a9:4b:32
ge-0/0/3.0            mxD-R2-1       2  Up                   20  0:50:56:a9:3d:65
```

Question: Has the R3-1 router formed an IS-IS neighbor relationship with R2-1 and R5-1? At what level are the adjacencies?

_____

Answer: Yes, an level 1 IS-IS neighbor relationship exists between R3-1 and R5-1 while a level 2 IS-IS neighbor relationship exists between R3-1 and R2-1.

Step 1.13

Issue the **show isis adjacency logical-system *R4-1*** command to ensure the R4-1 router has established an IS-IS neighbor relationship with R2-1 and R5-1.

```
lab@mxA> show isis adjacency logical-system R4-1
Interface             System         L State        Hold (secs) SNPA
ge-0/0/1.0            mxE-R5-1       1  Up                   22  0:50:56:a9:c:9e
ge-0/0/4.0            mxD-R2-1       2  Up                   18  0:50:56:a9:57:d1
```

Question: Has the R4-1 router formed an IS-IS neighbor relationship with R2-1 and R5-1? At what level are the adjacencies?

_____

Answer: Yes, a level 1 IS-IS neighbor relationship exists between R4-1 and R5-1 while a level 2 IS-IS neighbor relationship exists between R4-1 and R2-1.

Step 1.14

Access the CLI on your mxC device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab7-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.3
Password: lab123
Last login: Fri Dec  6 12:59:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxC> configure
Entering configuration mode
```

```
[edit]
lab@mxC# load override ajspr/lab7-start.config
load complete

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

lab@mxC>
```

Step 1.15

Issue the **show configuration** command. Use the lab diagram to verify that the mxC router has two logical systems, R3-2 and R4-2, that have the correct interface and IS-IS configuration.

```
lab@mxC> show configuration
## Last changed: 2019-12-06 14:46:02 UTC
version 20190829.221548_builder.r1052644;
system {
    host-name mxC;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
```

```
R3-2 {
    interfaces {
        ge-0/0/1 {
            unit 0 {
                family inet {
                    address 10.0.16.1/24;
                }
                family iso;
            }
        }
        ge-0/0/3 {
            unit 0 {
                family inet {
                    address 10.0.14.2/24;
                }
                family iso;
            }
        }
        lo0 {
            unit 1 {
                family inet {
                    address 172.16.2.2/32;
                }
                family iso {
                    address 49.0002.1720.1600.2002.00;
                }
            }
        }
    }
    protocols {
        isis {
            interface ge-0/0/1.0 {
                level 2 disable;
            }
            interface ge-0/0/3.0 {
                level 1 disable;
            }
            interface lo0.1;
        }
    }
}
R4-2 {
    interfaces {
        ge-0/0/2 {
            unit 0 {
                family inet {
                    address 10.0.17.1/24;
                }
                family iso;
            }
        }
        ge-0/0/4 {
            unit 0 {
                family inet {
                    address 10.0.15.2/24;
```

```
                    }
                    family iso;
                }
            }
            lo0 {
                unit 2 {
                    family inet {
                        address 172.16.2.3/32;
                    }
                    family iso {
                        address 49.0002.1720.1600.2003.00;
                    }
                }
            }
        }
        protocols {
            isis {
                interface ge-0/0/2.0 {
                    level 2 disable;
                }
                interface ge-0/0/4.0 {
                    level 1 disable;
                }
                interface lo0.2;
            }
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.3/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxC router. If the configuration is incorrect, please notify your instructor.

Step 1.16

Issue the **`show isis adjacency logical-system R3-2`** command to ensure the R3-2 router has established an IS-IS neighbor relationship with R2-2 and R5-2.

```
lab@mxC> show isis adjacency logical-system R3-2
Interface             System         L State        Hold (secs) SNPA
ge-0/0/1.0            mxE-R5-2       1  Up                    7  0:50:56:a9:78:33
ge-0/0/3.0            mxD-R2-2       2  Up                   21  0:50:56:a9:56:53
```

Question: Has the R3-2 router formed an IS-IS neighbor relationship with R2-2 and R5-2? At what level are the adjacencies?

_____

Answer: Yes, a level 1 IS-IS neighbor relationship exists between R3-2 and R5-2 while a level 2 IS-IS neighbor relationship exists between R3-2 and R2-2.

Step 1.17

Issue the **`show isis adjacency logical-system R4-2`** command to ensure the R4-2 router has established an IS-IS neighbor relationship with R2-2 and R5-2.

```
lab@mxC> show isis adjacency logical-system R4-2
Interface             System         L State        Hold (secs) SNPA
ge-0/0/2.0            mxE-R5-2       1  Up                    7  0:50:56:a9:7c:c1
ge-0/0/4.0            mxD-R2-2       2  Up                   26  0:50:56:a9:5c:39
```

Question: Has the R4-2 router formed an IS-IS neighbor relationship with R2-2 and R5-2? At what level are the adjacencies?

_____

Answer: Yes, a level 1 IS-IS neighbor relationship exists between R4-2 and R5-2 while a level 2 IS-IS neighbor relationship exists between R4-2 and R2-2.

Part 2: Examining the IS-IS Multilevel Flooding Scope

In this lab part, you examine the default flooding scope of a multilevel IS-IS topology and then modify that default flooding scope. The IS-IS multilevel topology supports multiple IS-IS databases, a Level 1 database and a Level 2 database. Because the R3-*X* and R4-*X* logical systems have interfaces connected to both a Level 1 and Level 2, they have both databases. Information that is passed between these databases is known as the IS-IS flooding scope. IS-IS, by default, floods all Level 1 internal information into the Level 2 database but does not flood Level 2 internal information into the Level 1 database. Also, by default, IS-IS does not flood any external information between levels.

Step 2.1

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, change the perspective of the CLI to that of the R2-1 logical system by issuing the **`set cli logical-system R2-1`** command.

```
lab@mxD> set cli logical-system R2-1
Logical system: R2-1

lab@mxD:R2-1>
```

Step 2.2

The R2-1 router only has IS-IS Level 2 adjacencies. Use the **`show route 172.16/16`** command to display all active loopback addresses in the Level 2 database.

```
lab@mxD:R2-1> show route 172.16/16

inet.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.1/32       *[Direct/0] 00:34:06
                     > via lo0.1
172.16.1.2/32       *[IS-IS/18] 00:31:39, metric 10
                     > to 10.0.10.2 via ge-0/0/1.0
172.16.1.3/32       *[IS-IS/18] 00:31:38, metric 10
                     > to 10.0.11.2 via ge-0/0/2.0
172.16.1.4/32       *[IS-IS/18] 00:31:38, metric 20
                       to 10.0.10.2 via ge-0/0/1.0
                     > to 10.0.11.2 via ge-0/0/2.0
172.16.2.1/32       *[IS-IS/18] 00:32:54, metric 20
                     > to 172.22.121.2 via ge-0/0/0.0
172.16.2.2/32       *[IS-IS/18] 00:28:02, metric 30
                     > to 172.22.121.2 via ge-0/0/0.0
172.16.2.3/32       *[IS-IS/18] 00:28:03, metric 30
                     > to 172.22.121.2 via ge-0/0/0.0
172.16.2.4/32       *[IS-IS/18] 00:28:02, metric 40
                     > to 172.22.121.2 via ge-0/0/0.0
```

Question: Within the multilevel IS-IS topology, are any loopback addresses missing from the R2-1 router's routing table?

_____

_____

Answer: No. All loopback addresses are active in the routing table as IS-IS Level 2 routes with a routing preference of 18. The R5-1 and R5-2 loopback addresses have been leaked into the Level 2 database through the L1/L2 attached routers (R3-1, R4-1, R3-2, and R4-2).

Step 2.3

Return to the SSH session that you have open to the mxE router.

From the SSH session with the mxE router, change the perspective of the CLI to that of the R5-1 logical system by issuing the **`set cli logical-system R5-1`** command.

```
lab@mxE> set cli logical-system R5-1
Logical system: R5-1

lab@mxE:R5-1>
```

Step 2.4

The R5-1 router only has IS-IS Level 1 adjacencies. Use the **`show route 172.16/16`** command to display all active loopback address in the Level 1 database.

```
lab@mxE:R5-1> show route 172.16/16

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.2/32       *[IS-IS/15] 00:00:57, metric 10
                     >  to 10.0.12.1 via ge-0/0/6.0
172.16.1.3/32       *[IS-IS/15] 00:00:57, metric 10
                     >  to 10.0.13.1 via ge-0/0/7.0
172.16.1.4/32       *[Direct/0] 00:13:32
                     >  via lo0.1
```

Question: How many loopback addresses are active IS-IS routes in the routing table? Do any of the loopback addresses from Level 2 exist in the routing table? Why?

_____

_____

Answer: The routing table only has three loopback addresses listed and none of the routes are from the IS-IS Level 2 database. By default, routes are not leaked from the Level 2 database into the Level 1 database.

Step 2.5

Because the internal Level 2 IS-IS routes are not leaked into the Level 1 database, access from the R5-1 router to the Level 2 networks requires a default route. Use the **`show route 0/0 exact`** command to display the active default route in the R5-1 router's routing table.

```
lab@mxE:R5-1> show route 0/0 exact

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0           *[IS-IS/15] 00:01:40, metric 10
                     >  to 10.0.12.1 via ge-0/0/6.0
                        to 10.0.13.1 via ge-0/0/7.0
```

Question: Is the default route an active IS-IS route in the routing table? Is it an internal or external route?

_____

_____

Answer: The default route is an active internal Level 1 IS-IS route with a routing preference of 15.

Step 2.6

All L1/L2 attached routers (R3-1, R4-1, R3-2, and R4-2) enable the attach bit on their Level 1 LSP if they are able to detect that a level 2 router exists, P1 for example, in a different IS-IS area than their own. Display the Level 1 database using the **`show isis database level 1`** command. The command output may vary.

```
lab@mxE:R5-1> show isis database level 1
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
mxA-R3-1.00-00                 0x30d   0x837f     1185 L1 L2 Attached
mxA-R4-1.00-00                 0x307   0x5e9d      996 L1 L2 Attached
mxA-R4-1.02-00                   0x4   0xeadf      878 L1 L2
mxE-R5-1.00-00                 0x30e   0x32a3     1049 L1 L2
mxE-R5-1.02-00                   0x7   0x8048      941 L1 L2
  5 LSPs
```

Question: Which routers have enabled the attach bit in their Level 1 LSP?

_____

_____

Answer: The R3-1 and R4-1 routers have enabled the attach bit in their Level 1 LSP. Because of this, every router in the Level 1 domain creates a default route in its own routing table to the closest L1/L2 attached router.

Question: What is the primary LSP ID for the R5-1 router in the IS-IS Level 1 database?

_____

_____

Answer: Based on the previous output, the LSP-ID for the R5-1 router is `mxE-R5-1.00-00`.

Step 2.7

Use the **`show route protocol rip`** command to verify that RIP routes are being received from the P3 RIP network.

```
lab@mxE:R5-1> show route protocol rip

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/24       *[RIP/100] 00:06:57, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.1.0/24       *[RIP/100] 00:06:57, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.2.0/24       *[RIP/100] 00:06:57, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
20.20.3.0/24       *[RIP/100] 00:06:57, metric 2, tag 0
                    > to 172.22.125.2 via ge-0/0/0.0
224.0.0.9/32       *[RIP/100] 00:07:02, metric 1
                      MultiRecv

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: How many active RIP routes exist in the R5-1 router's routing table?

_____

_____

Answer: There are four /24 routes and one RIP muliticast route in the R5-1 router's routing table.

Step 2.8

Use the primary LSP-ID determined previously to display the IS-IS database for the R5-1 router's link-state PDU to verify that the RIP routes have been redistributed into IS-IS as external Type 130 TLVs. Use the **`show isis database level 1 mxE-R5-1.00-00 extensive | find tlv`** command to display these TLVs.

```
lab@mxE:R5-1> show isis database level 1 mxE-R5-1.00-00 extensive | find tlv
TLVs:
    Area address: 49.0001 (3)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 172.16.1.4
    IP address: 172.16.1.4
    Hostname: mxE-R5-1
    IP prefix: 10.0.12.0/24, Internal, Metric: default 10, Up
    IP prefix: 10.0.13.0/24, Internal, Metric: default 10, Up
    IP prefix: 172.16.1.4/32, Internal, Metric: default 0, Up
    IP extended prefix: 10.0.12.0/24 metric 10 up
    IP extended prefix: 10.0.13.0/24 metric 10 up
    IP extended prefix: 172.16.1.4/32 metric 0 up
    Router Capability:  Router ID 172.16.1.4, Flags: 0x00
    IP external prefix: 20.20.0.0/24, Internal, Metric: default 2, Up
    IP external prefix: 20.20.1.0/24, Internal, Metric: default 2, Up
    IP external prefix: 20.20.2.0/24, Internal, Metric: default 2, Up
    IP external prefix: 20.20.3.0/24, Internal, Metric: default 2, Up
    IP extended prefix: 20.20.0.0/24 metric 2 up
    IP extended prefix: 20.20.1.0/24 metric 2 up
    IP extended prefix: 20.20.2.0/24 metric 2 up
    IP extended prefix: 20.20.3.0/24 metric 2 up
    IS neighbor: mxE-R5-1.02, Internal, Metric: default 10
    IS neighbor: mxE-R5-1.03, Internal, Metric: default 10
    Extended IS Reachability TLV, Type: 22, Length: 54
    IS extended neighbor: mxE-R5-1.02, Metric: default 10 SubTLV len: 16
      IP address: 10.0.12.2
      Local interface index: 333, Remote interface index: 0
    IS extended neighbor: mxE-R5-1.03, Metric: default 10 SubTLV len: 16
      IP address: 10.0.13.2
      Local interface index: 337, Remote interface index: 0
  No queued transmissions
```

Question: Are the RIP routes installed in the R5-1 router's IS-IS LSP as external Type 130 TLVs?

_____

_____

Answer: The four /24 RIP routes are installed in the R5-1 router's LSP as Type 130 TLVs (IP external prefix) and as Type 135 TLVs (IP extended prefix). Because both the external and extended TLVs exist, only the TLV 130 values are used in the SPF algorithm.

Step 2.9

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, display the RIP routes in the R2-1 router's routing table using the **show route 20.20/22** command.

```
lab@mxD:R2-1> show route 20.20/22

lab@mxD:R2-1>
```

Question: Are the RIP routes present in the R2-1 router's routing table? Why?

_____

_____

Answer: The RIP routes are not present in the routing table because, by default, external routes are not leaked between the Level 1 database and the Level 2 database.

Part 3: Modifying the Default Flooding Scope

The default IS-IS flooding scope can be modified using routing policy. Because the R3-1, R4-1, R3-2, and R4-2 routers provide connectivity to both the Level 1 and Level 2 databases, export policy can be applied to the IS-IS protocol to direct routes into a specific database. In this lab, you write and apply IS-IS export policy for the R3-1, R4-1, R3-2, and R4-2 routers to change the default IS-IS flooding scope.

Step 3.1

IS-IS Level 1 internal routes are redistributed into the Level 2 database by default.

Use the **show route 10.0.12.0/23** command to display the Level 1 interface routes that have been learned by the R2-1 router due to this default rule of IS-IS operation.

```
lab@mxD:R2-1> show route 10.0.12.0/23

inet.0: 22 destinations, 22 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.12.0/24        *[IS-IS/18] 01:21:13, metric 20
                     > to 10.0.10.2 via ge-0/0/1.0
10.0.13.0/24        *[IS-IS/18] 01:21:12, metric 20
                     > to 10.0.11.2 via ge-0/0/2.0
```

Question: What routes have been leaked from Level 1 to Level 2?

_____

_____

Answer: The 10.0.12/24 and 10.0.13/24 routes have been leaked from Level 1 to Level 2.

Step 3.2

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxA> set cli logical-system R3-1
Logical system: R3-1

lab@mxA:R3-1>
```

Step 3.3

Enter configuration mode and navigate to the `[edit routing-options]` hierarchy. Configure an aggregate route of 10.0.12.0/23.

```
lab@mxA:R3-1> configure
Entering configuration mode

[edit]
lab@mxA:R3-1# edit routing-options

[edit routing-options]
lab@mxA:R3-1# set aggregate route 10.0.12.0/23

[edit routing-options]
lab@mxA:R3-1#
```

Step 3.4

Navigate to the `[edit policy-options policy-statement summarize-level-1]` hierarchy. Create a routing policy to summarize the two Level 1 /24 routes, identified in step 3.1, into a single /23 route, send them to Level 2 and, finally, suppress the more specific routes.

```
[edit routing-options]
lab@mxA:R3-1# top edit policy-options policy-statement summarize-level-1

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1# set term 1 from protocol aggregate

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1# set term 1 from route-filter 10.0.12.0/23 exact
```

```
[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1# set term 1 to level 2

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1# set term 1 then accept

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1# set term 2 from route-filter 10.0.12.0/23 longer

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1# set term 2 to level 2

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1# set term 2 then reject

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1# show
term 1 {
    from {
        protocol aggregate;
        route-filter 10.0.12.0/23 exact;
    }
    to level 2;
    then accept;
}
term 2 {
    from {
        route-filter 10.0.12.0/23 longer;
    }
    to level 2;
    then reject;
}

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1#
```

Step 3.5

Navigate to the `[edit protocols isis]` hierarchy and apply the policy named *`summarize-level-1`* as an export policy. Commit the changes and return to operational mode when completed.

```
[edit policy-options policy-statement summarize-level-1]
lab@mxA:R3-1# top edit protocols isis

[edit protocols isis]
lab@mxA:R3-1# set export summarize-level-1

[edit protocols isis]
lab@mxA:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R3-1>
```

Step 3.6

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, use the **`show route 10.0.12.0/23`** command to verify that the Level 1 internal routes have been summarized into the aggregate route and advertised into Level 2.

```
lab@mxD:R2-1> show route 10.0.12.0/23

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

10.0.12.0/23       *[IS-IS/165] 00:01:18, metric 20
                    > to 10.0.10.2 via ge-0/0/1.0
10.0.12.0/24       *[IS-IS/18] 00:01:18, metric 30
                    > to 10.0.11.2 via ge-0/0/2.0
10.0.13.0/24       *[IS-IS/18] 09:40:08, metric 20
                    > to 10.0.11.2 via ge-0/0/2.0
```

Question: Are the Level 1 routes being summarized in the Level 2 database?

_____

_____

Answer: Yes, the /23 summary route being advertised by the R3-1 router is an active IS-IS external route in the R2-1 router's routing table. The more specific /24 internal Level 1 routes have been suppressed by the R3-1 router.

Question: Why do you still see the /24 routes in the R2-1 router's routing table if the R3-1 is suppressing those routes?

_____

_____

Answer: The R4-1 router has not been configured to suppress the /24 routes. If you look at the next hops of the /24 routes, you will notice that all traffic to the /24 networks will be forwarded through the R4-1 router. You will configure the R4-1 router to suppress the /24 routes in the next few steps.

Step 3.7

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, change the perspective of the CLI to that of the R4-1 logical system by issuing the **`set cli logical-system R4-1`** command.

```
lab@mxA:R3-1> set cli logical-system R4-1
Logical system: R4-1

lab@mxA:R4-1>
```

Step 3.8

Enter configuration mode and navigate to the `[edit routing-options]` hierarchy. Configure an aggregate route of 10.0.12.0/23.

```
lab@mxA:R4-1> configure
Entering configuration mode

[edit]
lab@mxA:R4-1# edit routing-options

[edit routing-options]
lab@mxA:R4-1# set aggregate route 10.0.12.0/23

[edit routing-options]
lab@mxA:R4-1#
```

Step 3.9

Navigate to the `[edit policy-options policy-statement summarize-level-1]` hierarchy. Create a routing policy to summarize the two Level 1 /24 routes, identified in step 3.1, into a single /23 route, send them to Level 2 and, finally, suppress the more specific routes.

```
[edit routing-options]
lab@mxA:R4-1# top edit policy-options policy-statement summarize-level-1

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1# set term 1 from protocol aggregate

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1# set term 1 from route-filter 10.0.12.0/23 exact

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1# set term 1 to level 2

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1# set term 1 then accept

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1# set term 2 from route-filter 10.0.12.0/23 longer

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1# set term 2 to level 2

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1# set term 2 then reject

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1# show
term 1 {
    from {
        protocol aggregate;
        route-filter 10.0.12.0/23 exact;
    }
    to level 2;
    then accept;
}
```

```
term 2 {
    from {
        route-filter 10.0.12.0/23 longer;
    }
    to level 2;
    then reject;
}

[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1#
```

Step 3.10

Navigate to the `[edit protocols isis]` hierarchy and apply the policy named *`summarize-level-1`* as an export policy. Commit the changes when completed.

```
[edit policy-options policy-statement summarize-level-1]
lab@mxA:R4-1# top edit protocols isis

[edit protocols isis]
lab@mxA:R4-1# set export summarize-level-1

[edit protocols isis]
lab@mxA:R4-1# commit
commit complete

[edit protocols isis]
lab@mxA:R4-1#
```

Step 3.11

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, use the **`show route 10.0.12.0/23`** command to verify that the Level 1 internal routes have been summarized into the aggregate route and advertised into Level 2.

```
lab@mxD:R2-1> show route 10.0.12.0/23

inet.0: 21 destinations, 21 routes (21 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.12.0/23       *[IS-IS/165] 00:00:24, metric 20
                    > to 10.0.10.2 via ge-0/0/1.0
                      to 10.0.11.2 via ge-0/0/2.0
```

Question: Are the Level 1 routes being summarized in the Level 2 database?

_____

Answer: Yes, the /23 summary route being advertised by the R3-1 and R4-1 router is an active IS-IS external route in the R2-1 router's routing table as seen in the output and the more specific /24 internal Level 1 routes have been suppressed by both the R3-1 and R4-1 router.

Step 3.12

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, navigate to the `[edit policy-options policy-statement rip-to-level-2]` hierarchy. Configure the *`rip-to-level-2`* policy to accept all routes more specific than 20.20.0.0/22 into the Level 2 database.

```
[edit protocols isis]
lab@mxA:R4-1# top edit policy-options policy-statement rip-to-level-2

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R4-1# set term 1 from route-filter 20.20.0.0/22 longer

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R4-1# set term 1 to level 2

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R4-1# set term 1 then accept

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R4-1# show
term 1 {
    from {
        route-filter 20.20.0.0/22 longer;
    }
    to level 2;
    then accept;
}

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R4-1#
```

Step 3.13

Navigate to the `[edit protocols isis]` hierarchy and apply the *`rip-to-level-2`* policy as an export policy. Commit your configuration and exit to operational mode when finished.

```
[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R4-1# top edit protocols isis

[edit protocols isis]
lab@mxA:R4-1# set export rip-to-level-2

[edit protocols isis]
lab@mxA:R4-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R4-1>
```

Step 3.14

Change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system`** ***`R3-1`*** command.

```
lab@mxA:R4-1> set cli logical-system R3-1
Logical system: R3-1

lab@mxA:R3-1>
```

Step 3.15

Enter configuration mode and navigate to the `[edit policy-options policy-statement rip-to-level-2]` hierarchy. Configure the *`rip-to-level-2`* policy to accept all routes more specific than 20.20.0.0/22 into the Level 2 database.

```
lab@mxA:R3-1> configure
Entering configuration mode

[edit]
lab@mxA:R3-1# edit policy-options policy-statement rip-to-level-2

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R3-1# set term 1 from route-filter 20.20.0.0/22 longer

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R3-1# set term 1 to level 2

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R3-1# set term 1 then accept

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R3-1# show
term 1 {
    from {
        route-filter 20.20.0.0/22 longer;
    }
    to level 2;
    then accept;
}

[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R3-1#
```

Step 3.16

Navigate to the `[edit protocols isis]` hierarchy and apply the *`rip-to-level-2`* policy as an export policy. Commit your configuration and exit to operational mode when finished.

```
[edit policy-options policy-statement rip-to-level-2]
lab@mxA:R3-1# top edit protocols isis

[edit protocols isis]
lab@mxA:R3-1# set export rip-to-level-2

[edit protocols isis]
lab@mxA:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R3-1>
```

Step 3.17

Return to the SSH session that you have open to the mxD router.

From the SSH session with the mxD router, use the **`show route 20.20/22`** command to verify that the RIP routes have been injected into the Level 2 database and SPF installed them in the R2-1 router's routing table.

```
lab@mxD:R2-1> show route 20.20/22

inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/24       *[IS-IS/165] 00:04:38, metric 22
                    > to 10.0.10.2 via ge-0/0/1.0
                      to 10.0.11.2 via ge-0/0/2.0
20.20.1.0/24       *[IS-IS/165] 00:04:38, metric 22
                    > to 10.0.10.2 via ge-0/0/1.0
                      to 10.0.11.2 via ge-0/0/2.0
20.20.2.0/24       *[IS-IS/165] 00:04:38, metric 22
                    > to 10.0.10.2 via ge-0/0/1.0
                      to 10.0.11.2 via ge-0/0/2.0
20.20.3.0/24       *[IS-IS/165] 00:04:38, metric 22
                    > to 10.0.10.2 via ge-0/0/1.0
                      to 10.0.11.2 via ge-0/0/2.0
```

Question: Are the RIP routes active routes in the routing table? What protocol preference is assigned to the routes? Why?

_____

_____

Answer: Yes. The RIP routes are active IS-IS routes in the routing table. The routes have a protocol preference of 165, the preference assigned to external Level 2 routes.

Step 3.18

Return to the SSH session that you have open to the mxE router.

From the SSH session with the mxE router, change the perspective of the CLI to that of the R5-2 logical system by issuing the **`set cli logical-system R5-2`** command.

```
lab@mxE:R5-1> set cli logical-system R5-2
Logical system: R5-2

lab@mxE:R5-2>
```

Step 3.19

Use the **`show route 20.20/22`** command to verify that the RIP routes have been injected into the Level 1 database and SPF installed them in the R5-2 router's routing table.

```
lab@mxE:R5-2> show route 20.20/22

lab@mxE:R5-2>
```

Question: Are the RIP routes active routes in the routing table? Why?

_____

_____

Answer: No, the RIP routes are not active IS-IS routes in the routing table. By default, the L1/L2 routers for area 49.0002 (R3-2 and R4-2) do not leak Level 2 learned routes into the Level 1 database.

Step 3.20

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, change the perspective of the CLI to that of the R3-2 logical system by issuing the **`set cli logical-system R3-2`** command.

```
lab@mxC> set cli logical-system R3-2
Logical system: R3-2

lab@mxC:R3-2>
```

Step 3.21

Enter configuration mode and navigate to the `[edit policy-options policy-statement level2-to-level1]` hierarchy. Configure the *`level2-to-level1`* policy to accept all IS-IS Level 2 routes into Level 1.

```
lab@mxC:R3-2> configure
Entering configuration mode

[edit]
lab@mxC:R3-2# edit policy-options policy-statement level2-to-level1

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R3-2# set term 1 from protocol isis

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R3-2# set term 1 from level 2

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R3-2# set term 1 to level 1

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R3-2# set term 1 then accept

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R3-2# show
term 1 {
    from {
        protocol isis;
        level 2;
    }
    to level 1;
    then accept;
```

```
}

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R3-2#
```

Step 3.22

Navigate to the `[edit protocols isis]` hierarchy and apply the *`level2-to-level1`* policy as an export policy. Commit the configuration and return to operational mode.

```
[edit policy-options policy-statement level2-to-level1]
lab@mxC:R3-2# top edit protocols isis

[edit protocols isis]
lab@mxC:R3-2# set export level2-to-level1

[edit protocols isis]
lab@mxC:R3-2# commit and-quit
commit complete
Exiting configuration mode

lab@mxC:R3-2>
```

Step 3.23

Return to the SSH session that you have open to the mxE router.

From the SSH session with the mxE router, use the **`show route 20.20/22`** command to verify that the RIP routes have been injected into the Level 1 database and SPF installed them in the R5-2 router's routing table.

```
lab@mxE:R5-2> show route 20.20/22

inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/24       *[IS-IS/165] 00:01:29, metric 62
                    > to 10.0.16.1 via ge-0/0/2.0
20.20.1.0/24       *[IS-IS/165] 00:01:29, metric 62
                    > to 10.0.16.1 via ge-0/0/2.0
20.20.2.0/24       *[IS-IS/165] 00:01:29, metric 62
                    > to 10.0.16.1 via ge-0/0/2.0
20.20.3.0/24       *[IS-IS/165] 00:01:29, metric 62
                    > to 10.0.16.1 via ge-0/0/2.0
```

Question: Are the RIP routes active routes in the routing table?

_____

_____

Answer: Yes, the RIP routes are active IS-IS routes in the routing table.

Question: What is the next hop for each of the /24 routes?

_____

_____

Answer: The next hop for each of the routes is ge-0/0/2.0 which is through the R3-2 router.

Step 3.24

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, change the perspective of the CLI to that of the R4-2 logical system by issuing the **`set cli logical-system R4-2`** command.

```
lab@mxC:R3-2> set cli logical-system R4-2
Logical system: R4-2

lab@mxC:R4-2>
```

Step 3.25

Enter configuration mode and navigate to the `[edit policy-options policy-statement level2-to-level1]` hierarchy. Configure the *`level2-to-level1`* policy to accept all IS-IS Level 2 routes into Level 1.

```
lab@mxC:R4-2> configure
Entering configuration mode

[edit]
lab@mxC:R4-2# edit policy-options policy-statement level2-to-level1

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R4-2# set term 1 from protocol isis

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R4-2# set term 1 from level 2

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R4-2# set term 1 to level 1

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R4-2# set term 1 then accept

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R4-2# show
term 1 {
    from {
        protocol isis;
        level 2;
    }
    to level 1;
    then accept;
}

[edit policy-options policy-statement level2-to-level1]
lab@mxC:R4-2#
```

Step 3.26

Navigate to the `[edit protocols isis]` hierarchy and apply the *`level2-to-level1`* policy as an export policy. Commit the configuration and return to operational mode.

```
[edit policy-options policy-statement level2-to-level1]
lab@mxC:R4-2# top edit protocols isis

[edit protocols isis]
lab@mxC:R4-2# set export level2-to-level1

[edit protocols isis]
lab@mxC:R4-2# commit and-quit
commit complete
Exiting configuration mode

lab@mxC:R4-2>
```

Step 3.27

Return to the SSH session that you have open to the mxE router.

From the SSH session with the mxE router, verify that the RIP routes have been leaked into the IS-IS Level 1 database and installed on the R5-2 router. Use the **`show route 20.20/22`** command to display the routes learned by the R5-2 router.

```
lab@mxE:R5-2> show route 20.20/22

inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.0.0/24       *[IS-IS/165] 00:02:57, metric 62
                    > to 10.0.16.1 via ge-0/0/2.0
                      to 10.0.17.1 via ge-0/0/3.0
20.20.1.0/24       *[IS-IS/165] 00:02:57, metric 62
                    > to 10.0.16.1 via ge-0/0/2.0
                      to 10.0.17.1 via ge-0/0/3.0
20.20.2.0/24       *[IS-IS/165] 00:02:57, metric 62
                    > to 10.0.16.1 via ge-0/0/2.0
                      to 10.0.17.1 via ge-0/0/3.0
20.20.3.0/24       *[IS-IS/165] 00:02:57, metric 62
                    > to 10.0.16.1 via ge-0/0/2.0
                      to 10.0.17.1 via ge-0/0/3.0
```

Question: Have the RIP routes been leaked into the Level 1 database and two next hops installed in the R5-2 router's routing table? What is the protocol preference of the leaked routes in the routing table?

_____

_____

Answer: Yes. The RIP routes have been leaked into the Level 1 database and SPF has installed the routes in the R5-2 router's routing table. The protocol preference of the leaked routes is 165, the preference for external routes.

Step 3.28

Use the **`show isis database level 1`** command to locate the LSP-ID for the R3-2 router in the Level 1 IS-IS database. Output may vary.

```
lab@mxE:R5-2> show isis database level 1
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
mxC-R3-2.00-00             0x334   0x6ab3     1171 L1 L2 Attached
mxC-R4-2.00-00             0x331   0x918e      843 L1 L2 Attached
mxE-R5-2.00-00             0x332   0xfc49      526 L1 L2
mxE-R5-2.02-00              0x33    0x545      742 L1 L2
mxE-R5-2.03-00              0x34   0xb594     1044 L1 L2
  5 LSPs
```

Question: What is the LSP-ID of the primary LSP for the R3-2 router?

_____

_____

Answer: Based on the previous output, the primary LSP-ID of the R3-2 router is `mxC-R3-2.00-00`.

Step 3.29

Use the LSP-ID located from the previous step to display the TLVs inserted by the *`level2-to-level1`* export policy. Enter the **`show isis database level 1`** ***`mxC-R3-2.00-00`*** **`extensive | find tlv`** command to display the TLVs in the R3-2 router's LSP.

```
lab@mxE:R5-2> show isis database level 1 mxC-R3-2.00-00 extensive | find tlv
TLVs:
    Area address: 49.0002 (3)
    LSP Buffer Size: 1492
    Speaks: IP
    Speaks: IPV6
    IP router id: 172.16.2.2
    IP address: 172.16.2.2
    Hostname: mxC-R3-2
    IS neighbor: mxE-R5-2.02, Internal, Metric: default 10
    Extended IS Reachability TLV, Type: 22, Length: 27
    IS extended neighbor: mxE-R5-2.02, Metric: default 10 SubTLV len: 16
      IP address: 10.0.16.1
      Local interface index: 342, Remote interface index: 0
    IP prefix: 10.0.16.0/24, Internal, Metric: default 10, Up
    IP prefix: 172.16.2.2/32, Internal, Metric: default 0, Up
    IP prefix: 10.0.10.0/24, Internal, Metric: default 40, Down
    IP prefix: 10.0.11.0/24, Internal, Metric: default 40, Down
    IP prefix: 10.0.15.0/24, Internal, Metric: default 20, Down
    IP prefix: 172.16.1.1/32, Internal, Metric: default 30, Down
    IP prefix: 172.16.1.2/32, Internal, Metric: default 40, Down
    IP prefix: 172.16.1.3/32, Internal, Metric: default 40, Down
    IP prefix: 172.16.1.4/32, Internal, Metric: default 50, Down
    IP prefix: 172.16.2.1/32, Internal, Metric: default 10, Down
    IP prefix: 172.22.121.0/24, Internal, Metric: default 30, Down
```

```
    IP prefix: 172.22.122.0/24, Internal, Metric: default 20, Down
    IP prefix: 172.31.100.1/32, Internal, Metric: default 20, Down
    IP extended prefix: 10.0.16.0/24 metric 10 up
    IP extended prefix: 172.16.2.2/32 metric 0 up
    IP extended prefix: 10.0.10.0/24 metric 40 down
    IP extended prefix: 10.0.11.0/24 metric 40 down
    IP extended prefix: 10.0.15.0/24 metric 20 down
    IP extended prefix: 172.16.1.1/32 metric 30 down
    IP extended prefix: 172.16.1.2/32 metric 40 down
    IP extended prefix: 172.16.1.3/32 metric 40 down
    IP extended prefix: 172.16.1.4/32 metric 50 down
    IP extended prefix: 172.16.2.1/32 metric 10 down
    IP extended prefix: 172.22.121.0/24 metric 30 down
    IP extended prefix: 172.22.122.0/24 metric 20 down
    IP extended prefix: 172.31.100.1/32 metric 20 down
    Router Capability:  Router ID 172.16.2.2, Flags: 0x00
    IP external prefix: 10.0.12.0/23, Internal, Metric: default 50, Down
    IP external prefix: 20.20.0.0/24, Internal, Metric: default 52, Down
    IP external prefix: 20.20.1.0/24, Internal, Metric: default 52, Down
    IP external prefix: 20.20.2.0/24, Internal, Metric: default 52, Down
    IP external prefix: 20.20.3.0/24, Internal, Metric: default 52, Down
    IP extended prefix: 10.0.12.0/23 metric 50 down
    IP extended prefix: 20.20.0.0/24 metric 52 down
    IP extended prefix: 20.20.1.0/24 metric 52 down
    IP extended prefix: 20.20.2.0/24 metric 52 down
    IP extended prefix: 20.20.3.0/24 metric 52 down
No queued transmissions
```

Question: The Type 128 TLVs leaked into the Level 1 database have the “down” bit set. What is the function of this bit?

_____

_____

Answer: Because the default IS-IS flooding scope is to leak Level 1 internal Type 128 TLVs into the Level 2 database, the down bit is set in order to prevent routing loops. IS-IS Level 1 Type 128 TLVs with the down bit set are never leaked back into the Level 2 database as a loop detection mechanism.

Step 3.30

Log out of your assigned device using the **exit** command.

```
lab@mxE:R5-2> exit

mxE (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

Network Management Diagram

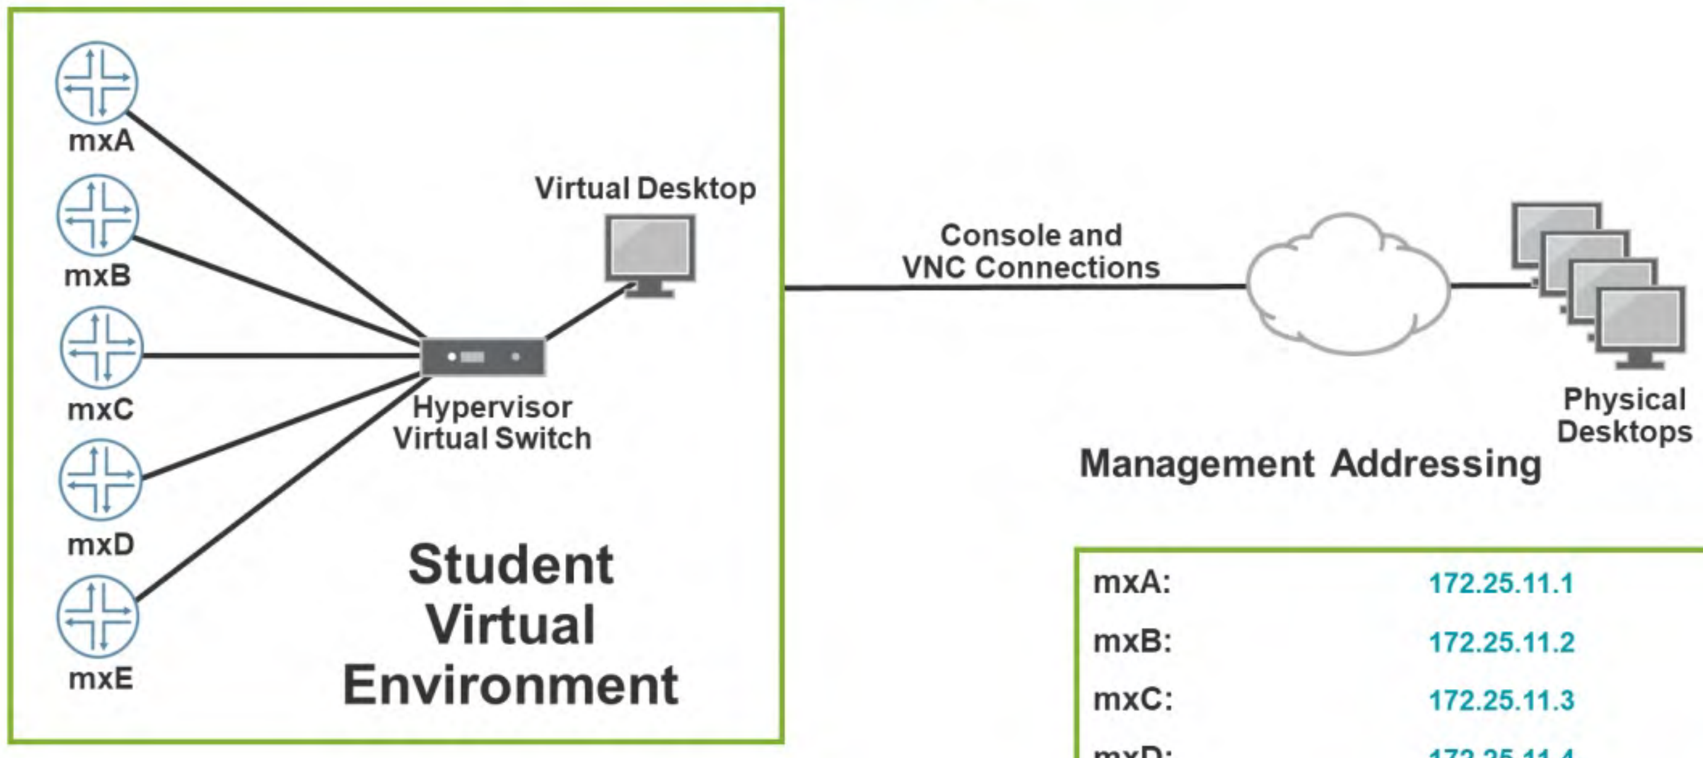

Management Addressing

mxA:	172.25.11.1
mxB:	172.25.11.2
mxC:	172.25.11.3
mxD:	172.25.11.4
mxE:	172.25.11.5
Virtual Desktop:	172.25.11.254

Note: Your instructor will provide address and access information.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab Network Diagram: Configuring a Multilevel IS-IS Network

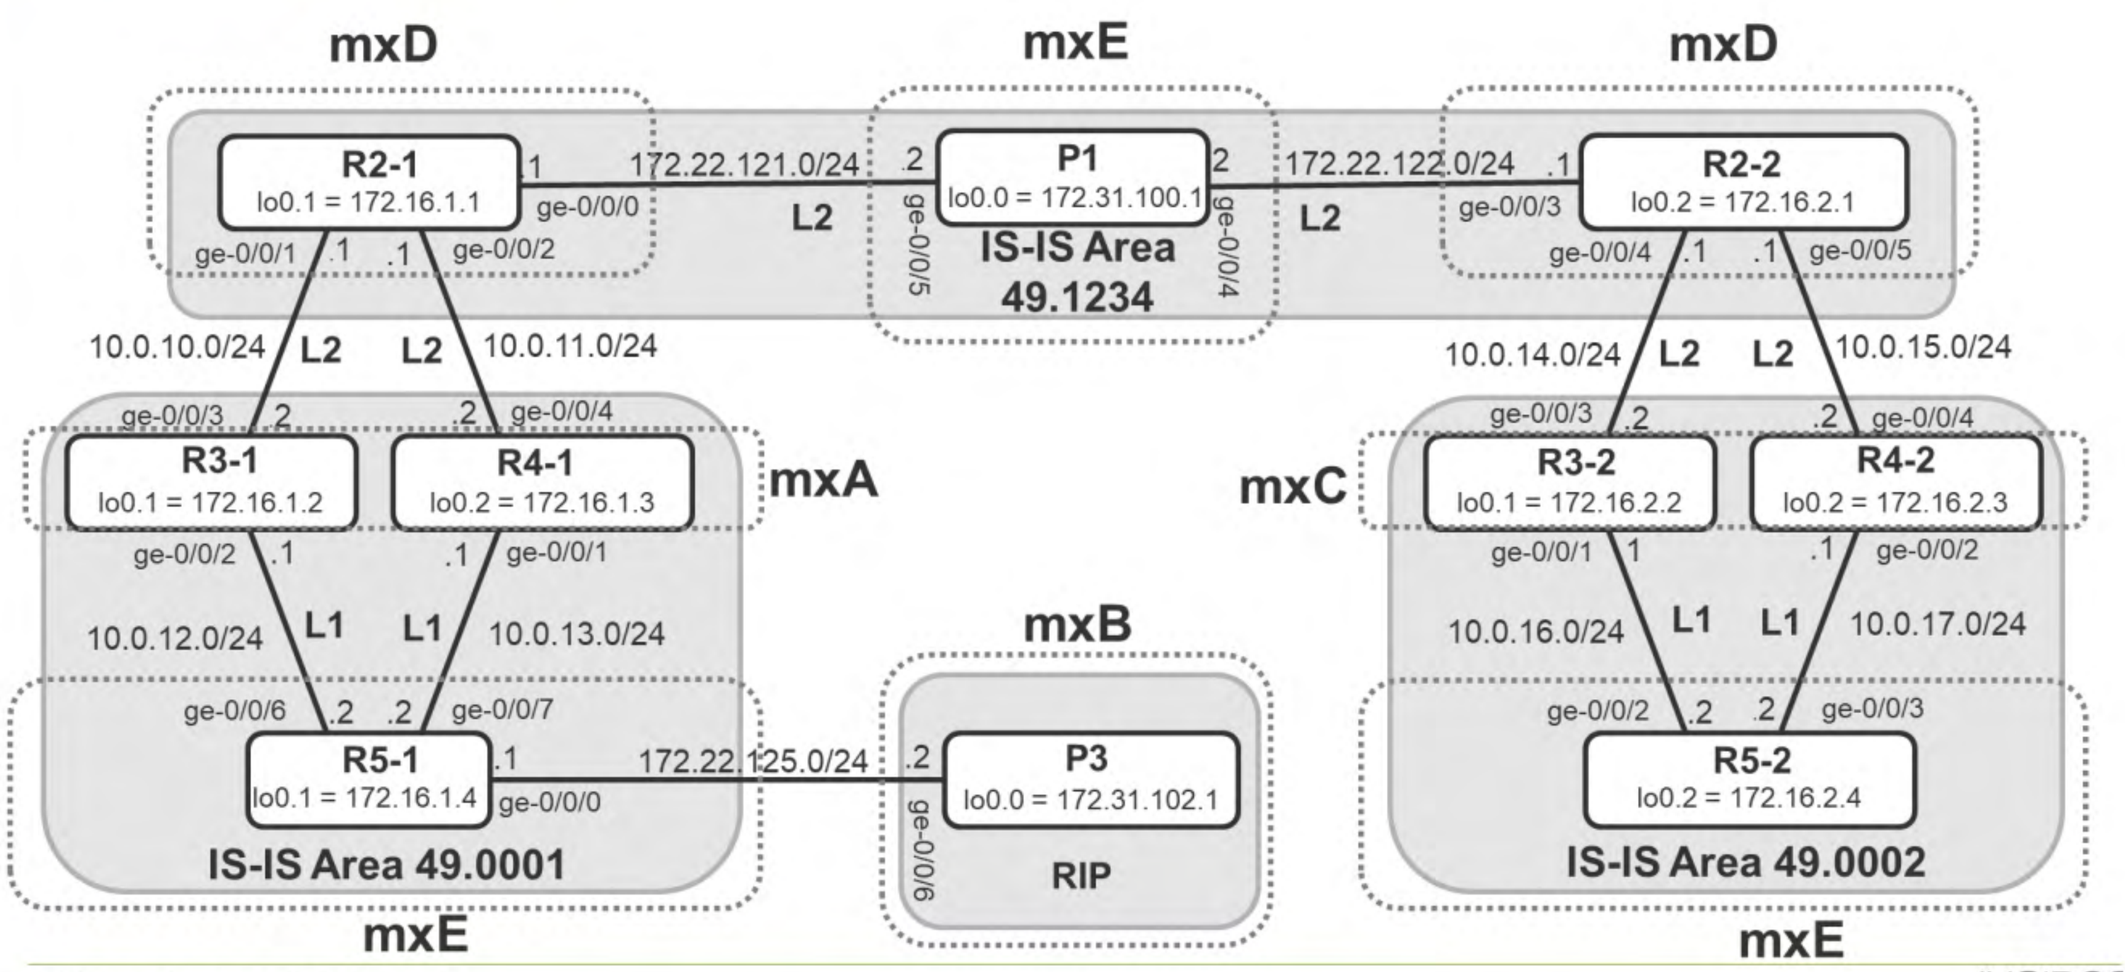

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab

Troubleshooting IS-IS

Overview

In this lab, you use the lab diagram titled "Lab Network Diagram: Troubleshooting IS-IS" to troubleshoot an IS-IS scenario. The goal is to get all adjacencies up and make sure IS-IS is advertising each routers directly attached networks, so that you will be able to send traffic through the backbone from the endpoints.

By completing this lab, you will perform the following tasks:

- Troubleshoot various IS-IS peering issues.
- Make appropriate configuration changes to rectify peering problems.
- Verify and validate routes being shared and ensure that all routes are present.

Part 1: Load Starting Configuration

In this lab part, you load the initial configuration of the routers.

Note

The instructor will tell you the nature of your access and will provide you with the necessary details to access your assigned device.

Note

The lab topology requires you to display information in the different logical systems of mxB and mxD. You will be changing the perspective of the CLI by issuing the **`set cli logical-system`** ***`ls-name`*** command. Refer to the lab diagram for the correct logical system name.

By changing the perspective of the CLI, you will be able to perform network commands such as **`ping`** or **`traceroute`** from the perspective of the pertinent logical system.

Step 1.1

Ensure that you know to which student device you have been assigned. Check with your instructor if you are not certain. Consult the management network diagram to determine the management address of the student devices.

Access the CLI on your mxD device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab8-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.4
Password: lab123
Last login: Fri Dec  6 12:59:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab8-start.config
load complete

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

lab@mxD>
```

Step 1.2

Access the CLI on your mxC device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab8-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.3
Password: lab123
Last login: Fri Dec  6 12:59:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# load override ajspr/lab8-start.config
load complete

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

lab@mxC>
```

Step 1.3

Access the CLI on your mxB device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab8-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.2
Password: lab123
Last login: Fri Dec  6 12:59:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxB> configure
Entering configuration mode

[edit]
lab@mxB# load override ajspr/lab8-start.config
load complete

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

lab@mxB>
```

Step 1.4

Access the CLI on your mxA device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab8-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.1
Password: lab123
Last login: Fri Dec  6 12:59:36 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# load override ajspr/lab8-start.config
load complete

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

lab@mxA>
```

Part 2: Troubleshooting IS-IS

In this lab part, you check IS-IS adjacencies to troubleshoot the link-state database.

Step 2.1

Issue the **show isis interface** command to view the interfaces that have been configured for IS-IS.

```
lab@mxA> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/0.0            2   0x1 Disabled          mxA.00                 10/10
ge-0/0/3.0            2   0x1 Passive           Passive                10/10
lo0.0                 2   0x1 Passive           Passive                 0/0
```

Question: Is IS-IS running on the correct interfaces?

_____

Answer: Yes, IS-IS is running on the correct interfaces. The ge-0/0/3.0 interface should be running in the passive mode of operation. System ID names may vary due to IS-IS DIS elections.

Step 2.2

Issue the **show isis adjacency** command to view the status of the IS-IS neighbor relationships.

```
lab@mxA> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/0.0            mxB-R2         2  Up                    6  0:50:56:a9:5f:d4
```

Question: Are the expected IS-IS neighbors up?

_____

Answer: Yes, the mxA router should have an IS-IS adjacency with the R2 router. System ID names may vary due to IS-IS DIS elections.

Step 2.3

Return to the SSH session that you have open to the mxB router.

From the SSH session with the mxB router, change the perspective of the CLI to that of the R2 logical system by issuing the **`set cli logical-system R2`** command.

```
lab@mxB> set cli logical-system R2
Logical system: R2

lab@mxB:R2>
```

Step 2.4

Issue the **`show isis interface`** command to view the interfaces that have been configured for IS-IS.

```
lab@mxB:R2> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/0.0            2   0x2 Disabled          mxB-R2.02              10/10
ge-0/0/2.0            2   0x1 Disabled          mxB-R3.02              10/10
lo0.2                 2   0x1 Passive           Passive                 0/0
lt-0/0/10.0           2   0x1 Disabled          mxB-R4.02              10/10
```

Question: Is IS-IS running on the correct interfaces?

_____

Answer: Yes, IS-IS is running on the correct interfaces. System ID names may vary due to IS-IS DIS elections.

Step 2.5

Issue the **`show isis adjacency`** command to view the status of the IS-IS neighbor relationships.

```
lab@mxB:R2> show isis adjacency
Interface             System          L State        Hold (secs) SNPA
ge-0/0/0.0            mxC             2  Up                   24  0:50:56:a9:60:de
ge-0/0/2.0            mxB-R3          2  Up                    8  0:50:56:a9:56:d8
lt-0/0/10.0           mxB-R4          2  Up                    7  0:50:56:a9:62:51
```

Question: Are the expected IS-IS neighbors up?

_____

Answer: Yes, the R2 router should have an IS-IS adjacency with all three directly attached neighbors. If all adjacencies do not appear, you can try the **clear isis adjacency all** command. System ID names may vary due to IS-IS DIS elections.

Step 2.6

Change the perspective of the CLI to that of the R3 logical system by issuing the **`set cli logical-system R3`** command.

```
lab@mxB:R2> set cli logical-system R3
Logical system: R3

lab@mxB:R3>
```

Step 2.7

Issue the **`show isis interface`** command to view the interfaces that have been configured for IS-IS.

```
lab@mxB:R3> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/3.0            2   0x2 Disabled          mxB-R3.02              10/10
ge-0/0/4.0            2   0x3 Disabled          mxB-R3.03              10/10
lo0.3                 2   0x1 Passive           Passive                 0/0
```

Question: Is IS-IS running on the correct interfaces?

_____

Answer: Yes, IS-IS is running on the correct interfaces. System ID names may vary due to IS-IS DIS elections.

Step 2.8

Issue the **`show isis adjacency`** command to view the status of the IS-IS neighbor relationships.

```
lab@mxB:R3> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/3.0            mxB-R2         2  Up                   22  0:50:56:a9:a:3a
ge-0/0/4.0            mxB-R4         2  Up                   20  0:50:56:a9:1e:eb
```

Question: Are the expected IS-IS neighbors up?

_____

_____

Answer: Yes, the R3 router should have an IS-IS adjacency with both attached neighbors. System ID names may vary due to IS-IS DIS elections.

Step 2.9

Change the perspective of the CLI to that of the R4 logical system by issuing the **`set cli logical-system R4`** command.

```
lab@mxB:R3> set cli logical-system R4
Logical system: R4

lab@mxB:R4>
```

Step 2.10

Issue the **`show isis interface`** command to view the interfaces that have been configured for IS-IS.

```
lab@mxB:R4> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/1.0            2   0x1 Disabled          mxC.02                 10/10
ge-0/0/5.0            2   0x1 Disabled          mxB-R3.03              10/10
lo0.4                 2   0x1 Passive           Passive                 0/0
lt-0/0/10.1           2   0x2 Disabled          mxB-R4.02              10/10
```

Question: Is IS-IS running on the correct interfaces?

_____

Answer: Yes, IS-IS is running on the correct interfaces. System ID names may vary due to IS-IS DIS elections.

Step 2.11

Issue the **`show isis adjacency`** command to view the status of the IS-IS neighbor relationships.

```
lab@mxB:R4> show isis adjacency
Interface             System          L State        Hold (secs) SNPA
ge-0/0/1.0            mxA             2  Up                   21  0:50:56:a9:51:76
ge-0/0/5.0            mxB-R3          2  Up                    8  0:50:56:a9:6c:92
lt-0/0/10.1           mxB-R2          2  Up                   19  0:50:56:a9:5f:d4
```

Question: Are the expected IS-IS neighbors up?

_____

_____

Answer: Yes, the R4 router should have an IS-IS adjacency with all three directly attached neighbors. System ID names may vary due to IS-IS DIS elections.

Step 2.12

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, issue the **`show isis interface`** command to view the interfaces that have been configured for IS-IS.

```
lab@mxC> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR        Level 2 DR        L1/L2 Metric
ge-0/0/0.0            2   0x2 Disabled          mxC.02                 10/10
ge-0/0/3.0            2   0x1 Disabled          mxC.00                 10/10
lo0.0                 2   0x1 Passive           Passive                 0/0
```

Question: Is IS-IS running on the correct interfaces?

_____

_____

Answer: Yes, IS-IS is running on the correct interfaces. System ID names may vary due to IS-IS DIS elections.

Step 2.13

Issue the **`show isis adjacency`** command to view the status of the IS-IS neighbor relationships.

```
lab@mxC> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/0.0            mxB-R4         2  Up                    6  0:50:56:a9:70:61
```

Question: Are the expected IS-IS neighbors up?

_____

_____

Answer: Yes, the mxC router should have an IS-IS adjacency with the R4 router.

Step 2.14

Return to the SSH session that you have open to the mxB router.

From the SSH session with the mxB router, change the perspective of the CLI to that of the R3 logical system by issuing the **`set cli logical-system R3`** command.

```
lab@mxB:R4> set cli logical-system R3
Logical system: R3

lab@mxB:R3>
```

Step 2.15

Check that you are receiving routes for all loopbacks of the five routers in the lab network.

```
lab@mxB:R3> show route 192.168/16

inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.71.1/32    *[IS-IS/18] 00:02:00, metric 20
                    > to 172.22.131.9 via ge-0/0/3.0
192.168.71.2/32    *[IS-IS/18] 1d 14:21:58, metric 10
                    > to 172.22.131.9 via ge-0/0/3.0
192.168.71.3/32    *[Direct/0] 1d 14:22:09
                    > via lo0.3
192.168.71.4/32    *[IS-IS/18] 1d 14:17:31, metric 10
                    > to 172.22.131.14 via ge-0/0/4.0
```

Question: Do you see all loopback interfaces?

_____

_____

Answer: No, the mxC router's loopback interface is missing.

Note

In some situations, you might also see that the mxA router's loopback address is missing from the routing table. This is related to how and when mxA's loopback is placed in the database. If you are missing mxA's loopback your outputs will vary slightly through the next few steps. Do not be alarmed this is normal and will be resolved in future steps when you resolve mxC's missing address.

Step 2.16

Examine the link-state database using the **show isis database** command.

```
lab@mxB:R3> show isis database
IS-IS level 1 link-state database:
  0 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
mxA.00-00                     0x6560   0x2259     1194 L1 L2
mxB-R2.00-00                    0xba   0x5543      848 L1 L2
mxB-R2.02-00                    0xb3   0x3b61      848 L1 L2
mxB-R3.00-00                    0xb4   0x6fe0      803 L1 L2
mxB-R3.02-00                    0xb1   0xadeb     1092 L1 L2
mxB-R3.03-00                    0xb0   0x3560      872 L1 L2
mxB-R4.00-00                    0xb4   0x6ad0      800 L1 L2
mxB-R4.02-00                    0xb1   0xd5c0      452 L1 L2
mxB-R4.03-00                    0xb1   0x880f      801 L1 L2
  9 LSPs
```

Question: Do you see a link-state PDU for all the routers? Do you also see any difference in the magnitude of the sequence numbers between the PDUs generate by different routers?

_____

_____

Answer: One link-state PDU seems to be missing (which one can vary). Moreover, one of the routers seems to have a link-state PDU with much higher sequence number from the rest. Take note of the router which is missing from the database, and the one whose PDU has a higher-than-expected sequence number: since we are troubleshooting an anomaly, anything unusual can help.

Step 2.17

Connect to the missing router, and check the content of the link-state database. The example will assume that it is the mxC router that is missing.

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, issue the **`show isis database`** command.

```
lab@mxC> show isis database
IS-IS level 1 link-state database:
  0 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
mxC.00-00                     0x6589    0x556     1192 L1 L2
mxB-R2.00-00                    0xba   0x5543      642 L1 L2
mxB-R2.02-00                    0xb3   0x3b61      641 L1 L2
mxB-R3.00-00                    0xb4   0x6fe0      595 L1 L2
mxB-R3.02-00                    0xb1   0xadeb      884 L1 L2
mxB-R3.03-00                    0xb0   0x3560      664 L1 L2
mxB-R4.00-00                    0xb4   0x6ad0      596 L1 L2
mxB-R4.02-00                    0xb2   0xd3c1     1053 L1 L2
mxB-R4.03-00                    0xb1   0x880f      596 L1 L2
  9 LSPs
```

Question: Do you notice any missing link-state PDU, or any link-state PDU with unusually high sequence number?

_____

_____

Answer: There is a missing link-state PDU. It was the router whose PDU had unusually high sequence number. On this router, their own link-state PDU shows a high (and increasing) sequence number.

Step 2.18

Enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Enable traceoptions using a file called ***isis*** and the **`flag error detail`** options. Commit your configuration when finished.

```
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# edit protocols isis

[edit protocols isis]
lab@mxC# set traceoptions file isis

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

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

[edit protocols isis]
lab@mxC#
```

Step 2.19

Wait a minute, then check the logs generated because of the traceoptions settings by issuing the **`run show log isis | last 4`** command.

```
[edit protocols isis]
lab@mxC# run show log isis | last 4
Dec 26 15:48:50.598673 Notification not sent due to 5-second throttle as per
   rfc4444: isisSequenceNumberSkip
Dec 26 15:48:55.609199 Notification not sent due to 5-second throttle as per
   rfc4444: isisSequenceNumberSkip
Dec 26 15:48:55.661548 Notification not sent due to 5-second throttle as per
   rfc4444: isisSequenceNumberSkip
Dec 26 15:48:55.712103 ERROR: Possible sysid collision

[edit protocols isis]
lab@mxC#
```

Question: Do you see anything that could indicate an error?

_____

Answer: There are messages about throttling (delaying) purging your own LSP; you may also see a message that is more explicit: `ERROR: Possible sysid collision`.

Step 2.20

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, enter configuration mode and navigate to the `[edit protocols isis]` hierarchy. Enable traceoptions using a file called ***isis*** and the **flag error detail** options. Commit your configuration when finished.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit protocols isis

[edit protocols isis]
lab@mxA# set traceoptions file isis

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

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

[edit protocols isis]
lab@mxA#
```

Step 2.21

Wait a minute, then check the logs generated because of the traceoptions settings by issuing the **run show log isis | last 4** command.

```
[edit protocols isis]
lab@mxA# run show log isis | last 4
May 17 17:36:01.965884 Notification not sent due to 5-second throttle as per
   rfc4444: isisSequenceNumberSkip
May 17 17:36:02.225742 Notification not sent due to 5-second throttle as per
   rfc4444: isisSequenceNumberSkip
May 17 17:36:07.247911 Notification not sent due to 5-second throttle as per
   rfc4444: isisSequenceNumberSkip
May 17 17:36:07.906711 Notification not sent due to 5-second throttle as per
   rfc4444: isisSequenceNumberSkip
```

Question: Do you see anything that could indicate an error?

_____

_____

Answer: There are messages about throttling (delaying) purging your own LSP; you may also see a message that is more explicit: `ERROR: Possible sysid collision`.

Step 2.22

Return to the SSH session that you have open to the mxC router.

From the SSH session with the mxC router, navigate to the `[edit interfaces lo0]` hierarchy. Issue the **show** command and review the ISO address on assigned to the loopback interface.

```
[edit protocols isis]
lab@mxC# top edit interfaces lo0

[edit interfaces lo0]
lab@mxC# show
unit 0 {
    family inet {
        address 192.168.71.5/32;
    }
    family iso {
        address 99.0001.0000.0000.7801.00;
    }
    family inet6 {
        address 2001:db8:101::5/128 {
            primary;
        }
        address ::ffff:192.168.71.5/128;
    }
}

[edit interfaces lo0]
lab@mxC#
```

Question: Has the correct ISO address been configured on the loopback interface?

_____

_____

Answer: No, mxC has been configured with an incorrect address. Based on the lab diagram, mxC should be configured with an ISO address of 49.0001.0000.0000.7105.00.

Step 2.23

Fix the incorrectly configured address. Commit and exit to operational mode when you are finished.

```
[edit interfaces lo0]
lab@mxC# rename unit 0 family iso address 99.0001.0000.0000.7801.00 to address
   49.0001.0000.0000.7105.00

[edit interfaces lo0]
lab@mxC# show
unit 0 {
    family inet {
        address 192.168.71.5/32;
    }
    family iso {
        address 49.0001.0000.0000.7105.00;
    }
    family inet6 {
        address 2001:db8:101::5/128 {
            primary;
        }
```

```
        address ::ffff:192.168.71.5/128;
    }
}

[edit interfaces lo0]
lab@mxC# commit and-quit
commit complete
Exiting configuration mode

lab@mxC>
```

Step 2.24

Use the **show isis database** to verify that the link-state database is now complete. You might need to wait a minute to allow the database to synchronize.

```
lab@mxC> show isis database
IS-IS level 1 link-state database:
  0 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime Attributes
mxC.00-00                        0x2   0xd5a9     1177 L1 L2
mxC.02-00                        0x2   0xe56e     1177 L1 L2
mxA.00-00                      0x16d   0xebb8     1171 L1 L2
mxA.02-00                      0x169        0        0 L1 L2
mxB-R2.00-00                       0x7   0xaca9      940 L1 L2
mxB-R2.02-00                       0x3   0x9cb0      940 L1 L2
mxB-R3.00-00                       0x6   0xbc3e      819 L1 L2
mxB-R3.02-00                       0x3   0x97ac      768 L1 L2
mxB-R3.03-00                       0x3    0x443      819 L1 L2
mxB-R4.00-00                       0x8   0x25d5     1165 L1 L2
mxB-R4.02-00                       0x3   0x3312      814 L1 L2
  11 LSPs
```

Question: Are there any missing IS-IS LSPs?

_____

Answer: No, all five router now have at least one LSP to represent itself in the IS-IS link state database.

Step 2.25

Return to the SSH session that you have open to the mxB router.

From the SSH session with the mxB router, verify that from the R3 router's point of view, routes to all loopbacks are present. For IPv6, use **show route 2001::/16**; for IPv4, use **show route 192.168/16**.

```
lab@mxB:R3> show route 2001::/16

inet6.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
2001:db8:101::1/128*[IS-IS/18] 00:03:23, metric 20
                    > to fe80::250:56ff:fea9:a3a via ge-0/0/3.0
2001:db8:101::2/128*[IS-IS/18] 1d 14:48:30, metric 10
                    > to fe80::250:56ff:fea9:a3a via ge-0/0/3.0
2001:db8:101::3/128*[Direct/0] 1d 14:48:42
                    > via lo0.3
2001:db8:101::4/128*[IS-IS/18] 1d 14:44:04, metric 10
                    > to fe80::250:56ff:fea9:1eeb via ge-0/0/4.0
2001:db8:101::5/128*[IS-IS/18] 00:03:56, metric 20
                    > to fe80::250:56ff:fea9:1eeb via ge-0/0/4.0
2001:db8:101:1::/64*[IS-IS/18] 00:03:23, metric 30
                    > to fe80::250:56ff:fea9:a3a via ge-0/0/3.0
2001:db8:101:2::/64*[IS-IS/18] 00:03:56, metric 30
                    > to fe80::250:56ff:fea9:1eeb via ge-0/0/4.0

lab@mxB:R3> show route 192.168/16

inet.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.71.1/32    *[IS-IS/18] 00:03:38, metric 20
                    > to 172.22.131.9 via ge-0/0/3.0
192.168.71.2/32    *[IS-IS/18] 1d 14:48:46, metric 10
                    > to 172.22.131.9 via ge-0/0/3.0
192.168.71.3/32    *[Direct/0] 1d 14:48:57
                    > via lo0.3
192.168.71.4/32    *[IS-IS/18] 1d 14:44:19, metric 10
                    > to 172.22.131.14 via ge-0/0/4.0
192.168.71.5/32    *[IS-IS/18] 00:04:11, metric 20
                    > to 172.22.131.14 via ge-0/0/4.0
```

Question: Are all of the loopback interface routes active in the routing table?

_____

_____

Answer: Yes, each of the five loopback addresses can be found in the routing table.

Step 2.26

Log out of your assigned devices using the **exit** command.

```
lab@mxB:R3> exit

mxB (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

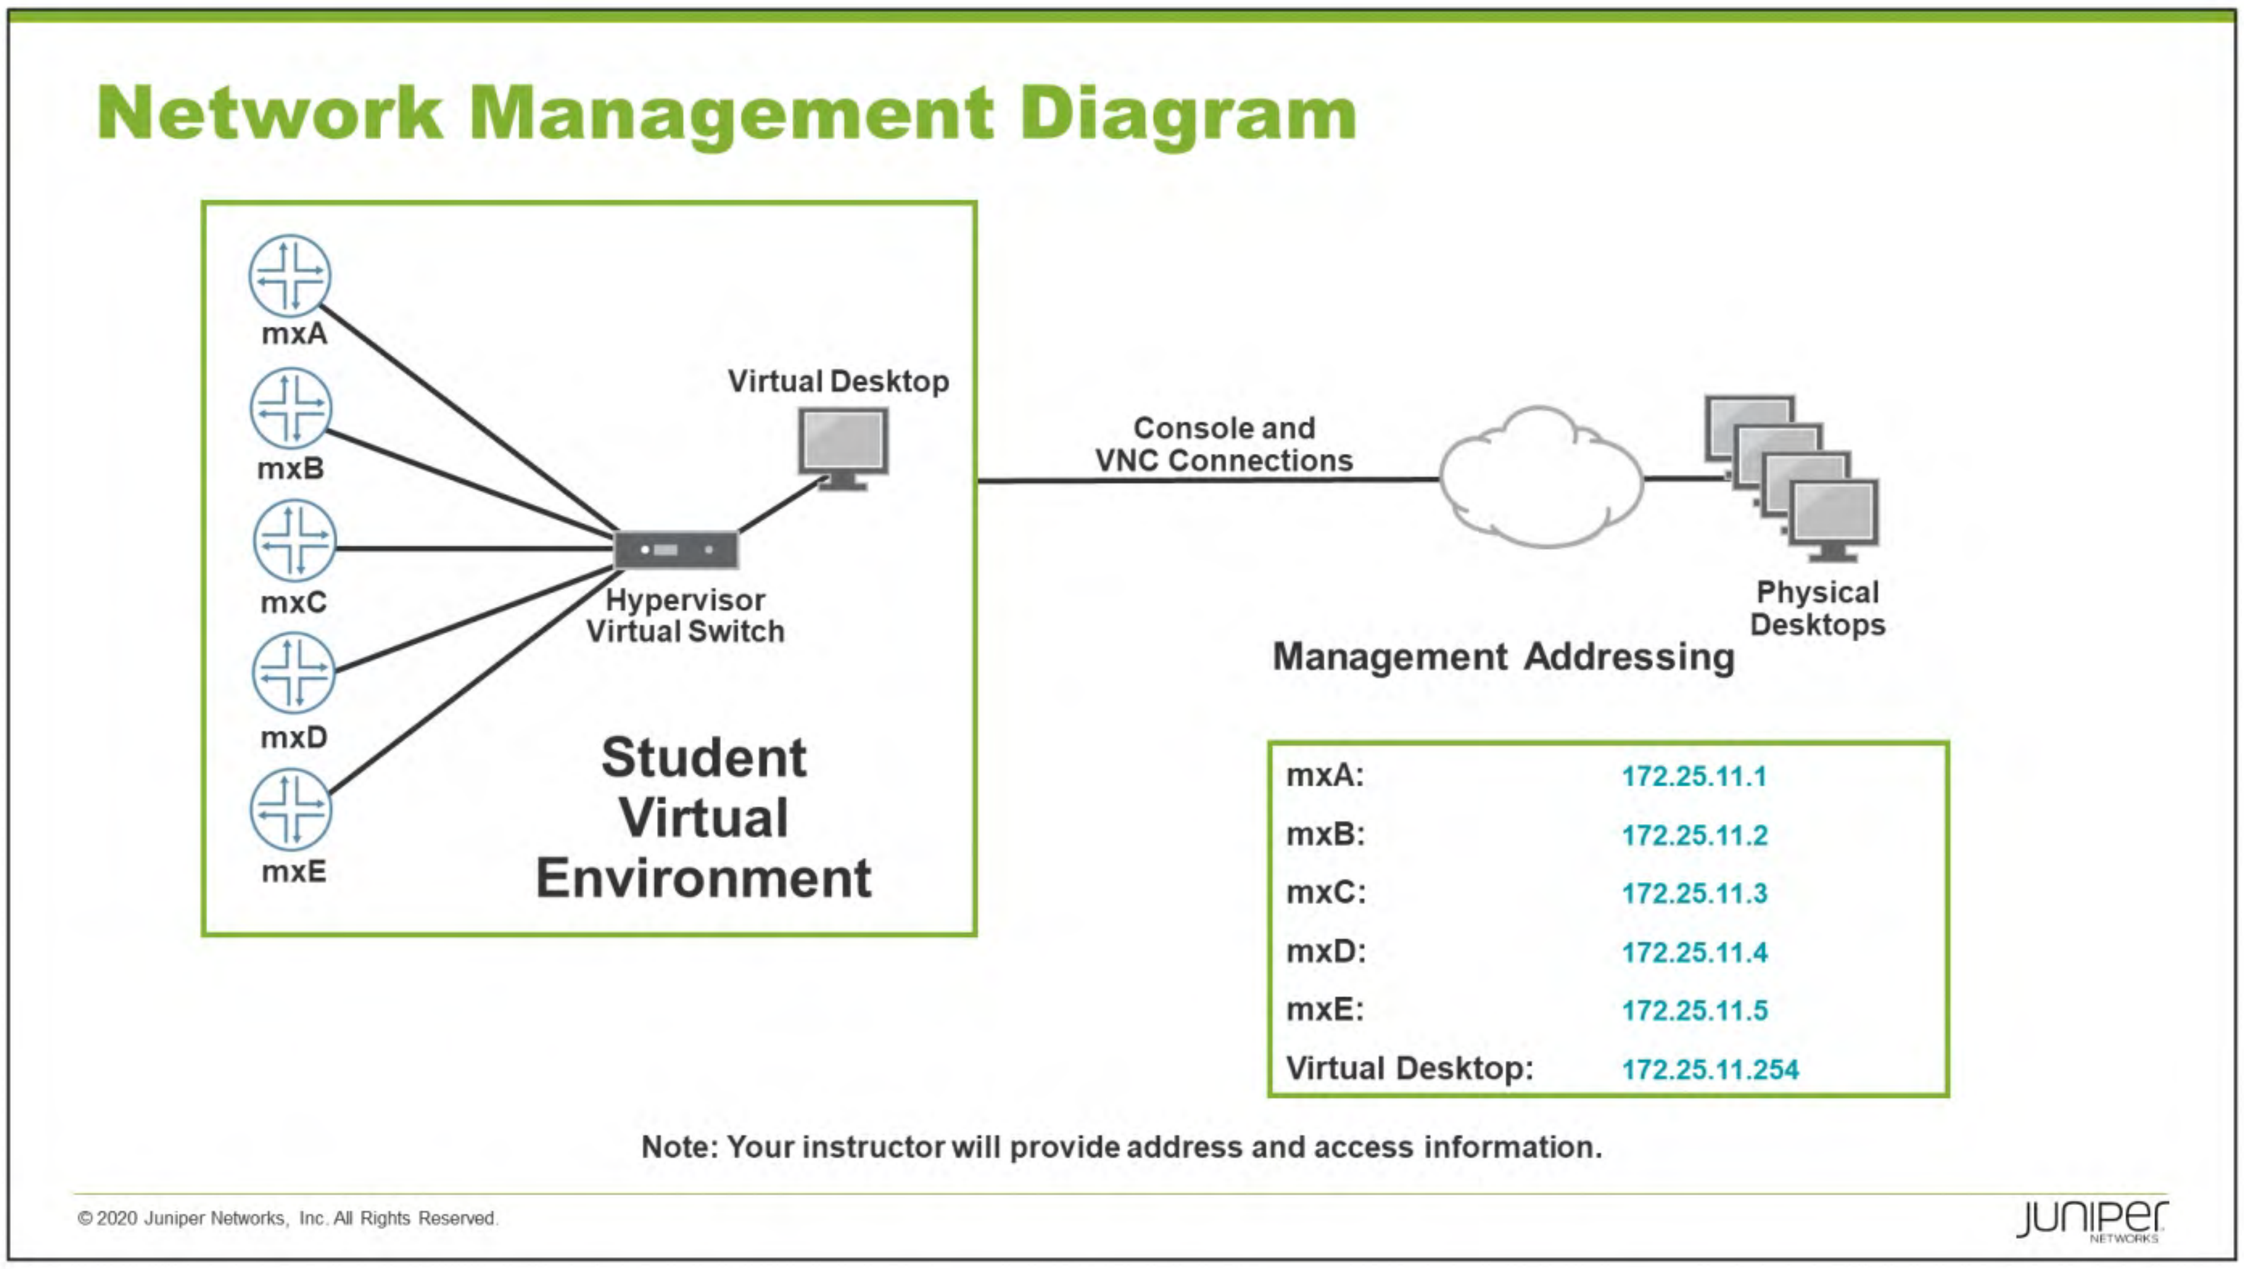
Network Management Diagram
mxA
mxB
mxC
mxD
mxE
Virtual Desktop
Hypervisor
Virtual Switch
Student
Virtual
Environment
Console and
VNC Connections
Physical
Desktops
Management Addressing
mxA: 172.25.11.1
mxB: 172.25.11.2
mxC: 172.25.11.3
mxD: 172.25.11.4
mxE: 172.25.11.5
Virtual Desktop: 172.25.11.254
Note: Your instructor will provide address and access information.
© 2020 Juniper Networks, Inc. All Rights Reserved.
juniper NETWORKS

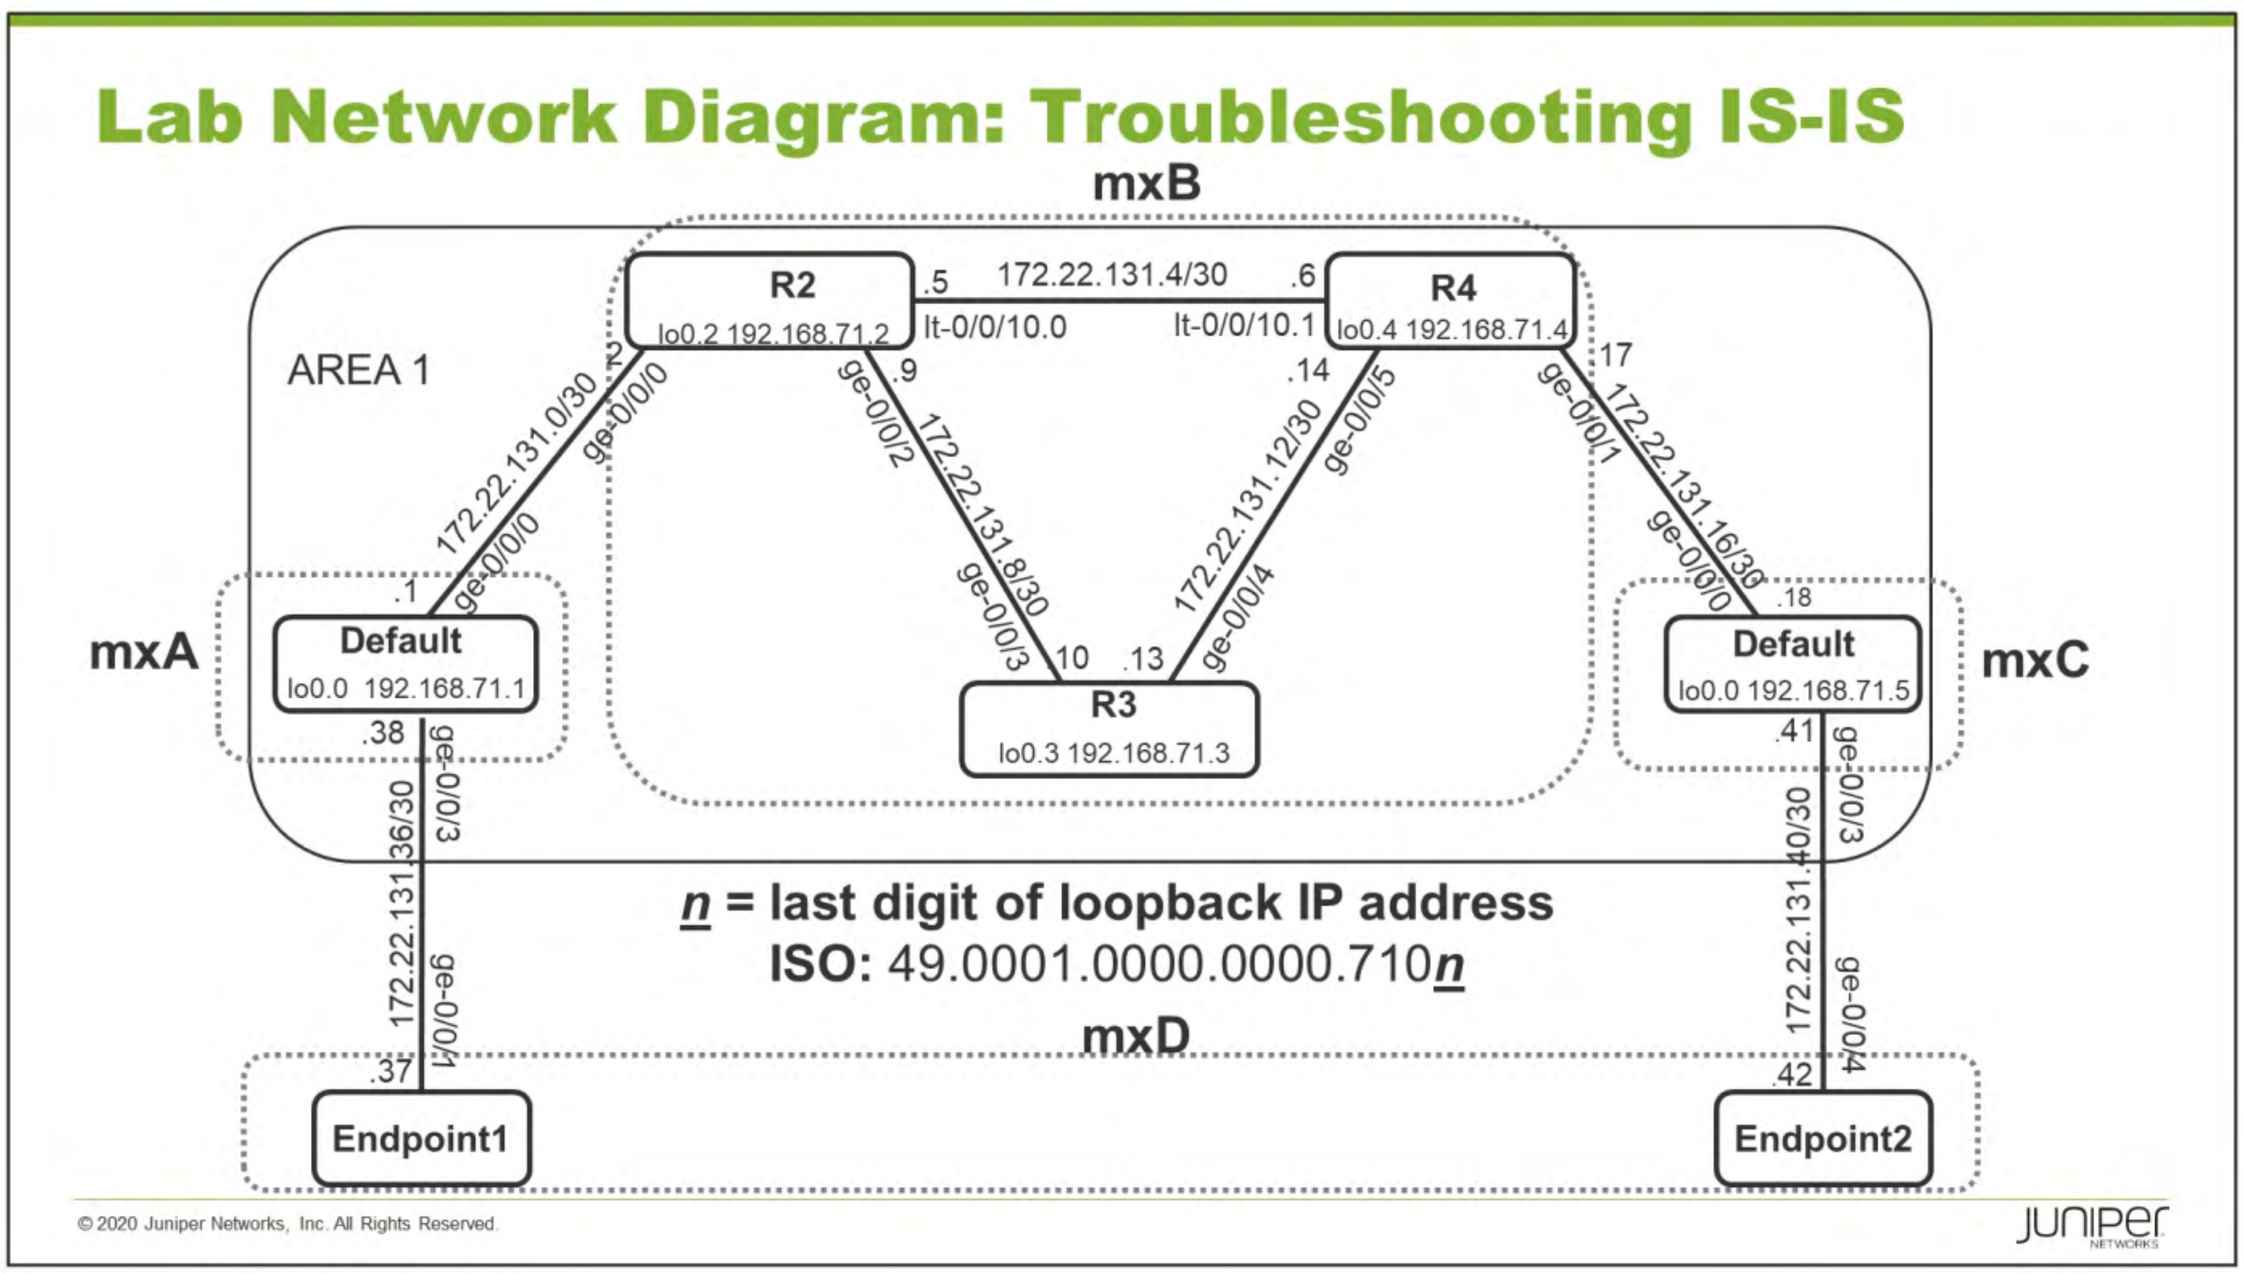
Lab Network Diagram: Troubleshooting IS-IS
mxB
AREA 1
R2
lo0.2 192.168.71.2
.5
172.22.131.4/30
.6
lt-0/0/10.0
lt-0/0/10.1
R4
lo0.4 192.168.71.4
172.22.131.0/30
ge-0/0/0
.2
.1
ge-0/0/0
ge-0/0/2
.9
172.22.131.8/30
ge-0/0/3
.10
.13
R3
lo0.3 192.168.71.3
.14
172.22.131.12/30
ge-0/0/5
ge-0/0/4
ge-0/0/1
.17
172.22.131.16/30
ge-0/0/0
.18
mxA
Default
lo0.0 192.168.71.1
.38
ge-0/0/3
172.22.131.36/30
Default
lo0.0 192.168.71.5
mxC
.41
ge-0/0/3
172.22.131.40/30
n = last digit of loopback IP address
ISO: 49.0001.0000.0000.710n
mxD
ge-0/0/1
.37
Endpoint1
ge-0/0/4
.42
Endpoint2
© 2020 Juniper Networks, Inc. All Rights Reserved.
juniper NETWORKS

Lab

Configuring BGP

Overview

In this lab, you use the lab diagram titled, “Lab Network Diagram: Configuring BGP” to establish a BGP network. You will be configuring the devices in AS 65001 which will initially be running OSPF. After verifying the OSPF Area 0.0.0.0 adjacency, an IBGP session must be established between the mxA router and the R3-1 router. The P1 and P2 routers are in AS 65412, and the P3 router is in AS 65020. You must establish EBGP peering sessions to all three of the P routers. The P3 EBGP peering session peers to the loopback address.

This lab requires the configuration of both internal and EBGP peering sessions.

By completing this lab, you perform the following tasks:

- Load the default configuration.
- Establish an IBGP peering session.
- Establish an EBGP peering session with multipath.
- Establish an EBGP peering session with multihop.
- Use policy to summarize IBGP routes.

Part 1: Load Starting Configuration

In this lab part, you verify the initial configuration of the routers. You then verify that the interfaces are operational and OSPF and BGP neighbor relationships have formed.

Note

The instructor will tell you the nature of your access and will provide you with the necessary details to access your assigned device.

Note

The lab topology requires you to display information in the different logical systems of mxA, mxB, mxC, mxD, and mxE. You will be changing the perspective of the CLI by issuing the **`set cli logical-system`** ***`ls-name`*** command. Refer to the lab diagram for the correct logical system name.

By changing the perspective of the CLI, you will be able to perform network commands such as **`ping`** or **`traceroute`** from the perspective of the pertinent logical system.

Step 1.1

Ensure that you know to which student device you have been assigned. Check with your instructor if you are not certain. Consult the management network diagram to determine the management address of the student devices.

Step 1.2

Access the CLI on your mxB device using Secure Shell (SSH) or as directed by your instructor. Log in as user **`lab`** with the password **`lab123`**.

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab9-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.2
Password: lab123
Last login: Sun Nov 24 11:44:57 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxB> configure
Entering configuration mode

[edit]
lab@mxB# load override ajspr/lab9-start.config
load complete

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

lab@mxB>
```

Step 1.3

Issue the **show configuration** command. Use the lab diagram to verify that the mxB router has the correct logical system and interface configuration. Verify that BGP has been enabled.

```
lab@mxB> show configuration
## Last commit: 2019-11-25 16:50:31 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxB;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    P3 {
        interfaces {
            ge-0/0/3 {
                unit 0 {
                    family inet {
                        address 172.22.125.2/24;
                    }
                }
            }
            ge-0/0/4 {
                unit 0 {
```

```
                    family inet {
                        address 172.22.126.2/24;
                    }
                }
            }
            lo0 {
                unit 0 {
                    family inet {
                        address 172.31.102.1/32;
                    }
                }
            }
        }
        protocols {
            bgp {
                group R3-1 {
                    type external;
                    multihop;
                    local-address 172.31.102.1;
                    export ajspr-bgp-export-p3;
                    neighbor 172.16.1.2 {
                        peer-as 65001;
                    }
                }
                group R3-2 {
                    type external;
                    multihop;
                    local-address 172.31.102.1;
                    export ajspr-bgp-export-p3;
                    neighbor 172.16.2.2 {
                        peer-as 65002;
                    }
                }
            }
        }
        policy-options {
            policy-statement ajspr-bgp-export-p3 {
                term 1 {
                    from {
                        protocol static;
                        route-filter 40.40.0.0/22 orlonger;
                    }
                    then accept;
                }
                term 2 {
                    from {
                        protocol bgp;
                        route-filter 30.30.0.0/22 longer;
                    }
                    then reject;
                }
            }
        }
        routing-options {
            static {
```

```
            route 20.20.0.0/24 reject;
            route 20.20.1.0/24 reject;
            route 20.20.2.0/24 reject;
            route 20.20.3.0/24 reject;
            route 20.20.4.0/25 reject;
            route 20.20.4.128/25 reject;
            route 20.20.5.0/26 reject;
            route 20.20.5.64/26 reject;
            route 20.20.5.128/26 reject;
            route 20.20.5.192/26 reject;
            route 40.40.0.0/24 reject;
            route 40.40.1.0/24 reject;
            route 40.40.2.0/24 reject;
            route 40.40.3.0/24 reject;
            route 172.16.1.2/32 next-hop 172.22.125.1;
            route 172.16.2.2/32 next-hop 172.22.126.1;
            route 172.16.1.4/32 next-hop 172.22.125.1;
            route 172.16.2.4/32 next-hop 172.22.126.1;
            route 10.0.20.0/22 next-hop 172.22.125.1;
            route 10.0.24.0/22 next-hop 172.22.126.1;
        }
        autonomous-system 65020;
        aggregate {
            route 20.20.0.0/21;
        }
    }
}
R3-1 {
    interfaces {
        ge-0/0/0 {
            unit 0 {
                family inet {
                    address 10.0.10.2/24;
                }
            }
        }
        ge-0/0/2 {
            unit 0 {
                family inet {
                    address 172.22.125.1/24;
                }
            }
        }
        lo0 {
            unit 1 {
                family inet {
                    address 172.16.1.2/32;
                }
            }
        }
    }
    protocols {
        ospf {
            area 0.0.0.0 {
                interface lo0.1;
```

```
                    interface ge-0/0/0.0;
                }
            }
        }
        routing-options {
            static {
                route 172.16.1.128/26 reject;
                route 172.16.1.192/26 reject;
            }
        }
    }
    R3-2 {
        interfaces {
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 10.0.14.2/24;
                    }
                }
            }
            ge-0/0/5 {
                unit 0 {
                    family inet {
                        address 172.22.126.1/24;
                    }
                }
            }
            lo0 {
                unit 2 {
                    family inet {
                        address 172.16.2.2/32;
                    }
                }
            }
        }
        protocols {
            ospf {
                area 0.0.0.0 {
                    interface lo0.2;
                    interface ge-0/0/1.0;
                }
            }
            bgp {
                group ibgp {
                    type internal;
                    local-address 172.16.2.2;
                    export redistribute-statics;
                    neighbor 172.16.2.1;
                }
                group P3 {
                    type external;
                    multihop;
                    local-address 172.16.2.2;
                    neighbor 172.31.102.1 {
                        peer-as 65020;
```

```
                }
            }
        }
    }
    policy-options {
        policy-statement redistribute-statics {
            term 1 {
                from protocol static;
                then accept;
            }
        }
    }
    routing-options {
        static {
            route 172.16.2.128/26 reject;
            route 172.16.2.192/26 reject;
            route 172.31.102.1/32 {
                next-hop 172.22.126.2;
                no-readvertise;
            }
        }
        autonomous-system 65002;
        aggregate {
            route 172.16.2.0/24;
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.2/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxB router. If the configuration is incorrect, please notify your instructor.

Question: Which routers are P3 configured to peer with using BGP?

_____

_____

Answer: The P3 router is configured to have an EBGP peering session with the R3-1 and R3-2 router.

Question: What routes will the P3 router be advertising to its EBGP peers?

_____

_____

Answer: The P3 router is configured with an export policy called *ajspr-bgp-export-p3* that is applied to its EBGP peers. This policy causes the P3 router to advertise static routes in the 40.40/22 range.

Question: Which routers are R3-1 configured to peer with using BGP?

_____

_____

Answer: The R3-1 router is not configured for BGP. You will configured BGP on the R3-1 router later in this lab.

Question: Which routers are R3-2 configured to peer with using BGP?

_____

_____

Answer: The R3-2 router is configured for and EBGP session with the P3 router and an IBGP session with the mxC router.

Question: Which interfaces are configured to run OSPF on the R3-2 router? Why is this necessary?

_____

_____

Answer: The R3-2 router is configured to run OSPF on the ge-0/0/1 interface and the lo0.2 interface. This configuration is necessary to allow the R3-2 router and the mxC router to learn each other's loopback address since they will use the loopback interfaces for IBGP peering.

Step 1.4

Issue the **show bgp summary logical-system *P3*** command to view the state of the P3 router's EBGP sessions.

```
lab@mxB> show bgp summary logical-system P3

lab@mxB> show bgp summary logical-system P3
Threading mode: BGP I/O
Groups: 2 Peers: 2 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2            65001          0          0       0       0        7:26 Connect
172.16.2.2            65002         17         18       0       0        7:13 Establ
  inet.0: 0/0/0/0
```

Question: What is the state of the P3 router's BGP neighbor relationships?

_____

_____

Answer: The P3 router should have an established EBGP session with the R3-2 router and an BGP session in the active or connect state with the R3-1 router.

Step 1.5

Access the CLI on your mxC device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab9-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.3
Password: lab123
Last login: Sun Nov 24 13:04:48 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# load override ajspr/lab9-start.config
load complete

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

Step 1.6

Issue the **show configuration** command. Use the lab diagram to verify that the mxC router has the correct interface, OSPF, and BGP configuration.

```
lab@mxC> show configuration
## Last commit: 2019-11-25 16:57:20 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxC;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 10.0.14.1/24;
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.22.124.1/24;
```

```
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 172.22.122.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.3/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.2.1/32;
            }
        }
    }
}
policy-options {
    policy-statement export-aggregate {
        term 1 {
            from {
                protocol aggregate;
                route-filter 172.16.2.0/24 exact;
            }
            then accept;
        }
        term 2 {
            from {
                route-filter 172.16.2.0/24 longer;
            }
            then reject;
        }
    }
    policy-statement pfe-load-balance {
        term 1 {
            from {
                protocol bgp;
                route-filter 30.30.0.0/22 longer;
            }
            then {
                load-balance per-packet;
            }
        }
    }
    policy-statement redistribute-statics {
        term 1 {
            from protocol static;
```

```
                then accept;
            }
        }
    }
    routing-options {
        static {
            route 172.16.2.0/26 reject;
            route 172.16.2.64/26 reject;
        }
        forwarding-table {
            export pfe-load-balance;
        }
        autonomous-system 65002;
        aggregate {
            route 172.16.2.0/24;
        }
    }
    protocols {
        ospf {
            area 0.0.0.0 {
                interface lo0.0;
                interface ge-0/0/0.0;
            }
        }
        bgp {
            group ibgp {
                type internal;
                local-address 172.16.2.1;
                export redistribute-statics;
                neighbor 172.16.2.2;
            }
            group P1-P2 {
                type external;
                export export-aggregate;
                peer-as 65412;
                multipath;
                neighbor 172.22.122.2;
                neighbor 172.22.124.2;
            }
        }
    }
```

Question: Does the configuration match the details of the lab diagram?

_____

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxC router. If the configuration is incorrect, please notify your instructor.

Question: Which routers are mxC configured to peer with using BGP?

_____

_____

Answer: The mxC router is configured with an IBGP session with the R3-2 router and an EBGP session with the P1 and P2 routers.

Question: Which interfaces are configured to run OSPF on the mxC router? Why is this necessary?

_____

_____

Answer: The mxC router is configured to run OSPF on the ge-0/0/0 interface and the lo0.0 interface. This configuration is necessary to allow the R3-2 router and the mxC router to learn each other's loopback address since they will use the loopback interfaces for IBGP peering.

Step 1.7

Issue the **show ospf neighbor** command to ensure the mxC router has established an OSPF neighbor relationship with the R3-2 router.

```
lab@mxC> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.0.14.2        ge-0/0/0.0             Full      172.16.2.2       128    33
```

Question: Has the mxC router formed an OSPF neighbor relationship with the R3-2 router?

_____

Answer: Yes, an OSPF neighbor relationship exists between the mxC and R3-2 routers.

Step 1.8

Issue the **show bgp summary** command to view the state of the mxC router's BGP sessions.

```
lab@mxC> show bgp summary
Threading mode: BGP I/O
Groups: 2 Peers: 3 Down peers: 2
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     6          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.2.2            65002          7          6       0       0        1:35 Establ
  inet.0: 2/6/6/0
172.22.122.2          65412          0          0       0       0        1:45 Connect
172.22.124.2          65412          0          0       0       0        1:45 Connect
```

Question: What is the state of the mxC router's BGP neighbor relationships?

_____

_____

Answer: The mxC router should have an established IBGP session with the R3-2 router and an EBGP session in the active or connect state with the P1 and P2 routers.

Step 1.9

Access the CLI on your mxE device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab9-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.5
Password: lab123
Last login: Sun Nov 24 11:46:46 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxE> configure
Entering configuration mode

[edit]
lab@mxE# load override ajspr/lab9-start.config
load complete

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

lab@mxE>
```

Step 1.10

Issue the **show configuration** command. Use the lab diagram to verify that the mxE router has the correct logical system and interface configuration. Verify that OSPF and BGP has been enabled.

```
lab@mxE> show configuration
## Last commit: 2019-12-27 11:56:28 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxE;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
```

```
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    P2 {
        interfaces {
            ge-0/0/2 {
                unit 0 {
                    family inet {
                        address 172.22.124.2/24;
                    }
                }
            }
            ge-0/0/5 {
                unit 0 {
                    family inet {
                        address 172.22.252.2/30;
                    }
                }
            }
            ge-0/0/7 {
                unit 0 {
                    family inet {
                        address 172.22.123.2/24;
                    }
                }
            }
            lo0 {
                unit 0 {
                    family inet {
                        address 172.31.101.1/32;
                    }
                }
```

```
            }
        }
        protocols {
            ospf {
                area 0.0.0.0 {
                    interface lo0.0;
                    interface ge-0/0/5.0;
                }
            }
            bgp {
                group ibgp {
                    type internal;
                    local-address 172.31.101.1;
                    export ajspr-bgp-nhs;
                    neighbor 172.31.100.1;
                }
                group mxA {
                    type external;
                    export ajspr-bgp-export;
                    neighbor 172.22.123.1 {
                        peer-as 65001;
                    }
                }
                group mxC {
                    type external;
                    export ajspr-bgp-export;
                    neighbor 172.22.124.1 {
                        peer-as 65002;
                    }
                }
            }
        }
        policy-options {
            policy-statement ajspr-bgp-export {
                term 1 {
                    from {
                        protocol static;
                        route-filter 30.30.0.0/22 orlonger;
                    }
                    then accept;
                }
                term 2 {
                    from {
                        protocol bgp;
                        route-filter 40.40.0.0/22 longer;
                    }
                    then reject;
                }
            }
            policy-statement ajspr-bgp-nhs {
                term 1 {
                    from {
                        protocol bgp;
                        route-type external;
                    }
```

```
                    then {
                        next-hop self;
                    }
                }
            }
        }
        routing-options {
            static {
                route 30.30.0.0/24 reject;
                route 30.30.1.0/24 reject;
                route 30.30.2.0/24 reject;
                route 30.30.3.0/24 reject;
            }
            autonomous-system 65412;
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.5/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxE router. If the configuration is incorrect, please notify your instructor.

Question: Based on the configuration, which routers (logical systems) will mxE represent in this lab?

_____

_____

Answer: The mxE device will, using logical systems, represent the P2 router.

Question: Which interfaces are configured to run OSPF on the P2 router? Why is this necessary?

_____

_____

Answer: The P2 router is configured to run OSPF on the ge-0/0/5 interface and the lo0.0 interface. This configuration is necessary to allow the P2 router and the P1 router to learn each other's loopback address since they will use the loopback interfaces for IBGP peering.

Question: Which routers are P2 configured to peer with using BGP?

_____

_____

Answer: The P2 router is configured with an IBGP session with the P1 router and an EBGP session with the mxA and mxC routers.

Question: What routes will the P2 router be advertising to its EBGP peers?

_____

_____

Answer: The P2 router is configured with an export policy called *ajspr-bgp-export* that is applied to its EBGP peers. This policy causes the P2 router to advertise any static routes in the 30.30/22 range.

Step 1.11

Issue the **show bgp summary logical-system *P2*** command to view the state of the P2 router's BGP sessions.

```
lab@mxE> show bgp summary logical-system P2
Threading mode: BGP I/O
Groups: 3 Peers: 3 Down peers: 2
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...
172.22.123.1          65001          0          0       0       0         49 Connect
172.22.124.1          65002          4          3       0       0         37 Establ
  inet.0: 1/1/1/0
172.31.100.1          65412          0          0       0       0         49 Active
```

Question: What is the state of the P2 router's BGP neighbor relationships?

_____

_____

Answer: The P2 router should have an established EBGP session with the mxC router, an IBGP session in the active or connect state with the P1 router, and an EBGP session in the active or connect state with the mxA router.

Step 1.12

Access the CLI on your mxD device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123.**

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab9-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.4
Password: lab123
Last login: Sun Nov 24 13:00:47 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab9-start.config
load complete

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

lab@mxD>
```

Step 1.13

Issue the **show configuration** command. Use the lab diagram to verify that the mxD router has been configured for a logical system called P1 that has the correct interface and OSPF configuration.

```
lab@mxD> show configuration
## Last commit: 2019-11-25 17:05:10 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxD;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
```

```
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
logical-systems {
    P1 {
        interfaces {
            ge-0/0/0 {
                unit 0 {
                    family inet {
                        address 172.22.252.1/30;
                    }
                }
            }
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 172.22.121.2/24;
                    }
                }
            }
            ge-0/0/4 {
                unit 0 {
                    family inet {
                        address 172.22.122.2/24;
                    }
                }
            }
            lo0 {
                unit 0 {
                    family inet {
                        address 172.31.100.1/32;
                    }
```

```
            }
        }
    }
    protocols {
        ospf {
            area 0.0.0.0 {
                interface lo0.0;
                interface ge-0/0/0.0;
            }
        }
        bgp {
            group ibgp {
                type internal;
                local-address 172.31.100.1;
                export ajspr-bgp-nhs;
                neighbor 172.31.101.1;
            }
            group mxA {
                type external;
                export ajspr-bgp-export;
                neighbor 172.22.121.1 {
                    peer-as 65001;
                }
            }
            group mxC {
                type external;
                export ajspr-bgp-export;
                neighbor 172.22.122.1 {
                    peer-as 65002;
                }
            }
        }
    }
    policy-options {
        policy-statement ajspr-bgp-export {
            term 1 {
                from {
                    protocol static;
                    route-filter 30.30.0.0/22 orlonger;
                }
                then accept;
            }
            term 2 {
                from {
                    protocol bgp;
                    route-filter 40.40.0.0/22 longer;
                }
                then reject;
            }
        }
        policy-statement ajspr-bgp-nhs {
            term 1 {
                from {
                    protocol bgp;
                    route-type external;
```

```
                    }
                    then {
                        next-hop self;
                    }
                }
            }
        }
        routing-options {
            static {
                route 30.30.0.0/24 reject;
                route 30.30.1.0/24 reject;
                route 30.30.2.0/24 reject;
                route 30.30.3.0/24 reject;
            }
            autonomous-system 65412;
        }
    }
}
interfaces {
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.4/24;
            }
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications. If the configuration is incorrect, please notify your instructor.

Question: Based on the configuration, which routers (logical systems) will mxD represent in this lab?

_____

_____

Answer: The mxD device will, using logical systems, represent the P1 router.

Question: Which interfaces are configured to run OSPF on the P1 router? Why is this necessary?

_____

_____

Answer: The P1 router is configured to run OSPF on the ge-0/0/0 interface and the lo0.0 interface. This configuration is necessary to allow the P1 router and the P2 router to learn each other's loopback address since they will use the loopback interfaces for IBGP peering.

Question: Which routers are P1 configured to peer with using BGP?

_____

_____

Answer: The P1 router is configured with an IBGP session with the P2 router and an EBGP session with the mxA and mxC routers.

Question: What routes will the P1 router be advertising to its EBGP peers?

_____

_____

Answer: The P1 router is configured with an export policy called *ajspr-bgp-export* that is applied to its EBGP peers. This policy causes the P1 router to advertise any static routes in the 30.30/22 range.

Step 1.14

Issue the **show ospf neighbor logical-system *P1*** command to ensure the P1 router has established an OSPF neighbor relationship with the P2 router.

```
lab@mxD> show ospf neighbor logical-system P1
Address          Interface              State     ID               Pri  Dead
172.22.252.2     ge-0/0/0.0             Full      172.31.101.1     128    32
```

Question: Has the P1 router formed an OSPF neighbor relationship with the P2 router?

_____

Answer: Yes, an OSPF neighbor relationship exists between the P1 and P2 routers.

Step 1.15

Issue the **show bgp summary logical-system *P1*** command to view the state of the P1 router's BGP sessions.

```
lab@mxD> show bgp summary logical-system P1
Threading mode: BGP I/O
Groups: 3 Peers: 3 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       2          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.121.1          65001          0          0       0       0        1:02 Connect
172.22.122.1          65002          4          4       0       0          51 Establ
  inet.0: 1/1/1/0
172.31.101.1          65412          4          3       0       0          44 Establ
  inet.0: 0/1/1/0
```

Question: What is the state of the P1 router's BGP neighbor relationships?

Answer: The P1 router should have an established IBGP session with the P2 router, an established EBGP session with the mxC router, and an EBGP session in the active or connect state with the mxA router.

Step 1.16

Access the CLI on your mxA device using Secure Shell (SSH) or as directed by your instructor. Log in as user **lab** with the password **lab123**.

Enter configuration mode and load the device's starting configuration by issuing the **load override *ajspr/lab9-start.config*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
[lab@desktop ~]$ ssh lab@172.25.11.1
Password: lab123
Last login: Sun Nov 24 11:44:57 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# load override ajspr/lab9-start.config
load complete

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

lab@mxA>
```

Step 1.17

Issue the **show configuration** command. Use the lab diagram to verify that the mxA router has the correct interface and OSPF configuration.

```
lab@mxA> show configuration
## Last commit: 2019-11-25 17:08:14 UTC by lab
version 20190829.221548_builder.r1052644;
system {
    host-name mxA;
    root-authentication {
        encrypted-password
   "$6$GBq0FqVe$MXFb2Y.AmNd1hw1a9MApGk.md1pusC7vWkzduzRuuafxYgZT2XJzpU40sb/
   .3tQ0bCdrevo0v635R7ReHTL0c."; ## SECRET-DATA
    }
    login {
        user lab {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$6$JEnFYM1n$C6pjHzEv3cK/iovqkiJywOgyrmgNnX/
   U0r3B3kfaeXa4ygKFE1l7De7YsKVbjdJnab3PfylSLmDHjgPVMiilP1"; ## SECRET-DATA
            }
        }
    }
    services {
        ssh;
        netconf {
            ssh;
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                address 10.0.10.1/24;
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.22.123.1/24;
            }
        }
    }
    ge-0/0/3 {
```

```
        unit 0 {
            family inet {
                address 172.22.121.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 172.25.11.1/24;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.16.1.1/32;
            }
        }
    }
}
routing-options {
    static {
        route 172.16.1.0/26 reject;
        route 172.16.1.64/26 reject;
    }
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-0/0/0.0;
        }
    }
}
```

Question: Does the configuration match the details of the lab diagram?

_____

Answer: Yes, the configuration should match the lab diagram specifications for the mxA router. If the configuration is incorrect, please notify your instructor.

Step 1.18

Issue the **`show ospf neighbor`** command to ensure the mxA router has established an OSPF neighbor relationship with the R3-1 router.

```
lab@mxA> show ospf neighbor
Address          Interface              State     ID               Pri  Dead
10.0.10.2        ge-0/0/0.0             Full      172.16.1.2       128    34
```

Question: Has the mxA router formed an OSPF neighbor relationship with the R3-1 router?

_____

Answer: Yes, an OSPF neighbor relationship exists between the mxA and R3-1 routers.

Step 1.19

Issue the **show bgp summary** command to view the status of the mxA router's BGP neighbor relationships.

```
lab@mxA> show bgp summary
BGP is not running

lab@mxA>
```

Question: What is the status of the mxA router's BGP neighbor relationships? Why?

_____

_____

Answer: The mxA router does not have any BGP neighbor relationships because BGP has not been configured on the mxA router. You will configure BGP on the mxA router in the following lab parts.

Step 1.20

Use the `ping` utility to verify reachability to the directly connected interfaces of the P1 and P2 routers.

```
lab@mxA> ping 172.22.121.2 count 5
PING 172.22.121.2 (172.22.121.2): 56 data bytes
64 bytes from 172.22.121.2: icmp_seq=0 ttl=64 time=1.484 ms
64 bytes from 172.22.121.2: icmp_seq=1 ttl=64 time=1.355 ms
64 bytes from 172.22.121.2: icmp_seq=2 ttl=64 time=1.222 ms
64 bytes from 172.22.121.2: icmp_seq=3 ttl=64 time=1.185 ms
64 bytes from 172.22.121.2: icmp_seq=4 ttl=64 time=11.899 ms

--- 172.22.121.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.185/3.429/11.899/4.236 ms

lab@mxA> ping 172.22.123.2 count 5
PING 172.22.123.2 (172.22.123.2): 56 data bytes
64 bytes from 172.22.123.2: icmp_seq=0 ttl=64 time=1.024 ms
64 bytes from 172.22.123.2: icmp_seq=1 ttl=64 time=1.235 ms
64 bytes from 172.22.123.2: icmp_seq=2 ttl=64 time=1.307 ms
64 bytes from 172.22.123.2: icmp_seq=3 ttl=64 time=1.110 ms
64 bytes from 172.22.123.2: icmp_seq=4 ttl=64 time=1.028 ms
```

```
--- 172.22.123.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.024/1.141/1.307/0.113 ms
```

Question: Are the ping tests successful?

_____

Answer: Yes, the ping tests should be successful at this time. If your tests are not successful, check with your instructor.

Step 1.21

Use the **show route table inet.0** command to display the routing table.

```
lab@mxA> show route table inet.0

inet.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.10.0/24       *[Direct/0] 00:06:50
                    > via ge-0/0/0.0
10.0.10.1/32       *[Local/0] 00:06:50
                      Local via ge-0/0/0.0
172.16.1.0/26      *[Static/5] 00:06:50
                      Reject
172.16.1.1/32      *[Direct/0] 00:06:50
                    > via lo0.0
172.16.1.2/32      *[OSPF/10] 00:06:45, metric 1
                    > to 10.0.10.2 via ge-0/0/0.0
172.16.1.64/26     *[Static/5] 00:06:50
                      Reject
172.22.121.0/24    *[Direct/0] 00:06:50
                    > via ge-0/0/3.0
172.22.121.1/32    *[Local/0] 00:06:50
                      Local via ge-0/0/3.0
172.22.123.0/24    *[Direct/0] 00:06:50
                    > via ge-0/0/1.0
172.22.123.1/32    *[Local/0] 00:06:50
                      Local via ge-0/0/1.0
172.25.11.0/24     *[Direct/0] 02:15:27
                    > via fxp0.0
172.25.11.1/32     *[Local/0] 02:15:27
                      Local via fxp0.0
224.0.0.5/32       *[OSPF/10] 00:06:50, metric 1
                      MultiRecv
```

Question: Are the loopback addresses of the mxA and R3-1 routers active in the routing table?

_____

Answer: Yes. The loopback addresses are active in the routing table. The R3-1 router's loopback is an active OSPF route and the mxA router's loopback is a direct route.

Part 2: Establishing an IBGP Peering Session

In this lab part, you configure the Internal BGP (IBGP) session between the mxA and R3-1 routers. The IBGP session should use loopback addresses for peering between the two routers. After configuring the IBGP group, verify that the session is established and redistribute the two static routes configured under the `[edit routing-options]` hierarchy.

Step 2.1

The autonomous system (AS) number must be configured on both routers.

Enter configuration mode, navigate to the `[edit routing-options]` and configure the AS number for your *`default`* routing instance. Refer to your lab diagram as needed.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit routing-options

[edit routing-options]
lab@mxA# set autonomous-system 65001

[edit routing-options]
lab@mxA#
```

Step 2.2

Navigate to the `[edit protocols bgp]` hierarchy and configure a BGP group named *`ibgp`* that establishes an IBGP peering session with the R3-1 loopback address. Commit the configuration when completed.

```
[edit routing-options]
lab@mxA# top edit protocols bgp

[edit protocols bgp]
lab@mxA# set group ibgp type internal

[edit protocols bgp]
lab@mxA# set group ibgp local-address 172.16.1.1

[edit protocols bgp]
lab@mxA# set group ibgp neighbor 172.16.1.2

[edit protocols bgp]
lab@mxA# commit
commit complete

[edit protocols bgp]
lab@mxA#
```

Step 2.3

Return to the SSH session that you have open to the mxB router.

From the SSH session with the mxB router, change the perspective of the CLI to that of the R3-1 logical system by issuing the **`set cli logical-system R3-1`** command.

```
lab@mxB> set cli logical-system R3-1
Logical system: R3-1

lab@mxB:R3-1>
```

Step 2.4

Enter configuration mode and navigate to the `[edit routing-options]` hierarchy and configure the AS number for the R3-1 routing instance (it should be the same as the one configured for the mxA router).

```
lab@mxB:R3-1> configure
Entering configuration mode

[edit]
lab@mxB:R3-1# edit routing-options

[edit routing-options]
lab@mxB:R3-1# set autonomous-system 65001

[edit routing-options]
lab@mxB:R3-1#
```

Step 2.5

Navigate to the `[edit protocols bgp]` hierarchy and configure a BGP group named *ibgp* that establishes an IBGP peering session with the mxA router. Commit the configuration when completed.

```
[edit routing-options]
lab@mxB:R3-1# top edit protocols bgp

[edit protocols bgp]
lab@mxB:R3-1# set group ibgp type internal

[edit protocols bgp]
lab@mxB:R3-1# set group ibgp local-address 172.16.1.2

[edit protocols bgp]
lab@mxB:R3-1# set group ibgp neighbor 172.16.1.1

[edit protocols bgp]
lab@mxB:R3-1# commit
commit complete

[edit protocols bgp]
lab@mxB:R3-1#
```

Step 2.6

Use the **`run show bgp summary`** command to verify that the IBGP session has been established.

```
[edit protocols bgp]
lab@mxB:R3-1# run show bgp summary
Threading mode: BGP I/O
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...
172.16.1.1            65001          2          2       0       0          6 Establ
  inet.0: 0/0/0/0
```

Question: Is the IBGP session established? Are any BGP routes being exchanged between the peers?

_____

_____

Answer: The IBGP session is established between the two routing instances. If the peering session is not established, check your configuration and if necessary consult with your instructor. The 0/0/0/0 identifies the Active/Received/Accepted/Damped routes; therefore, no active BGP routes are being exchanged at this time.

Step 2.7

Issue the **`run show route protocol static`** command to view the static route's that exist in the R3-1 router's routing table.

```
[edit protocols bgp]
lab@mxB:R3-1# run show route protocol static

inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.128/26    *[Static/5] 4d 18:50:30
                      Reject
172.16.1.192/26    *[Static/5] 4d 18:50:30
                      Reject

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

[edit protocols bgp]
lab@mxB:R3-1#
```

Question: What static routes exist in the R3-1 router's routing table?

_____

_____

Answer: There are two static routes in the R3-1 router's routing table; 172.16.1.128/26 and 172.16.1.192/26. Each static route has a next hop of reject.

Step 2.8

Navigate to the `[edit policy-options policy-statement redistribute-statics]` hierarchy and create a policy that accepts all static routes.

```
[edit protocols bgp]
lab@mxB:R3-1# top edit policy-options policy-statement redistribute-statics

[edit policy-options policy-statement redistribute-statics]
lab@mxB:R3-1# set term 1 from protocol static

[edit policy-options policy-statement redistribute-statics]
lab@mxB:R3-1# set term 1 then accept

[edit policy-options policy-statement redistribute-statics]
lab@mxB:R3-1#
```

Step 2.9

Navigate to the `[edit protocols bgp]` hierarchy and apply the *redistribute-statics* policy as an export policy in the protocols bgp group *ibgp* hierarchy. Commit the configuration when completed.

```
[edit policy-options policy-statement redistribute-statics]
lab@mxB:R3-1# top edit protocols bgp

[edit protocols bgp]
lab@mxB:R3-1# set group ibgp export redistribute-statics

[edit protocols bgp]
lab@mxB:R3-1# commit
commit complete

[edit protocols bgp]
lab@mxB:R3-1#
```

Step 2.10

Verify that the R3-1 router is advertising its static routes to the mxA router by issuing the **`run show route advertising-protocol bgp 172.16.1.1`** command.

```
[edit protocols bgp]
lab@mxB:R3-1# run show route advertising-protocol bgp 172.16.1.1

inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.128/26         Self                         100        I
* 172.16.1.192/26         Self                         100        I
```

Question: How many routes are being advertised to the mxA router?

_____

Answer: The R3-1 router should be advertising two routes to the mxA router.

Step 2.11

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, use the **`run show bgp summary`** command to verify that the mxA router is receiving the routes being advertised by the R3-1 router.

```
[edit protocols bgp]
lab@mxA# run show bgp summary
Threading mode: BGP I/O
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       2          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2            65001          9          7       0       0        2:21 Establ
  inet.0: 2/2/2/0
```

Question: Is the IBGP session established? Are any BGP routes being exchanged between the peers?

_____

_____

Answer: The IBGP session is established between the two routers. The 2/2/2/0 identifies the Active/Received/Accepted/Damped routes; therefore, two active BGP routes are being received from the R3-1 router.

Step 2.12

Issue the **`run show route protocol static`** command to view the static route's that exist in the mxA router's routing table.

```
[edit protocols bgp]
lab@mxA# run show route protocol static

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/26      *[Static/5] 00:45:16
                      Reject
172.16.1.64/26     *[Static/5] 00:45:16
                      Reject

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: What static routes exist in the mxA router's routing table?

_____

_____

Answer: Two static routes exist in the mxA router's routing table; 172.16.1.0/26 and 172.16.1.64/26. Each static route has a next hop of reject.

Step 2.13

Navigate to the `[edit policy-options policy-statement redistribute-statics]` hierarchy and create a policy that accepts all static routes.

```
[edit protocols bgp]
lab@mxA# top edit policy-options policy-statement redistribute-statics

[edit policy-options policy-statement redistribute-statics]
lab@mxA# set term 1 from protocol static

[edit policy-options policy-statement redistribute-statics]
lab@mxA# set term 1 then accept

[edit policy-options policy-statement redistribute-statics]
lab@mxA#
```

Step 2.14

Navigate to the `[edit protocols bgp]` hierarchy and apply the *redistribute-statics* policy as an export policy in the protocols bgp group *ibgp* hierarchy. Commit the configuration when completed.

```
[edit policy-options policy-statement redistribute-statics]
lab@mxA# top edit protocols bgp

[edit protocols bgp]
lab@mxA# set group ibgp export redistribute-statics

[edit protocols bgp]
lab@mxA# commit
commit complete

[edit protocols bgp]
lab@mxA#
```

Step 2.15

Verify that the mxA router is advertising its static routes to the R3-1 router by issuing the **`run show route advertising-protocol bgp 172.16.1.2`** command.

```
[edit protocols bgp]
lab@mxA# run show route advertising-protocol bgp 172.16.1.2

inet.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.0/26           Self                         100        I
* 172.16.1.64/26          Self                         100        I
```

Question: How many routes are being advertised to the R3-1 router?

_____

Answer: The mxA router should be advertising two routes to the R3-1 router.

Part 3: Configuring the P1 and P2 EBGP Peers

In this lab part, you configure the mxA router to have two EBGP peering sessions with the P1 and P2 routers. These devices are both in AS 65412.

Step 3.1

Configure a BGP group named *P1-P2* that establishes an EBGP peering session with the directly connected interfaces of the P1 and P2 routers. Refer to the network diagram for this lab as necessary. Commit the configuration when completed.

```
[edit protocols bgp]
lab@mxA# set group P1-P2 type external

[edit protocols bgp]
lab@mxA# set group P1-P2 neighbor 172.22.121.2

[edit protocols bgp]
lab@mxA# set group P1-P2 neighbor 172.22.123.2

[edit protocols bgp]
lab@mxA# set group P1-P2 peer-as 65412

[edit protocols bgp]
lab@mxA# commit
commit complete

[edit protocols bgp]
lab@mxA#
```

Step 3.2

Use the **run show bgp summary** command to verify that the EBGP sessions to P1 and P2 routers are established.

```
[edit protocols bgp]
lab@mxA# run show bgp summary
Threading mode: BGP I/O
Groups: 2 Peers: 3 Down peers: 0
```

```
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                        12          7          0          0          0          0
Peer                       AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2              65001         15         16       0       0        5:10 Establ
  inet.0: 2/2/2/0
172.22.121.2            65412          4          3       0       0          10 Establ
  inet.0: 5/5/5/0
172.22.123.2            65412          4          3       0       0           6 Establ
  inet.0: 0/5/5/0
```

Question: Are the P1 and P2 peers established and how many routes are being received from the P1 and P2 routers?

_____

_____

Answer: The P1 and P2 BGP sessions should be established. The mxA router should be receiving 5 routes from both EBGP neighbors.

Step 3.3

The BGP `neighbor` command has a lot of valuable information. Use the **`run show bgp neighbor 172.22.121.2`** command to view the P1 EBGP peer.

```
[edit protocols bgp]
lab@mxA# run show bgp neighbor 172.22.121.2
Peer: 172.22.121.2+179 AS 65412 Local: 172.22.121.1+59615 AS 65001
  Group: P1-P2                    Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference PeerAS Refresh>
  Options: <GracefulShutdownRcv>
  Holdtime: 90 Preference: 170
  Graceful Shutdown Receiver local-preference: 0
  Number of flaps: 0
  Peer ID: 172.31.100.1    Local ID: 172.16.1.1        Active Holdtime: 90
  Keepalive Interval: 30        Group index: 1    Peer index: 0    SNMP index: 20
  I/O Session Thread: bgpio-0 State: Enabled
  BFD: disabled, down
  Local Interface: ge-0/0/3.0
  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
  Restart flag received from the peer: Notification
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
```

```
Peer does not support LLGR Restarter functionality
Peer supports 4 byte AS extension (peer-as 65412)
Peer does not support Addpath
Table inet.0 Bit: 20001
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:              5
  Received prefixes:            5
  Accepted prefixes:            5
  Suppressed due to damping:    0
  Advertised prefixes:          2
Last traffic (seconds): Received 22   Sent 21   Checked 49
Input messages:  Total 5      Updates 3       Refreshes 0     Octets 182
Output messages: Total 3      Updates 1       Refreshes 0     Octets 95
Output Queue[1]: 0            (inet.0, inet-unicast)
```

Question: On which TCP port is the peer established on? On which TCP port is the local peer established on?

_____

_____

Answer: In the above capture, the remote peer router is established on TCP port 179 and the local router is established on TCP port 59615. Your results might vary.

Question: Which peer established this session?

_____

Answer: In the example the remote peer router established the session. The local router initiated the TCP session by sending a TCP sync to 172.22.121.2 port 179 from 172.22.121.1 port 59615.

Step 3.4

Use the **`run show route receive-protocol bgp 172.22.121.2`** and **`run show route receive-protocol bgp 172.22.123.2`** commands to view the routes being received from the P1 and P2 routers.

```
[edit protocols bgp]
lab@mxA# run show route receive-protocol bgp 172.22.121.2

inet.0: 20 destinations, 25 routes (20 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  30.30.0.0/24            172.22.121.2                            65412 I
  30.30.1.0/24            172.22.121.2                            65412 I
  30.30.2.0/24            172.22.121.2                            65412 I
  30.30.3.0/24            172.22.121.2                            65412 I
  172.16.2.0/24           172.22.121.2                            65412 65002 I

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

[edit protocols bgp]
```

```
lab@mxA# run show route receive-protocol bgp 172.22.123.2

inet.0: 20 destinations, 25 routes (20 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 30.30.0.0/24            172.22.123.2                            65412 I
* 30.30.1.0/24            172.22.123.2                            65412 I
* 30.30.2.0/24            172.22.123.2                            65412 I
* 30.30.3.0/24            172.22.123.2                            65412 I
* 172.16.2.0/24           172.22.123.2                            65412 65002 I

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Are the same routes being received from both the P1 and P2 routers?

_____

Answer: Yes, the same routes are being received from the P1 and P2 routers. The routes marked active may vary.

Step 3.5

Display the 30.30.0.0/24 route using the **`run show route 30.30/24 detail`** command.

```
[edit protocols bgp]
lab@mxA# run show route 30.30/24 detail

inet.0: 20 destinations, 25 routes (20 active, 0 holddown, 0 hidden)
30.30.0.0/24 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 582
                Address: 0xc0f22d4
                Next-hop reference count: 10
                Source: 172.22.121.2
                Next hop: 172.22.121.2 via ge-0/0/3.0, selected
                Session Id: 0x187
                State: <Active Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 7:09
                Validation State: unverified
                Task: BGP_65412.172.22.121.2
                Announcement bits (3): 0-KRT 2-BGP_RT_Background 5-Resolve tree 5
                AS path: 65412 I
                Accepted
                Localpref: 100
                Router ID: 172.31.100.1
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0xc0f1b68
                Next-hop reference count: 5
                Source: 172.22.123.2
                Next hop: 172.22.123.2 via ge-0/0/1.0, selected
                Session Id: 0x0
                State: <NotBest Ext Changed>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65001 Peer AS: 65412
```

```
Age: 7:05
Validation State: unverified
Task: BGP_65412.172.22.123.2
AS path: 65412 I
Accepted
Localpref: 100
Router ID: 172.31.101.1
```

Question: The 30.30.0.0/24 route is being received from both the P1 and P2 peers. Which route is currently the active route? Why?

_____

_____

Answer: The P2 route is the active route. The inactive reason in the P1 route is "`Not Best in its group - Active preferred`", which indicates that the P1 route's selection criteria cannot override the current active route.

Question: How many next hops do you see for the active route?

_____

_____

Answer: At this point you should see only one next hop.

Step 3.6

Configure the BGP `multipath` option within the *`P1-P2`* group to install the P1 and P2 routes with two equal cost paths. Commit the configuration when completed.

```
[edit protocols bgp]
lab@mxA# set group P1-P2 multipath

[edit protocols bgp]
lab@mxA# commit
commit complete

[edit protocols bgp]
lab@mxA#
```

Step 3.7

Display the 30.30.0.0/24 route again using the **`run show route 30.30.0.0/24 detail`** command.

```
[edit protocols bgp]
lab@mxA# run show route 30.30/24 detail

inet.0: 20 destinations, 25 routes (20 active, 0 holddown, 0 hidden)
30.30.0.0/24 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0xcad0e4c
                Next-hop reference count: 6
```

```
                Source: 172.22.121.2
                Next hop: 172.22.123.2 via ge-0/0/1.0, selected
                Session Id: 0x0
                Next hop: 172.22.121.2 via ge-0/0/3.0
                Session Id: 0x0
                State: <Active Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 8
                Validation State: unverified
                Task: BGP_65412.172.22.121.2
                Announcement bits (4): 0-KRT 2-BGP_RT_Background 5-Resolve tree 5
8-BGP_RT_Background
                AS path: 65412 I
                Accepted Multipath
                Localpref: 100
                Router ID: 172.31.100.1
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 577
                Address: 0xc0f1b68
                Next-hop reference count: 9
                Source: 172.22.123.2
                Next hop: 172.22.123.2 via ge-0/0/1.0, selected
                Session Id: 0x188
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65001 Peer AS: 65412
                Age: 9:40
                Validation State: unverified
                Task: BGP_65412.172.22.123.2
                AS path: 65412 I
                Accepted MultipathContrib
                Localpref: 100
                Router ID: 172.31.101.1
```

Question: The active BGP route for 30.30.0.0/24 is marked with an *. How many next hops does the active route have installed?

_____

Answer: The 30.30/24 prefix now has two next hops.

Step 3.8

Use the **`run show route forwarding-table destination 30.30.0.0/24`** command to view the packet forwarding table.

```
[edit protocols bgp]
lab@mxA# run show route forwarding-table destination 30.30.0.0/24
Routing table: default.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
30.30.0.0/24       user     0 172.22.123.2       ucst      577     7 ge-0/0/1.0

Routing table: __pfe_private__.inet
Internet:
```

```
Enabled protocols: Bridging,
Destination          Type RtRef Next hop           Type Index    NhRef Netif
default              perm     0                    dscd      514     2

Routing table: __juniper_services__.inet
Internet:
Enabled protocols: Bridging,
Destination          Type RtRef Next hop           Type Index    NhRef Netif
default              perm     0                    dscd      523     2

Routing table: __master.anon__.inet
Internet:
Enabled protocols: Bridging, Dual VLAN,
Destination          Type RtRef Next hop           Type Index    NhRef Netif
default              perm     0                    rjct      538     1
```

Question: Are the two next hops to P1 and P2 installed in the forwarding table?

_____

_____

Answer: No. Only one next hop is installed in the forwarding table. The default forwarding behavior for two or more equal cost next-hops is to randomly pick one to be installed in the forwarding table. In this example, the route is using the P2 interface for the 30.30.0.0/24 route. Your results might vary.

Step 3.9

Navigate to the `[edit policy-options policy-statement pfe-load-balance]` hierarchy. Create a policy named *`pfe-load-balance`* that only load balances the 30.30/22 routes being received from the P1 and P2 routers.

```
[edit protocols bgp]
lab@mxA# top edit policy-options policy-statement pfe-load-balance

[edit policy-options policy-statement pfe-load-balance]
lab@mxA# set term 1 from protocol bgp

[edit policy-options policy-statement pfe-load-balance]
lab@mxA# set term 1 from route-filter 30.30/22 longer

[edit policy-options policy-statement pfe-load-balance]
lab@mxA# set term 1 then load-balance per-packet

[edit policy-options policy-statement pfe-load-balance]
lab@mxA# show
term 1 {
    from {
        protocol bgp;
        route-filter 30.30.0.0/22 longer;
    }
    then {
        load-balance per-packet;
```

```
    }
}

[edit policy-options policy-statement pfe-load-balance]
lab@mxA#
```

Step 3.10

After configuring the *pfe-load-balance* policy, apply it as an export policy under the `[edit routing-options forwarding-table]` hierarchy. Commit the changes when completed.

```
[edit policy-options policy-statement pfe-load-balance]
lab@mxA# top edit routing-options forwarding-table

[edit routing-options forwarding-table]
lab@mxA# set export pfe-load-balance

[edit routing-options forwarding-table]
lab@mxA# commit
commit complete

[edit routing-options forwarding-table]
lab@mxA#
```

Step 3.11

Use the **`run show route forwarding-table destination 30.30/24`** command to verify that the forwarding table now has two next-hop interfaces for the 30.30/24 route.

```
[edit routing-options forwarding-table]
lab@mxA# run show route forwarding-table destination 30.30/24
Routing table: default.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
30.30.0.0/24       user     0                    ulst  1048575     5
                              172.22.123.2       ucst      577     4 ge-0/0/1.0
                              172.22.121.2       ucst      582     5 ge-0/0/3.0

Routing table: __pfe_private__.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      514     2

Routing table: __juniper_services__.inet
Internet:
Enabled protocols: Bridging,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      523     2

Routing table: __master.anon__.inet
Internet:
Enabled protocols: Bridging, Dual VLAN,
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    rjct      538     1
```

Question: Is the forwarding table using both next-hop interfaces to reach the 30.30/24 route?

_____

Answer: Yes, the forwarding table now has two next-hop interfaces for the 30.30/24 route, one to P1 and the other to P2.

Part 4: Configuring the EBGP Session with the P3 Router

In this lab part, you configure the R3-1 router to EBGP peer with the P3 router. The peering is between the loopback interfaces of R3-1 and P3. EBGP loopback peering is a three step process. Because the external interfaces are not participating in the IGP, the first step requires a static route to the P3 loopback address. The second step requires configuring the local loopback as the source address of the BGP messages sent to the P3 router. Finally, the BGP `multihop` command is configured to override the physical connection requirement (i.e. default TTL of EBGP messages is 1) normally imposed of EBGP sessions.

Step 4.1

Return to the SSH session that you have open to the mxB router.

From the SSH session with the mxB router, navigate to the `[edit routing-options]` hierarchy. Configure a static route to P3's loopback address (172.31.102.1) with a next hop of 172.22.125.2. Use the `no-readvertise` setting to ensure that the route can not be redistributed into other protocols and commit the configuration when completed.

```
[edit protocols bgp]
lab@mxB:R3-1# top edit routing-options

[edit routing-options]
lab@mxB:R3-1# show
static {
    route 172.16.1.128/26 reject;
    route 172.16.1.192/26 reject;
}
autonomous-system 65001;

[edit routing-options]
lab@mxB:R3-1# set static route 172.31.102.1 next-hop 172.22.125.2

[edit routing-options]
lab@mxB:R3-1# set static route 172.31.102.1 no-readvertise

[edit routing-options]
lab@mxB:R3-1# commit
commit complete

[edit routing-options]
lab@mxB:R3-1#
```

Step 4.2

Use the **`run ping 172.31.102.1 source 172.16.1.2 count 5`** command to verify that connectivity between the loopback addresses has been established.

```
[edit routing-options]
lab@mxB:R3-1# run ping 172.31.102.1 source 172.16.1.2 count 5
PING 172.31.102.1 (172.31.102.1): 56 data bytes
64 bytes from 172.31.102.1: icmp_seq=0 ttl=64 time=1.854 ms
64 bytes from 172.31.102.1: icmp_seq=1 ttl=64 time=1.651 ms
64 bytes from 172.31.102.1: icmp_seq=2 ttl=64 time=1.280 ms
64 bytes from 172.31.102.1: icmp_seq=3 ttl=64 time=86.709 ms
64 bytes from 172.31.102.1: icmp_seq=4 ttl=64 time=1.635 ms

--- 172.31.102.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.280/18.626/86.709/34.042 ms
```

Question: Is the ping test successful?

_____

Answer: The ping test should be successful. If it is not successful, verify your configuration and consult your instructor.

Step 4.3

Navigate to the `[edit protocols bgp]` hierarchy. Configure a BGP group named *P3*. Configure the P3 loopback address as the peer and the R3-1 loopback address as the local-address. The peer-as is *65020*. Commit the configuration when completed.

```
[edit routing-options]
lab@mxB:R3-1# top edit protocols bgp

[edit protocols bgp]
lab@mxB:R3-1# set group P3 type external

[edit protocols bgp]
lab@mxB:R3-1# set group P3 local-address 172.16.1.2

[edit protocols bgp]
lab@mxB:R3-1# set group P3 neighbor 172.31.102.1

[edit protocols bgp]
lab@mxB:R3-1# set group P3 peer-as 65020

[edit protocols bgp]
lab@mxB:R3-1# commit
commit complete

[edit protocols bgp]
lab@mxB:R3-1#
```

Step 4.4

Check the state of the EBGP session using the **`run show bgp summary`** command.

```
[edit protocols bgp]
lab@mxB:R3-1# run show bgp summary
Threading mode: BGP I/O
```

```
Groups: 2 Peers: 2 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     7          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.1            65001         50         48       0       0       20:43 Establ
  inet.0: 2/7/7/0
172.31.102.1          65020          0          0       0       0           8 Idle
```

Question: What is the state of the P3 peering session?

_____

Answer: The P3 peering session is in `Idle` state. All EBGP peering sessions must be peered with the physical interface or a TCP session will not be established.

Step 4.5

To relax the EBGP requirement of physical interface peering and make it possible to EBGP peer between loopback addresses, apply the `multihop` statement to the *P3* BGP group. Commit the change when completed.

```
[edit protocols bgp]
lab@mxB:R3-1# set group P3 multihop

[edit protocols bgp]
lab@mxB:R3-1# commit
commit complete

[edit protocols bgp]
lab@mxB:R3-1#
```

Step 4.6

Check the status again of the P3 session with the **`run show bgp summary`** command.

```
[edit protocols bgp]
lab@mxB:R3-1# run show bgp summary
Threading mode: BGP I/O
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                    13          8          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.1            65001         69         70       0       0       29:21 Establ
  inet.0: 2/7/7/0
172.31.102.1          65020          5          3       0       0           3 Establ
  inet.0: 6/6/6/0
```

Question: What is the state of the P3 peering session after the multihop command is configured?

_____

Answer: The P3 peering session should now be established. If the session is not established, check your configuration or consult your instructor.

Step 4.7

Now that the P3 peering session is established, use the **`run show route receive-protocol bgp 172.31.102.1`** command to view the routes being received from the P3 router.

```
[edit protocols bgp]
lab@mxB:R3-1# run show route receive-protocol bgp 172.31.102.1

inet.0: 23 destinations, 23 routes (18 active, 0 holddown, 5 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 40.40.0.0/24            172.31.102.1                            65020 I
* 40.40.1.0/24            172.31.102.1                            65020 I
* 40.40.2.0/24            172.31.102.1                            65020 I
* 40.40.3.0/24            172.31.102.1                            65020 I
* 172.16.2.0/26           172.31.102.1                            65020 65002 I
* 172.16.2.64/26          172.31.102.1                            65020 65002 I

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Are routes being received from the P3 peering session?

_____

Answer: Yes. Six routes are being received from the P3 router.

Part 5: Summarizing the Internal Routes to the Peer Routers

In this lab part, you must create an aggregate route for the internal redistributed static routes. Use 172.16.1.0/24 as the aggregate route to use for both the mxA and R3-1 routers. Because the IBGP routes are advertised to the EBGP peers by default, an export EBGP routing policy is required to advertise the aggregate route and suppress the specific routes.

Step 5.1

Navigate to the `[edit routing-options]` hierarchy. Configure the 172.16.1.0/24 aggregate route as discussed above.

```
[edit protocols bgp]
lab@mxB:R3-1# top edit routing-options

[edit routing-options]
lab@mxB:R3-1# set aggregate route 172.16.1.0/24

[edit routing-options]
lab@mxB:R3-1#
```

Step 5.2

Navigate to the `[edit policy-options policy-statement export-aggregate]` hierarchy. Create a routing policy named *`export-aggregate`* that advertises the aggregate route and suppresses the more specific routes.

```
[edit routing-options]
lab@mxB:R3-1# top edit policy-options policy-statement export-aggregate

[edit policy-options policy-statement export-aggregate]
lab@mxB:R3-1# set term 1 from protocol aggregate

[edit policy-options policy-statement export-aggregate]
lab@mxB:R3-1# set term 1 from route-filter 172.16.1.0/24 exact

[edit policy-options policy-statement export-aggregate]
lab@mxB:R3-1# set term 1 then accept

[edit policy-options policy-statement export-aggregate]
lab@mxB:R3-1# set term 2 from route-filter 172.16.1.0/24 longer

[edit policy-options policy-statement export-aggregate]
lab@mxB:R3-1# set term 2 then reject

[edit policy-options policy-statement export-aggregate]
lab@mxB:R3-1# show
term 1 {
    from {
        protocol aggregate;
        route-filter 172.16.1.0/24 exact;
    }
    then accept;
}
term 2 {
    from {
        route-filter 172.16.1.0/24 longer;
    }
    then reject;
}

[edit policy-options policy-statement export-aggregate]
lab@mxB:R3-1#
```

Step 5.3

Navigate to the `[edit protocols bgp]` hierarchy. Apply the *`export-aggregate`* policy as an export policy for the *`P3`* group. Commit the configuration when complete.

```
[edit policy-options policy-statement export-aggregate]
lab@mxB:R3-1# top edit protocols bgp

[edit protocols bgp]
lab@mxB:R3-1# set group P3 export export-aggregate

[edit protocols bgp]
lab@mxB:R3-1# commit
```

```
commit complete

[edit protocols bgp]
lab@mxB:R3-1#
```

Step 5.4

Use the **`run show route advertising-protocol bgp 172.31.102.1`** command to view the routes being advertised to the P3 router.

```
[edit protocols bgp]
lab@mxB:R3-1# run show route advertising-protocol bgp 172.31.102.1

inet.0: 24 destinations, 24 routes (19 active, 0 holddown, 5 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.0/24           Self                                    I
```

Question: Is the aggregate route being advertised to the P3 router?

_____

_____

Answer: Yes, the aggregate route is being advertised.

Step 5.5

Return to the SSH session that you have open to the mxA router.

From the SSH session with the mxA router, navigate to the `[edit routing-options]` hierarchy. Configure the 172.16.1.0/24 aggregate route as discussed previously.

```
[edit routing-options forwarding-table]
lab@mxA# up

[edit routing-options]
lab@mxA# set aggregate route 172.16.1.0/24

[edit routing-options]
lab@mxA#
```

Step 5.6

Navigate to the `[edit policy-options policy-statement export-aggregate]` hierarchy. Create a routing policy named *`export-aggregate`* that advertises the aggregate route and suppresses the more specific routes.

```
[edit routing-options]
lab@mxA# top edit policy-options policy-statement export-aggregate

[edit policy-options policy-statement export-aggregate]
lab@mxA# set term 1 from protocol aggregate

[edit policy-options policy-statement export-aggregate]
lab@mxA# set term 1 from route-filter 172.16.1.0/24 exact

[edit policy-options policy-statement export-aggregate]
lab@mxA# set term 1 then accept
```

```
[edit policy-options policy-statement export-aggregate]
lab@mxA# set term 2 from route-filter 172.16.1.0/24 longer

[edit policy-options policy-statement export-aggregate]
lab@mxA# set term 2 then reject

[edit policy-options policy-statement export-aggregate]
lab@mxA# show
term 1 {
    from {
        protocol aggregate;
        route-filter 172.16.1.0/24 exact;
    }
    then accept;
}
term 2 {
    from {
        route-filter 172.16.1.0/24 longer;
    }
    then reject;
}

[edit policy-options policy-statement export-aggregate]
lab@mxA#
```

Step 5.7

Navigate to the `[edit protocols bgp group P1-P2]` hierarchy. Apply the *`export-aggregate`* policy as an export policy for the P1-P2 group. Commit the configuration and return to operational mode when complete.

```
[edit policy-options policy-statement export-aggregate]
lab@mxA# top edit protocols bgp group P1-P2

[edit protocols bgp group P1-P2]
lab@mxA# set export export-aggregate

[edit protocols bgp group P1-P2]
lab@mxA# commit and-quit
commit complete
Exiting configuration mode

lab@mxA>
```

Step 5.8

Use the **`show route advertising-protocol bgp 172.22.121.1`** and **`show route advertising-protocol bgp 172.22.123.2`** commands to view the routes being advertised to the P1 and P2 routers.

```
lab@mxA> show route advertising-protocol bgp 172.22.121.2

inet.0: 27 destinations, 32 routes (21 active, 0 holddown, 6 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.0/24           Self                                    I
```

```
lab@mxA> show route advertising-protocol bgp 172.22.123.2

inet.0: 27 destinations, 32 routes (21 active, 0 holddown, 6 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.0/24           Self                                    I
```

Question: Is the aggregate route being advertised to the P1 and P2 routers?

_____

_____

Answer: Yes, the aggregate route is being advertised.

Step 5.9

Log out of your assigned devices using the **exit** command.

```
lab@mxA> exit

mxA (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

Network Management Diagram

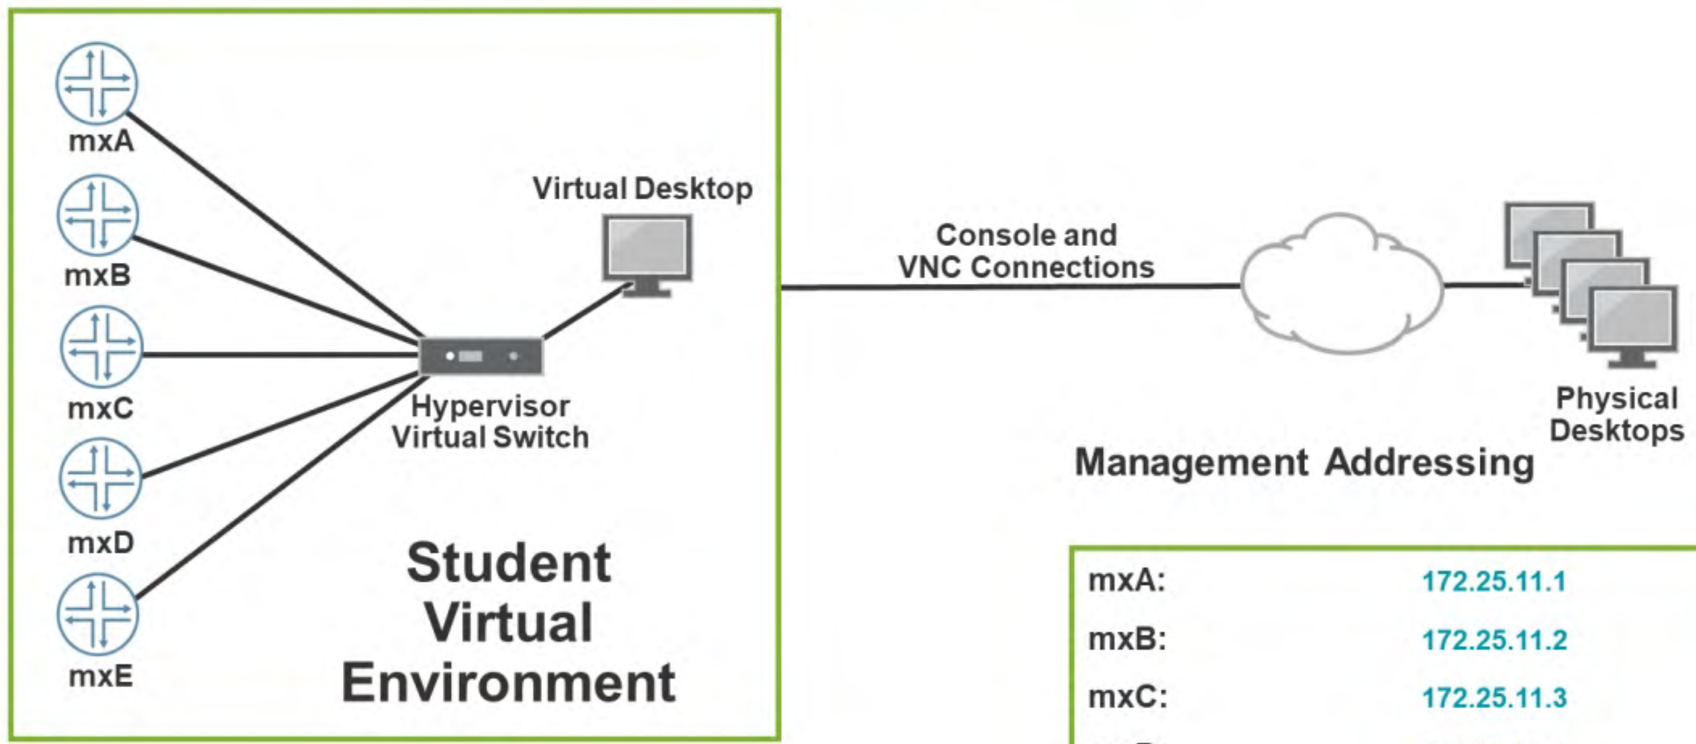

Management Addressing

mxA:	172.25.11.1
mxB:	172.25.11.2
mxC:	172.25.11.3
mxD:	172.25.11.4
mxE:	172.25.11.5
Virtual Desktop:	172.25.11.254

Note: Your instructor will provide address and access information.

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab Network Diagram: Configuring BGP

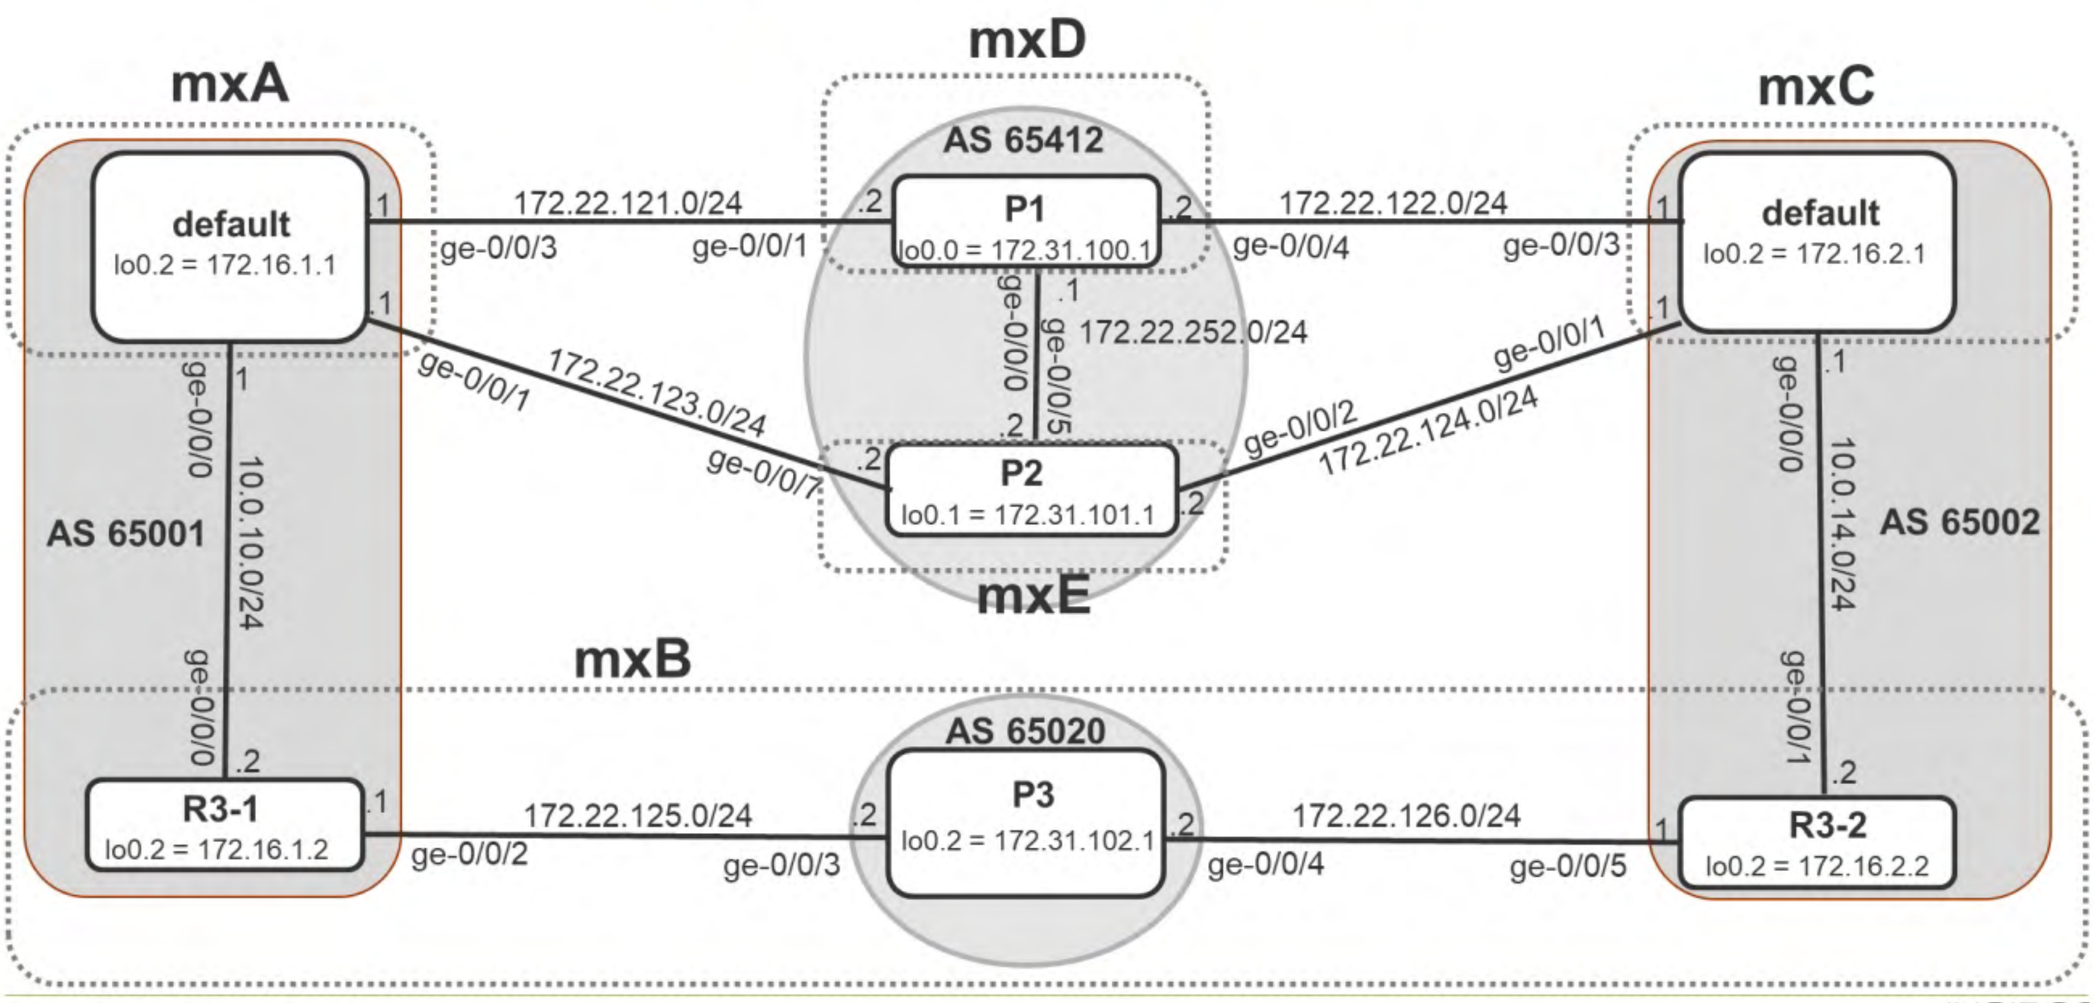

© 2020 Juniper Networks, Inc. All Rights Reserved.

Education Services

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1110 PZ SCHIPHOL-RIJK
Amsterdam, Netherlands
Phone: 31.0.207.125.700
Fax: 31.0.207.125.701

Copyright 2020
Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, services marks, registered marks, or registered services marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Engineering
Simplicity

EDU-JUN-AJSPR, Revision V19A