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 2 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 10: BGP Attributes: Part 1 10-1
Part 1: Load Starting Configuration 10-2
Part 2: Repairing Unusable Routes 10-30
Part 3: Modifying the Local-Preference Attribute 10-38
Part 4: Modifying the AS Path Attribute 10-46

Lab 11: BGP Attributes: Part 2 11-1
Part 2: Modifying the Origin Attribute 11-31
Part 3: Configuring the MED Attribute 11-38
Part 4: Configuring BGP Communities 11-43

Lab 12: Scaling BGP 12-1
Part 1: Load Starting Configuration 12-2
Part 2: Configuring Route Reflection 12-31
Part 3: Configuring Confederations 12-44

Lab 13: BGP FlowSpec 13-1
Part 1: Load Starting Configuration 13-2
Part 2: Configuring Customer-based FlowSpec 13-4
Part 3: Setting up the NOC to perform DDoS Attack Remediation 13-18

Lab 14: Troubleshooting BGP 14-1
Part 1: Load Starting Configuration 14-2
Part 2: Troubleshooting BGP 14-4

Lab 15: Troubleshooting Routing Policy 15-1
Part 1: Load Starting Configuration 15-2
Part 2: Troubleshooting Policies 15-4

Lab 16: BGP Route Damping 16-1
Part 1: Load Starting Configuration 16-2
Part 2: Configuring BGP Damping 16-32
Part 3: Modifying the BGP Damping Parameters 16-40

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

BGP Attributes: Part 1

Overview

In this lab, you use the lab diagram titled “Lab Network Diagram: BGP Attributes - Part 1” to repair unusable routes and influence the BGP route selection process. This lab utilizes pre-established internal Border Gateway Protocol (IBGP) and EBGP peering sessions in which certain peers are receiving "hidden" or unusable routes. The Local-Preference attribute is used in this lab to define a preferred exit point out for your AS for routes being received from AS 65412. Next, routes are advertised to the P1 and P2 routers using AS-path attributes.

By completing this lab you perform the following tasks:

- Repair the unusable routes.
- Influence traffic leaving the AS using the Local Preference attribute.
- Influence traffic entering the AS using the AS-path attribute.
- Use policy with AS Path regex expressions.

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/lab10-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: Sat Nov 30 23:52:44 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/lab10-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: 2019-12-27 12:08:29 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;
            }
        }
        bgp {
            group ibgp {
                type internal;
                local-address 172.16.1.2;
                export redistribute-statics;
                neighbor 172.16.1.1;
            }
            group P3 {
                type external;
                multihop;
                local-address 172.16.1.2;
                export export-aggregate;
                peer-as 65020;
                neighbor 172.31.102.1;
            }
        }
    }
    policy-options {
        policy-statement export-aggregate {
            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;
            }
        }
        policy-statement redistribute-statics {
            term 1 {
                from protocol static;
                then accept;
            }
        }
    }
    routing-options {
        static {
            route 172.31.102.1/32 {
                next-hop 172.22.125.2;
                no-readvertise;
            }
        }
        autonomous-system 65001;
        aggregate {
            route 172.16.1.0/24;
        }
    }
```

```
}
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 next-hop-self ];
                neighbor 172.16.2.1;
            }
            group P3 {
                type external;
                multihop;
                local-address 172.16.2.2;
                export export-aggregate;
                neighbor 172.31.102.1 {
                    peer-as 65020;
                }
            }
        }
    }
    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 next-hop-self {
            term 1 {
                from {
                    protocol bgp;
                    route-type external;
                }
                then {
                    next-hop self;
                }
            }
        }
        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: Based on the configuration, which routers (logical systems) will mxB represent in this lab?

Answer: The mxB device will, using logical systems, represent the P3, R3-1, and R3-2 routers.

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 any 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 configured for an IBGP session with the mxA router and an EBGP session with the P3 router.

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: What routes will the R3-2 router be advertising to its EBGP peer?

_____

Answer: The R3-2 router is configured with an export policy called *export-aggregate* that is applied to its EBGP peer. This policy causes the R3-2 router to advertise the configured aggregate route of 172.16.2.0/24.

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
Threading mode: BGP I/O
Groups: 2 Peers: 2 Down peers: 0
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     7          6          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2            65001        686        665       0       0     4:58:01 Establ
  inet.0: 1/2/2/0
172.16.2.2            65002        782        777       0       0     5:46:26 Establ
  inet.0: 5/5/5/0
```

Question: What is the state of the P3 router's BGP neighbor relationships?

_____

Answer: The P3 router should have an established EBGP session with both the R3-1 and R3-2 routers.

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/lab10-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: Sat Nov 30 23:52:44 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/lab10-start.config
load complete
```

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

lab@mxC>
```

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: 2019-12-01 15:38:27 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 import-P1 {
            term 1 {
                from {
                    protocol bgp;
                    as-path partner-as;
                }
                then accept;
```

```
        }
        term 2 {
            then reject;
        }
    }
    policy-statement next-hop-self {
        term 1 {
            from {
                protocol bgp;
                route-type external;
            }
            then {
                next-hop self;
            }
        }
    }
    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;
        }
    }
    as-path partner-as ".* 65001";
    as-path internal-as "()";
}
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 next-hop-self ];
            neighbor 172.16.2.2;
        }
        group P1-P2 {
            type external;
            export export-aggregate;
            peer-as 65412;
            neighbor 172.22.122.2 {
                import import-P1;
            }
            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: What routes will the mxC router be advertising to its EBGP peers?

_____

_____

Answer: The mxC router is configured with an export policy called `export-aggregate` that is applied to its EBGP peers. This policy causes the mxC router to advertise the configured aggregate route of 172.16.2.0/24.

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
                    17         11          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.2.2            65002        766        766       0       0     5:39:41 Establ
  inet.0: 6/7/7/0 0/0/0/0
172.22.122.2          65412          0          0       0       0        2:04 Connect
172.22.124.2          65412          0          0       0       0        2:04 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. Because of the state of the lab, you may see an established session with both the P1 and P2 routers. If you do, there is no problem.

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/lab10-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: Sat Nov 30 23:52:44 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/lab10-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

lab@mxE> show configuration |no-more
## Last commit: 2019-12-27 12:09:22 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
                      12          6          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        7:56 Connect
172.22.124.1          65002        603        607       0       1     4:28:40 Establ
  inet.0: 6/6/6/0
172.31.100.1          65412          0          0       0       0        7:56 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. Depending on the state of the topology, you may see the BGP sessions with the P1 and mxA router in the established state. If so, this is not a problem.

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/lab10-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: Sat Nov 30 23:52:44 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/lab10-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: 2019-12-27 12:09:34 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
                      12          6          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        1:21 Connect
172.22.122.1          65002        593        597       0       1     4:24:04 Establ
  inet.0: 6/6/6/0
172.31.101.1          65412        604        601       0       2     4:24:04 Establ
  inet.0: 0/6/6/0
```

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. Depending on the state of the lab, you may see that the BGP session with the mxA router is in the established state. If so, this is not a problem.

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/lab10-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 Dec  1 11:01:25 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/lab10-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, OSPF, and BGP configuration.

```
lab@mxA> show configuration
## Last commit: 2019-12-27 12:08:40 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;
            }
        }
    }
}
policy-options {
    policy-statement export-aggregate {
        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;
            }
        }
        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 {
        forwarding-table {
            export pfe-load-balance;
        }
        autonomous-system 65001;
        aggregate {
            route 172.16.1.0/24;
        }
    }
    protocols {
        ospf {
            area 0.0.0.0 {
                interface lo0.0;
                interface ge-0/0/0.0;
            }
        }
        bgp {
            group ibgp {
                type internal;
                local-address 172.16.1.1;
                export redistribute-statics;
                neighbor 172.16.1.2;
            }
            group P1-P2 {
                type external;
                export [ redistribute-statics export-aggregate ];
                peer-as 65412;
                neighbor 172.22.121.2;
                neighbor 172.22.123.2;
            }
        }
    }
```

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: Which routers are mxA configured to peer with using BGP?

_____

_____

Answer: The mxA router is configured with an IBGP session with the R3-1 router and an EBGP session with the P1 and P2 routers.

Question: What routes will the mxA router be advertising to its EBGP peers?

_____

_____

Answer: The mxA router is configured with an export policy called *export-aggregate* that is applied to its EBGP peers. This policy causes the mxA router to advertise the configured aggregate route of 172.16.1.0/24.

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
Threading mode: BGP I/O
Groups: 2 Peers: 3 Down peers: 0
```

```
Table            Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      15          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2            65001        575        573       0       2     4:10:03 Establ
  inet.0: 0/5/5/0
172.22.121.2          65412        563        569       0       3     4:10:03 Establ
  inet.0: 5/5/5/0
172.22.123.2          65412        563        570       0       3     4:10:03 Establ
  inet.0: 0/5/5/0
```

Question: What is the status of the mxA router's BGP neighbor relationships?

_____

_____

Answer: The mxA router has established BGP sessions with the R3-1, P1, and P2 routers.

Part 2: Repairing Unusable Routes

In this lab part, you identify unusable routes using the **`show route hidden`** command. After analyzing the hidden routes, and discovering the reason they are unusable, you write an IBGP export policy to change the next-hop attribute. After applying the IBGP export policy, you verify that the routes are now active and usable.

Step 2.1

Use the **`show route table inet.0`** command to display the routing table.

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

inet.0: 21 destinations, 27 routes (17 active, 0 holddown, 5 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.10.0/24       *[Direct/0] 00:47:43
                    >  via ge-0/0/0.0
10.0.10.1/32       *[Local/0] 00:47:43
                       Local via ge-0/0/0.0
30.30.0.0/24       *[BGP/170] 00:47:37, localpref 100
                      AS path: 65412 I, validation-state: unverified
                    >  to 172.22.123.2 via ge-0/0/1.0
                    [BGP/170] 00:05:53, localpref 100
                      AS path: 65412 I, validation-state: unverified
                    >  to 172.22.121.2 via ge-0/0/3.0
30.30.1.0/24       *[BGP/170] 00:47:37, localpref 100
                      AS path: 65412 I, validation-state: unverified
                    >  to 172.22.123.2 via ge-0/0/1.0
                    [BGP/170] 00:05:53, localpref 100
                      AS path: 65412 I, validation-state: unverified
                    >  to 172.22.121.2 via ge-0/0/3.0
```

```
30.30.2.0/24        *[BGP/170] 00:47:37, localpref 100
                      AS path: 65412 I, validation-state: unverified
                     >  to 172.22.123.2 via ge-0/0/1.0
                    [BGP/170] 00:05:53, localpref 100
                      AS path: 65412 I, validation-state: unverified
                     >  to 172.22.121.2 via ge-0/0/3.0
30.30.3.0/24        *[BGP/170] 00:47:37, localpref 100
                      AS path: 65412 I, validation-state: unverified
                     >  to 172.22.123.2 via ge-0/0/1.0
                    [BGP/170] 00:05:53, localpref 100
                      AS path: 65412 I, validation-state: unverified
                     >  to 172.22.121.2 via ge-0/0/3.0
172.16.1.0/24       *[Aggregate/130] 00:47:43
                       Reject
172.16.1.1/32       *[Direct/0] 00:47:43
                     >  via lo0.0
172.16.1.2/32       *[OSPF/10] 00:47:33, metric 1
                     >  to 10.0.10.2 via ge-0/0/0.0
172.16.2.0/24       *[BGP/170] 00:19:59, localpref 100
                      AS path: 65412 65002 I, validation-state: unverified
                     >  to 172.22.121.2 via ge-0/0/3.0
                    [BGP/170] 00:47:37, localpref 100
                      AS path: 65412 65002 I, validation-state: unverified
                     >  to 172.22.123.2 via ge-0/0/1.0
172.22.121.0/24     *[Direct/0] 00:47:43
                     >  via ge-0/0/3.0
172.22.121.1/32     *[Local/0] 00:47:43
                       Local via ge-0/0/3.0
172.22.123.0/24     *[Direct/0] 00:47:43
                     >  via ge-0/0/1.0
172.22.123.1/32     *[Local/0] 00:47:43
                       Local via ge-0/0/1.0
172.25.11.0/24      *[Direct/0] 5w4d 21:36:21
                     >  via fxp0.0
172.25.11.1/32      *[Local/0] 5w4d 21:36:21
                       Local via fxp0.0
224.0.0.5/32        *[OSPF/10] 00:47:43, metric 1
                       MultiRecv
```

Question: Are there currently any hidden routes learned by the mxA router?

_____

Answer: Yes. The output shows that there are currently 5 hidden routes learned by the mxA router.

Step 2.2

Use the **`show route hidden table inet.0`** command to identify the unusable routes in the default routing table.

```
lab@mxA> show route hidden table inet.0
```

```
inet.0: 23 destinations, 29 routes (19 active, 0 holddown, 5 hidden)
+ = Active Route, - = Last Active, * = Both

40.40.0.0/24           [BGP/170] 01:40:29, localpref 100, from 172.16.1.2
                         AS path: 65020 I, validation-state: unverified
                          Unusable
40.40.1.0/24           [BGP/170] 01:40:29, localpref 100, from 172.16.1.2
                         AS path: 65020 I, validation-state: unverified
                          Unusable
40.40.2.0/24           [BGP/170] 01:40:29, localpref 100, from 172.16.1.2
                         AS path: 65020 I, validation-state: unverified
                          Unusable
40.40.3.0/24           [BGP/170] 01:40:29, localpref 100, from 172.16.1.2
                         AS path: 65020 I, validation-state: unverified
                          Unusable
172.16.2.0/24          [BGP/170] 00:02:37, localpref 100, from 172.16.1.2
                         AS path: 65020 65002 I, validation-state: unverified
                          Unusable
```

Question: All of the hidden routes are BGP routes. Which BGP peer is advertising these routes?

_____

Answer: In the example output, all of the hidden routes are being advertised from the IBGP peer 172.16.1.2.

Step 2.3

Use the **show route 40.40.0.0/24 hidden extensive** command to display more information about the hidden routes.

```
lab@mxA> show route 40.40.0.0/24 hidden extensive

inet.0: 25 destinations, 31 routes (21 active, 0 holddown, 5 hidden)
40.40.0.0/24 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Next hop type: Unusable, Next hop index: 0
                Address: 0x9f2fec4
                Next-hop reference count: 5
                State: <Hidden Int Ext>
                Local AS: 65001 Peer AS: 65001
                Age: 8:04
                Validation State: unverified
                Task: BGP_65001.172.16.1.2
                AS path: 65020 I
                Accepted
                Localpref: 100
                Router ID: 172.16.1.2
                Indirect next hops: 1
                        Protocol next hop: 172.31.102.1
                        Indirect next hop: 0x0 - INH Session ID: 0x0
```

Question: What is the BGP next-hop attribute for this route?

_____

Answer: In this example, the `Protocol next hop` is 172.31.102.1.

Question: Because IBGP peers do not modify any attributes learned from EBGP peers by default, which EBGP peer must have advertised this route?

_____

Answer: The P3 router advertised this route.

Step 2.4

In the BGP route selection process, the BGP next-hop attribute must be resolved in the default routing table (inet.0). Use the **`show route 172.31.102.1 table inet.0`** command to verify that the protocol next-hop can be resolved.

```
lab@mxA> show route 172.31.102.1 table inet.0

lab@mxA>
```

Question: Can the BGP next-hop address be resolved in the default routing table?

_____

Answer: No, the BGP next-hop address cannot be resolved in the default routing table and therefore the route is marked as unusable.

Step 2.5

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 2.6

An IBGP export policy needs to be created to modify the next hop attribute that is being advertised by R3-1 to something that can be resolved in the default routing table.

Enter configuration mode and navigate to the `[edit policy-options policy-statement next-hop-self]` hierarchy. In the policy, change the next-hop attribute to the loopback address of the IBGP advertising peer. Ensure that you only change the next-hop attribute if the route is a EBGP route.

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

```
[edit]
lab@mxB:R3-1# edit policy-options policy-statement next-hop-self

[edit policy-options policy-statement next-hop-self]
lab@mxB:R3-1# set term 1 from protocol bgp

[edit policy-options policy-statement next-hop-self]
lab@mxB:R3-1# set term 1 from route-type external

[edit policy-options policy-statement next-hop-self]
lab@mxB:R3-1# set term 1 then next-hop self

[edit policy-options policy-statement next-hop-self]
lab@mxB:R3-1# show
term 1 {
    from {
        protocol bgp;
        route-type external;
    }
    then {
        next-hop self;
    }
}
```

Step 2.7

Navigate to the `[edit protocols bgp]` hierarchy and apply the *`next-hop-self`* export policy in the *`ibgp`* group. Commit the changes when completed.

```
[edit policy-options policy-statement next-hop-self]
lab@mxB:R3-1# top edit protocols bgp

[edit protocols bgp]
lab@mxB:R3-1# set group ibgp export next-hop-self

[edit protocols bgp]
lab@mxB:R3-1# commit
commit complete

[edit protocols bgp]
lab@mxB:R3-1#
```

Step 2.8

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

From the SSH session with the mxA router, issue the **`show route hidden`** command to verify that all of the hidden routes are gone from both routing tables.

```
lab@mxA> show route hidden

inet.0: 21 destinations, 27 routes (21 active, 0 holddown, 0 hidden)

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Are all of the hidden routes gone from all of the routing tables?

_____

Answer: Yes. By applying the *next-hop-self* policy on the routes being advertised by the R3-1 router, the hidden route problem on mxA has been solved.

Step 2.9

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

From the SSH session with the mxB router, issue the **run show route hidden** command to see if there are any hidden routes learned by the R3-1 router.

```
[edit protocols bgp]
lab@mxB:R3-1# run show route hidden

inet.0: 22 destinations, 23 routes (18 active, 0 holddown, 5 hidden)
+ = Active Route, - = Last Active, * = Both

30.30.0.0/24           [BGP/170] 00:22:19, localpref 100, from 172.16.1.1
                         AS path: 65412 I, validation-state: unverified
                         Unusable
30.30.1.0/24           [BGP/170] 00:22:19, localpref 100, from 172.16.1.1
                         AS path: 65412 I, validation-state: unverified
                         Unusable
30.30.2.0/24           [BGP/170] 00:22:19, localpref 100, from 172.16.1.1
                         AS path: 65412 I, validation-state: unverified
                         Unusable
30.30.3.0/24           [BGP/170] 00:22:19, localpref 100, from 172.16.1.1
                         AS path: 65412 I, validation-state: unverified
                         Unusable
172.16.2.0/24          [BGP/170] 00:22:19, localpref 100, from 172.16.1.1
                         AS path: 65412 65002 I, validation-state: unverified
                         Unusable

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Are there currently any hidden routes learned by the R3-1 router?

_____

Answer: Yes. The output shows that there are currently 5 hidden routes learned by the mxA router.

Question: All of the hidden routes are BGP routes. Which BGP peer is advertising these routes?

_____

Answer: In the example output, all of the hidden routes are being advertised from the IBGP peer 172.16.1.1.

Step 2.10

Use the **`run show route 30.30.0.0/24 hidden extensive`** command to display more information about the hidden routes.

```
[edit protocols bgp]
lab@mxB:R3-1# run show route 30.30.0.0/24 hidden extensive

inet.0: 22 destinations, 23 routes (18 active, 0 holddown, 5 hidden)
30.30.0.0/24 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Next hop type: Unusable, Next hop index: 0
                Address: 0x9f2fec4
                Next-hop reference count: 5
                State: <Hidden Int Ext>
                Local AS: 65001 Peer AS: 65001
                Age: 28:35
                Validation State: unverified
                Task: BGP_65001.172.16.1.1
                AS path: 65412 I
                Accepted
                Localpref: 100
                Router ID: 172.16.1.1
                Indirect next hops: 1
                        Protocol next hop: 172.22.123.2
                        Indirect next hop: 0x0 - INH Session ID: 0x0
```

Question: What is the BGP next-hop attribute for this route?

_____

Answer: In this example, the `Protocol next hop` is 172.22.123.2. **Your next hop may differ.** If the next hop is different remember the IP address as you will need it in the next step.

Question: Because IBGP peers do not modify any attributes learned from EBGP peers by default, which EBGP peer must have advertised this route?

_____

Answer: in the example, the P2 router advertised this route. In your topology, P1 may have been the router that advertised the route. If so, there is no problem.

Step 2.11

In the BGP route selection process, the BGP next-hop attribute must be resolved in the default routing table (inet.0). Use the **`run show route <IP from Step 2.10> table inet.0`** command to verify that the protocol next hop can be resolved.

```
[edit protocols bgp]
lab@mxB:R3-1# run show route <IP from Step 2.10> table inet.0

[edit protocols bgp]
lab@mxB:R3-1#
```

Question: Can the BGP next-hop address be resolved in the default routing table?

_____

Answer: No, the BGP next-hop address cannot be resolved in the default routing table and therefore the route is marked as unusable.

Step 2.12

An IBGP export policy needs to be created to modify the next hop attribute that is being advertised by mxA router to something that can be resolved in the default routing table.

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 policy-options policy-statement next-hop-self]` hierarchy. In the policy, change the next-hop attribute to the loopback address of the IBGP advertising peer. Ensure that you only change the next-hop attribute if the route is a EBGP route.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit policy-options policy-statement next-hop-self

[edit policy-options policy-statement next-hop-self]
lab@mxA# set term 1 from protocol bgp

[edit policy-options policy-statement next-hop-self]
lab@mxA# set term 1 from route-type external

[edit policy-options policy-statement next-hop-self]
lab@mxA# set term 1 then next-hop self

[edit policy-options policy-statement next-hop-self]
lab@mxA# show
term 1 {
    from {
        protocol bgp;
        route-type external;
    }
    then {
        next-hop self;
    }
}

[edit policy-options policy-statement next-hop-self]
lab@mxA#
```

Step 2.13

Navigate to the `[edit protocols bgp]` hierarchy and apply the *`next-hop-self`* export policy in the *`ibgp`* group. Commit the changes when completed.

```
[edit policy-options policy-statement next-hop-self]
lab@mxA# top edit protocols bgp

[edit protocols bgp]
lab@mxA# set group ibgp export next-hop-self

[edit protocols bgp]
lab@mxA# commit and-quit
commit complete

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

Step 2.14

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

From the SSH session with the mxB router, issue the **`run show route hidden`** command to verify that all of the hidden routes are gone from both routing tables.

```
[edit protocols bgp]
lab@mxB:R3-1# run show route hidden

inet.0: 18 destinations, 19routes (18 active, 0 holddown, 0 hidden)

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

[edit protocols bgp]
lab@mxB:R3-1#
```

Question: Are all of the hidden routes gone from all of the routing tables?

_____

Answer: Yes. By applying the *`next-hop-self`* policy on the routes being advertised by the mxA router, the hidden route problem on R3-1 has been solved.

Part 3: Modifying the Local-Preference Attribute

In this lab part, you will use the local-preference attribute to change the routing behavior within AS 65001.

Step 3.1

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

From the SSH session with the mxA router, use the **`show route 172.16.2.0/24 detail`** to display detailed information about the summary route being advertised from AS 65002.

```
lab@mxA> show route 172.16.2.0/24 detail
inet.0: 23 destinations, 29 routes (23 active, 0 holddown, 0 hidden)
```

```
172.16.2.0/24 (3 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 591
                Address: 0xb39d6f0
                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: 0x14c
                State: <Active Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 5:42
                Validation State: unverified
                Task: BGP_65412.172.22.121.2
                Announcement bits (3): 0-KRT 5-BGP_RT_Background 6-Resolve tree 1
                AS path: 65412 65002 I
                Aggregator: 65002 172.16.2.1
                Accepted
                Localpref: 100
                Router ID: 172.31.100.1
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0xb39d630
                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>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65001 Peer AS: 65412
                Age: 5:38
                Validation State: unverified
                Task: BGP_65412.172.22.123.2
                AS path: 65412 65002 I
                Aggregator: 65002 172.16.2.1
                Accepted
                Localpref: 100
                Router ID: 172.31.101.1
        BGP     Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39d3f0
                Next-hop reference count: 15
                Source: 172.16.1.2
                Next hop type: Router, Next hop index: 590
                Next hop: 10.0.10.2 via ge-0/0/0.0, selected
                Session Id: 0x14d
                Protocol next hop: 172.16.1.2
                Indirect next hop: 0xb2b2860 1048574 INH Session ID: 0x14e
                State: <Int Ext>
                Inactive reason: Interior > Exterior > Exterior via Interior
                Local AS: 65001 Peer AS: 65001
                Age: 5:02       Metric2: 1
                Validation State: unverified
                Task: BGP_65001.172.16.1.2
                AS path: 65020 65002 I
                Aggregator: 65002 172.16.2.2
```

```
Accepted
Localpref: 100
Router ID: 172.16.1.2
```

Question: Three BGP routes have been received by the mxA router from the R3-1, P1 and P2 routers. Of these three BGP updates, which route is active?

_____

_____

Answer: In the above capture, the route received from the P1 router is the active route. Your results might differ.

Question: What is the local-preference for the three BGP routes?

_____

_____

Answer: The local-preference value for all three BGP routes is at the default 100.

Step 3.2

Enter configuration mode and navigate to the `[edit policy-options policy-statement import-localpref]` hierarchy. Configure the *`import-localpref`* policy to set the `local-preference` on the summary route being received from the P1 router to 110. Ensure that the `local-preference` is only changed on the P1 neighbor. Use the **`show`** command to display the policy.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit policy-options policy-statement import-localpref

[edit policy-options policy-statement import-localpref]
lab@mxA# set term 1 from protocol bgp


[edit policy-options policy-statement import-localpref]
lab@mxA# set term 1 from neighbor 172.22.121.2

[edit policy-options policy-statement import-localpref]
lab@mxA# set term 1 from route-filter 172.16.2.0/24 exact

[edit policy-options policy-statement import-localpref]
lab@mxA# set term 1 then local-preference 110

[edit policy-options policy-statement import-localpref]
lab@mxA# show
term 1 {
    from {
        protocol bgp;
```

```
        neighbor 172.22.121.2;
        route-filter 172.16.2.0/24 exact;
    }
    then {
        local-preference 110;
    }
}

[edit policy-options policy-statement import-localpref]
lab@mxA#
```

Step 3.3

Navigate to the `[edit protocols bgp group P1-P2]` hierarchy and apply the *`import-localpref`* policy as an import policy under the group. Commit the configuration when completed.

```
[edit policy-options policy-statement import-localpref]
lab@mxA# top edit protocols bgp group P1-P2

[edit protocols bgp group P1-P2]
lab@mxA# set import import-localpref

[edit protocols bgp group P1-P2]
lab@mxA# commit
commit complete

[edit protocols bgp group P1-P2]
lab@mxA# #
```

Step 3.4

Use the **`run show route 172.16.2.0/24 detail`** command to display detailed information about the AS 65002 summary route.

```
[edit protocols bgp group P1-P2]
lab@mxA# run show route 172.16.2.0/24 detail

inet.0: 23 destinations, 28 routes (23 active, 0 holddown, 0 hidden)
172.16.2.0/24 (2 entries, 1 announced)
        *BGP    Preference: 170/-111
                Next hop type: Router, Next hop index: 591
                Address: 0xb39d6f0
                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: 0x14c
                State: <Active Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 2:20
                Validation State: unverified
                Task: BGP_65412.172.22.121.2
                Announcement bits (3): 0-KRT 5-BGP_RT_Background 6-Resolve tree 1
                AS path: 65412 65002 I
                Aggregator: 65002 172.16.2.1
                Accepted
                Localpref: 110
```

```
                Router ID: 172.31.100.1
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0xb39d630
                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: <Ext>
                Inactive reason: Local Preference
                Local AS: 65001 Peer AS: 65412
                Age: 47:50
                Validation State: unverified
                Task: BGP_65412.172.22.123.2
                AS path: 65412 65002 I
                Aggregator: 65002 172.16.2.1
                Accepted
                Localpref: 100
                Router ID: 172.31.101.1
```

Question: Now, only two BGP routes are in the mxA router's routing table. What happened to the route from the R3-1 router?

_____

_____

Answer: The R3-1 router is no longer advertising the 172.16.2.0/24 route to the mxA router. This is because the local-preference attribute is evaluated by the route selection algorithm on the R3-1 router before the AS Path attribute which means that the summary route that the R3-1 router receives from the mxA router (i.e. the route advertised by P1) is preferred over the route received from the P3 router. Only preferred routes are advertised to BGP peers.

Question: What do you notice about the local-preference value on the P1 route?

_____

_____

Answer: It has changed to 110 due to the *import-localpref* policy.

Step 3.5

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

Use the **show route 172.16.2.0/24 detail** command to display detail information about the summary route from AS 65002.

```
[edit protocols bgp]
lab@mxB:R3-1# run show route 172.16.2.0/24 detail
```

```
inet.0: 23 destinations, 24 routes (23 active, 0 holddown, 0 hidden)
172.16.2.0/24 (2 entries, 1 announced)
        *BGP    Preference: 170/-111
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39cd90
                Next-hop reference count: 14
                Source: 172.16.1.1
                Next hop type: Router, Next hop index: 766
                Next hop: 10.0.10.1 via ge-0/0/0.0, selected
                Session Id: 0x17f
                Protocol next hop: 172.16.1.1
                Indirect next hop: 0xb2b30e0 1048579 INH Session ID: 0x180
                State: <Active Int Ext>
                Local AS: 65001 Peer AS: 65001
                Age: 12:15      Metric2: 1
                Validation State: unverified
                Task: BGP_65001.172.16.1.1
                Announcement bits (3): 3-KRT 5-BGP_RT_Background 6-Resolve tree 1
                AS path: 65412 65002 I
                Aggregator: 65002 172.16.2.1
                Accepted
                Localpref: 110
                Router ID: 172.16.1.1
        BGP     Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39cc70
                Next-hop reference count: 9
                Source: 172.31.102.1
                Next hop type: Router, Next hop index: 692
                Next hop: 172.22.125.2 via ge-0/0/2.0, selected
                Session Id: 0x174
                Protocol next hop: 172.31.102.1
                Indirect next hop: 0xb2b2ca0 1048574 INH Session ID: 0x176
                State: <Ext>
                Inactive reason: Local Preference
                Local AS: 65001 Peer AS: 65020
                Age: 1d 5:01:07         Metric2: 0
                Validation State: unverified
                Task: BGP_65020.172.31.102.1
                AS path: 65020 65002 I
                Aggregator: 65002 172.16.2.2
                Accepted
                Localpref: 100
                Router ID: 172.31.102.1
```

Question: Two BGP routes are in the R3-1 router's routing table for the AS 65002 summary route. Why is the route from the mxA router active over the route received from P3?

_____

_____

Answer: The local preference attribute is evaluated before the AS Path length in the BGP route selection process. Because the higher local preference is preferred, the BGP route from the default routing instance with a local preference of 110 is preferred over the P3 route with a local preference of 100.

Step 3.6

Navigate to the `[edit policy-options policy-statement import-P3]` hierarchy. Configure the *import-p3* policy to set the `local-preference` on the AS 65002 summary route being received from the P3 router to 120. Ensure that the `local-preference` is only changed on the P3 neighbor. Also, recall that the P3 neighbor is a multihop peer to the P3 loopback address. Use the **show** command to display the policy.

```
[edit protocols bgp]
lab@mxB:R3-1# top edit policy-options policy-statement import-p3

[edit policy-options policy-statement import-p3]
lab@mxB:R3-1# set term 1 from protocol bgp

[edit policy-options policy-statement import-p3]
lab@mxB:R3-1# set term 1 from neighbor 172.31.102.1

[edit policy-options policy-statement import-p3]
lab@mxB:R3-1# set term 1 from route-filter 172.16.2.0/24 exact

[edit policy-options policy-statement import-p3]
lab@mxB:R3-1# set term 1 then local-preference 120

[edit policy-options policy-statement import-p3]
lab@mxB:R3-1# show
term 1 {
    from {
        protocol bgp;
        neighbor 172.31.102.1;
        route-filter 172.16.2.0/24 exact;
    }
    then {
        local-preference 120;
    }
}

[edit policy-options policy-statement import-p3]
lab@mxB:R3-1#
```

Step 3.7

Navigate to the `[edit protocols bgp group P3]` hierarchy and apply the *`import-p3`* policy as an import policy under the group. Commit the configuration and exit to operational mode when completed.

```
[edit policy-options policy-statement import-p3]
lab@mxB:R3-1# top edit protocols bgp group P3

[edit protocols bgp group P3]
lab@mxB:R3-1# set import import-p3

[edit protocols bgp group P3]
lab@mxB:R3-1# commit
commit complete
Exiting configuration mode

[edit protocols bgp group P3]
lab@mxB:R3-1#
```

Step 3.8

Use the **`show route 172.16.2.0/24 detail`** command to display detail information about the AS 65002 summary route.

```
[edit protocols bgp group P3]
lab@mxB:R3-1# run show route 172.16.2.0/24 detail

inet.0: 23 destinations, 23 routes (23 active, 0 holddown, 0 hidden)
172.16.2.0/24 (1 entry, 1 announced)
        *BGP    Preference: 170/-121
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39cc70
                Next-hop reference count: 10
                Source: 172.31.102.1
                Next hop type: Router, Next hop index: 692
                Next hop: 172.22.125.2 via ge-0/0/2.0, selected
                Session Id: 0x174
                Protocol next hop: 172.31.102.1
                Indirect next hop: 0xb2b2ca0 1048574 INH Session ID: 0x176
                State: <Active Ext>
                Local AS: 65001 Peer AS: 65020
                Age: 9:04       Metric2: 0
                Validation State: unverified
                Task: BGP_65020.172.31.102.1
                Announcement bits (3): 3-KRT 5-BGP_RT_Background 6-Resolve tree 1
                AS path: 65020 65002 I
                Aggregator: 65002 172.16.2.2
                Accepted
                Localpref: 120
                Router ID: 172.31.102.1
```

Question: Did the local preference for the AS 65002 summary route change to 120?

_____

Answer: Yes, as the above capture shows, the local preference on this route is now 120.

Step 3.9

To set up the next part of the lab, remove the *import-p3* statement from the R3-1 configuration. Commit the configuration when completed.

```
[edit protocols bgp group P3]
lab@mxB:R3-1# delete import

[edit protocols bgp group P3]
lab@mxB:R3-1# commit
commit complete
Exiting configuration mode

lab@mxB:R3-1>
```

Part 4: Modifying the AS Path Attribute

In this lab part, you modify the AS Path attribute. The AS Path attribute is a mandatory well-known attribute that must be included in every BGP update. The attribute is modified as routes are advertised between EBGP peers. The AS number of the advertising peer is prepended to the beginning of the attribute before it is advertised to the peer. If a BGP update is received from a peer and the AS number of the receiving peer is somewhere in the attribute, the update is considered to have looped and discarded. The AS Path attribute is also used in the route selection process, the shortest AS path length is preferred.#

Step 4.1

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

From the SSH session with the mxC router, issue the **show route 172.16.1.0/24** command to display the path of your summary route from AS 65001.

```
lab@mxC> show route 172.16.1.0/24

inet.0: 25 destinations, 31 routes (25 active, 0 holddown, 4 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24       *[BGP/170] 00:42:55, localpref 100
                       AS path: 65412 65001 I, validation-state: unverified
                     >  to 172.22.124.2 via ge-0/0/1.0
                     [BGP/170] 00:42:55, localpref 100
                       AS path: 65412 65001 I, validation-state: unverified
                     >  to 172.22.122.2 via ge-0/0/3.0
                     [BGP/170] 00:43:31, localpref 100, from 172.16.2.2
                       AS path: 65020 65001 I, validation-state: unverified
                     >  to 10.0.14.2 via ge-0/0/0.0
```

Question: How many BGP paths exist for your summary route?

_____

_____

Answer: Three possible paths exist to the summary route through P1, P2, and P3.

Question: What is the AS Path for the three BGP paths?

_____

_____

Answer: In the example output, the AS Path in two of the BGP updates is `65412 65001`. The third is `65020 65001`.

Step 4.2

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 policy-options]` hierarchy. Copy the *`export-aggregate`* policy to a new policy named *`export-p3`* and display the new policy with the **`show policy-statement export-p3`** command.

```
[[edit protocols bgp group P3]
lab@mxB:R3-1# top edit policy-options

[edit policy-options]
lab@mxB:R3-1# copy policy-statement export-aggregate to policy-statement export-p3

[edit policy-options]
lab@mxB:R3-1# show policy-statement export-p3
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]
lab@mxB:R3-1#
```

Step 4.3

Navigate to the `[edit policy-options policy-statement export-p3]` hierarchy. Using the `as-path-prepend` option, insert AS 65002 into the AS Path.

```
[edit policy-options]
lab@mxB:R3-1# edit policy-statement export-p3
```

```
[edit policy-options policy-statement export-p3]
lab@mxB:R3-1# set term 1 then as-path-prepend 65002

[edit policy-options policy-statement export-p3]
lab@mxB:R3-1#
```

Step 4.4

Navigate to the `[edit protocols bgp]` hierarchy and apply the *`export-p3`* policy as an export policy under the P3 neighbor hierarchy. Commit the changes when completed.

```
[edit policy-options policy-statement export-p3]
lab@mxB:R3-1# top edit protocols bgp

[edit protocols bgp]
lab@mxB:R3-1# set group P3 neighbor 172.31.102.1 export export-p3

[edit protocols bgp]
lab@mxB:R3-1# commit
commit complete

[edit protocols bgp]
lab@mxB:R3-1#
```

Step 4.5

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

From the SSH session with the mxC router, use the **`show route 172.16.1.0/24`** command to display the path of the summary route for AS 65001.

```
lab@mxC> show route 172.16.1.0/24

inet.0: 25 destinations, 30 routes (25 active, 0 holddown, 4 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24       *[BGP/170] 00:06:59, localpref 100
                       AS path: 65412 65001 I, validation-state: unverified
                     > to 172.22.124.2 via ge-0/0/1.0
                     [BGP/170] 13:58:43, localpref 100
                       AS path: 65412 65001 I, validation-state: unverified
                     > to 172.22.122.2 via ge-0/0/3.0
```

Question: What happened to the BGP update from the P3 router?

_____

_____

Answer: The BGP update was never sent from the P3 peer to R3-2. P3 determined that AS number 65002 matched the next router in the path and suppressed the advertisement. This is the default Juniper loop prevention mechanism.

Step 4.6

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 export-aggregate]` hierarchy. Display the policy with the **show** command. Using the **as-path-prepend** command insert AS 65001 three times in to the AS path before advertising the summary route. Commit the changes when completed.

```
[edit protocols bgp group P1-P2]
lab@mxA# top edit policy-options policy-statement export-aggregate

[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# set term 1 then as-path-prepend "65001 65001 65001"

[edit policy-options policy-statement export-aggregate]
lab@mxA# commit
commit complete

[edit policy-options policy-statement export-aggregate]
lab@mxA#
```

Step 4.7

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

From the SSH session with the mxC router, use the **show route 172.16.1.0/24** command to display the path of the summary route from AS 65001.

```
lab@mxC> show route 172.16.1.0/24

inet.0: 25 destinations, 30 routes (25 active, 0 holddown, 4 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24      *[BGP/170] 00:01:53, localpref 100
                      AS path: 65412 65001 65001 65001 65001 I, validation-state:
   unverified
                    > to 172.22.122.2 via ge-0/0/3.0
                    [BGP/170] 00:01:53, localpref 100
                      AS path: 65412 65001 65001 65001 65001 I, validation-state:
   unverified
                    > to 172.22.124.2 via ge-0/0/1.0
```

Question: What is the AS Path for the summary route now? Why are there four of your AS numbers?

_____

_____

Answer: In the example output, the AS path for the summary route is `65412 65001 65001 65001 65001`. The **`as-path-prepend`** command in the export policy inserted three of your AS numbers which is in addition to the single AS prepend that mxA performs automatically when advertising a route to its EBGP peer.

Step 4.8

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 policy-options policy-statement export-p3]` hierarchy. Display the policy using the **`show`** command and delete the `as-path-prepend` option. Commit the change when completed.

```
[edit protocols bgp]
lab@mxB:R3-1# top edit policy-options policy-statement export-p3

[edit policy-options policy-statement export-p3]
lab@mxB:R3-1# show
term 1 {
    from {
        protocol aggregate;
        route-filter 172.16.1.0/24 exact;
    }
    then {
        as-path-prepend 65002;
        accept;
    }
}
term 2 {
    from {
        route-filter 172.16.1.0/24 longer;
    }
    then reject;
}

[edit policy-options policy-statement export-p3]
lab@mxB:R3-1# delete term 1 then as-path-prepend

[edit policy-options policy-statement export-p3]
lab@mxB:R3-1# commit
commit complete

[edit policy-options policy-statement export-p3]
lab@mxB:R3-1#
```

Step 4.9

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

From the SSH session with the mxC router, use the **show route 172.16.1.0/24** command to display the path of the summary route from AS 65001.

```
lab@mxC> show route 172.16.1.0/24

inet.0: 25 destinations, 29 routes (25 active, 0 holddown, 4 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24       *[BGP/170] 00:01:26, localpref 100, from 172.16.2.2
                       AS path: 65020 65001 I, validation-state: unverified
                     > to 10.0.14.2 via ge-0/0/0.0
```

Question: Which of the three BGP routes is active and why?

_____

_____

Answer: The active route is the IBGP update from the R3-2 router. The AS path through the R3 router has a length of two while the AS path through P1 and P2 has a length of five.

Step 4.10

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 export-aggregate]` hierarchy. Display the *export-aggregate* policy using the **show** command and delete the `as-path-prepend` statement. Commit the change when completed.

```
[edit policy-options policy-statement export-aggregate]
lab@mxA# show
term 1 {
    from {
        protocol aggregate;
        route-filter 172.16.1.0/24 exact;
    }
    then {
        as-path-prepend "65001 65001 65001";
        accept;
    }
}
term 2 {
    from {
        route-filter 172.16.1.0/24 longer;
    }
    then reject;
}

[edit policy-options policy-statement export-aggregate]
lab@mxA# delete term 1 then as-path-prepend

[edit policy-options policy-statement export-aggregate]
lab@mxA# commit
commit complete
```

```
[edit policy-options policy-statement export-aggregate]
lab@mxA#
```

Step 4.11

Use the **`run show route receive-protocol bgp 172.22.121.2`** command to display the routes being received from the P1 router.

```
[edit policy-options policy-statement export-aggregate]
lab@mxA# run show route receive-protocol bgp 172.22.121.2

inet.0: 21 destinations, 26 routes (21 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)
```

Question: How many routes are being received from the P1 peer?

_____

Answer: Five routes are being received from the P1 peer.

Step 4.12

Use a regular expression in the **`run show route receive-protocol bgp 172.22.121.2 aspath-regex ".* 65002"`** command to only display routes that originate in AS 65002.

```
[edit policy-options policy-statement export-aggregate]
lab@mxA# run show route receive-protocol bgp 172.22.121.2 aspath-regex ".* 65002"

inet.0: 21 destinations, 26 routes (21 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.16.2.0/24           172.22.121.2                            65412 65002 I

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Which routes were displayed? Why?

_____

_____

Answer: Only the 172.16.2.0/24 route was displayed because it is the only route that originated in AS 65002. The route may or may not be active.

Step 4.13

In the following steps, you use regular expressions to create a BGP import policy to only accept your partner's summary route from the P1 peer.

Navigate to the `[edit policy-options]` hierarchy. Create an `as-path` named *`remote-as`* to match on all routes that originate in AS 65002.

```
[edit policy-options policy-statement export-aggregate]
lab@mxA# top edit policy-options

[edit policy-options]
lab@mxA# set as-path remote-as ".* 65002"

[edit policy-options]
```

Step 4.14

Navigate to the `[edit policy-options policy-statement import-P1]` hierarchy. Create a policy, using the `as-path` you just created, which accepts only routes that originate in AS 65002.

```
[edit policy-options]
lab@mxA# edit policy-statement import-P1

[edit policy-options policy-statement import-P1]
lab@mxA# set term 1 from protocol bgp

[edit policy-options policy-statement import-P1]
lab@mxA# set term 1 from as-path remote-as

[edit policy-options policy-statement import-P1]
lab@mxA# set term 1 then accept

[edit policy-options policy-statement import-P1]
lab@mxA# set term 2 then reject

[edit policy-options policy-statement import-P1
```

Step 4.15

Navigate to the `[edit protocols bgp group P1-P2]` hierarchy and apply the *`import-P1`* policy to the P1 neighbor as an import policy. Commit the configuration when completed.

```
[edit policy-options policy-statement import-P1]
lab@mxA# top edit protocols bgp group P1-P2

[edit protocols bgp group P1-P2]
lab@mxA# set neighbor 172.22.121.2 import import-P1

[edit protocols bgp group P1-P2]
lab@mxA# commit
commit complete

[edit protocols bgp group P1-P2]
lab@mxA#
```

Step 4.16

Use the **`run show route receive-protocol bgp 172.22.121.2`** command to display the routes being received from P1.

```
[edit protocols bgp group P1-P2]
lab@mxA# run show route receive-protocol bgp 172.22.121.2

inet.0: 25 destinations, 31 routes (25 active, 0 holddown, 4 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.16.2.0/24           172.22.121.2                            65412 65002 I

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Is the P1 import policy working?

_____

Answer: Yes, the import policy applied to the P1 peer is only accepting the summary route for AS 65002.

Step 4.17

Use the command **`run show route advertising-protocol bgp 172.22.121.2`** to display the routes being advertised to the P1 peer.

```
[edit protocols bgp group P1-P2]
lab@mxA# run show route advertising-protocol bgp 172.22.121.2

inet.0: 25 destinations, 31 routes (25 active, 0 holddown, 4 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 40.40.0.0/24            Self                                    65020 I
* 40.40.1.0/24            Self                                    65020 I
* 40.40.2.0/24            Self                                    65020 I
* 40.40.3.0/24            Self                                    65020 I
* 172.16.1.0/24           Self                                    I
```

Question: How many routes are being advertised to the P1 peer?

_____

Answer: The mxA router is advertising five routes to the P1 peer.

Step 4.18

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 of 172.16.10.0/24 in the R3-1 routing instance with a next-hop of `reject`. Commit the configuration when completed.

```
[edit policy-options policy-statement export-p3]
lab@mxB:R3-1# top edit routing-options

[edit routing-options]
lab@mxB:R3-1# set static route 172.16.10.0/24 reject

[edit routing-options]
lab@mxB:R3-1# commit
commit complete
```

```
[edit routing-options]
lab@mxB:R3-1#
```

Step 4.19

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

From the SSH session with the mxA router, use regular expressions in the **`run show route advertising-protocol bgp 172.22.121.2 aspath-regex "()"`** command to only display routes that originate in your autonomous system.

```
[edit protocols bgp group P1-P2]
lab@mxA# run show route advertising-protocol bgp 172.22.121.2 aspath-regex "()"

inet.0: 26 destinations, 32 routes (26 active, 0 holddown, 4 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.0/24           Self                                    I
* 172.16.10.0/24          Self                                    I
```

Step 4.20

Using regular expressions, you will now modify the BGP export policy for P1 and P2 to suppress all internal BGP routes from being advertised. Navigate to the `[edit policy-options]` hierarchy. Create an `as-path` named *`internal-as`* to match on all internal BGP routes.

```
[edit protocols bgp group P1-P2]
lab@mxA# top edit policy-options

[edit policy-options]
lab@mxA# set as-path internal-as "()"

[edit policy-options]
lab@mxA#
```

Step 4.21

Navigate to the `[edit policy-options policy-statement` *`export-aggregate`*`]` hierarchy and display the policy using the **`show`** command. Using the `as-path` configured in the previous step, create a new term to suppress the internal BGP routes from being advertised to the P1 and P2 peers. Commit the configuration and return to operational mode.

```
[edit policy-options]
lab@mxA# edit policy-statement export-aggregate

[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# set term 3 from protocol bgp

[edit policy-options policy-statement export-aggregate]
lab@mxA# set term 3 from as-path internal-as

[edit policy-options policy-statement export-aggregate]
lab@mxA# set term 3 then reject

[edit policy-options policy-statement export-aggregate]
lab@mxA# commit and-quit
commit complete
Exiting configuration mode

lab@mxA>
```

Step 4.22

Use regular expressions in the **`run show route advertising-protocol bgp 172.22.121.2 aspath-regex "()"`** command to only display routes that originate in your autonomous system.

```
lab@mxA> show route advertising-protocol bgp 172.22.121.2 aspath-regex "()"

inet.0: 26 destinations, 32 routes (26 active, 0 holddown, 4 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.16.1.0/24           Self                                    I
```

Question: Are the internal BGP routes being suppressed to the P1 peer?

_____

Answer: Yes. The remaining route is an internal route but it is being advertised due to term *1* of the *`export-aggregate`* policy.

Step 4.23

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

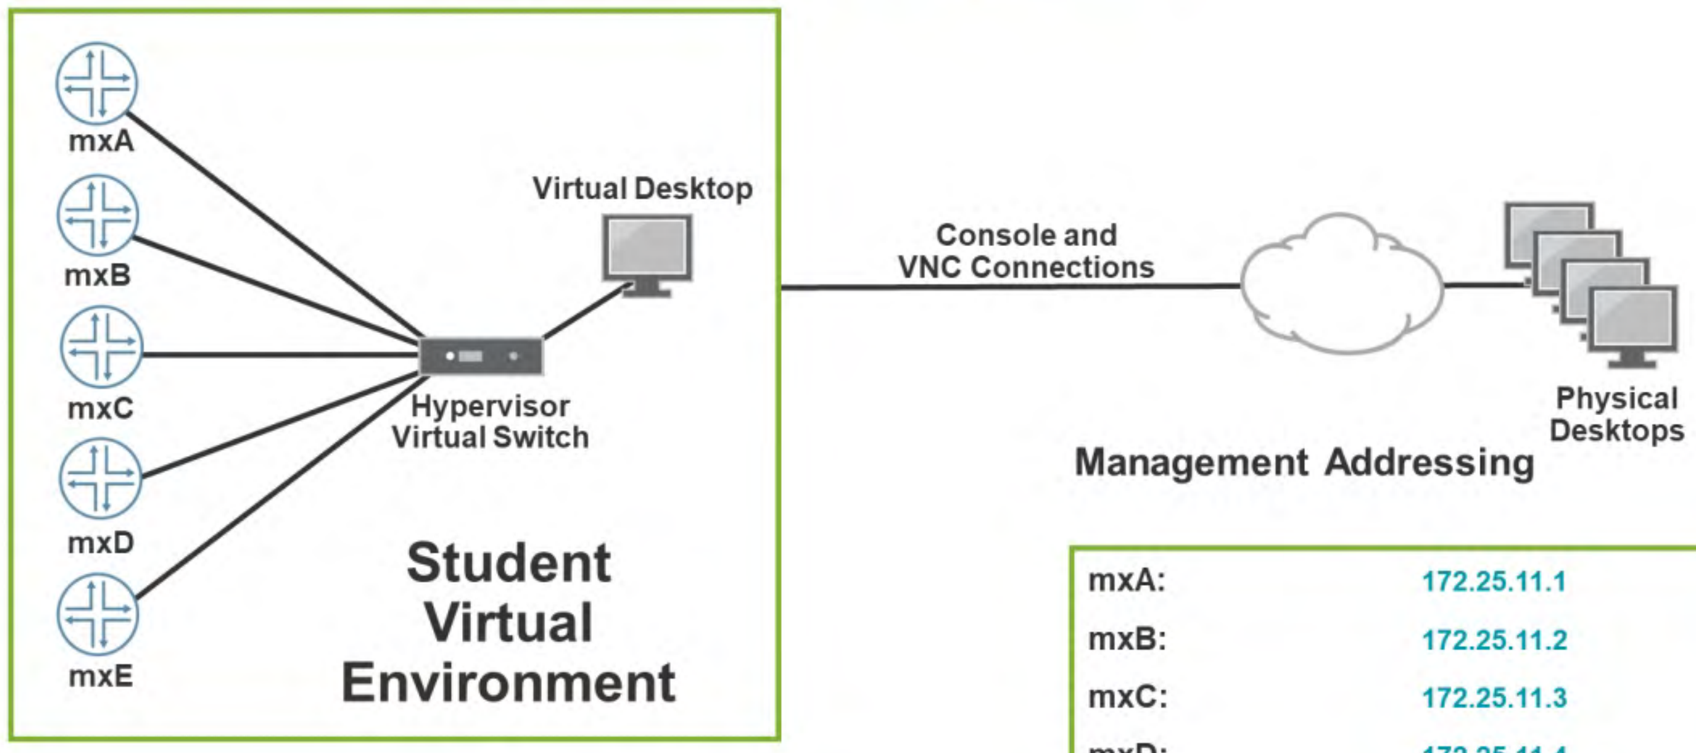

Management Addressing

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: BGP Attributes – Part 1

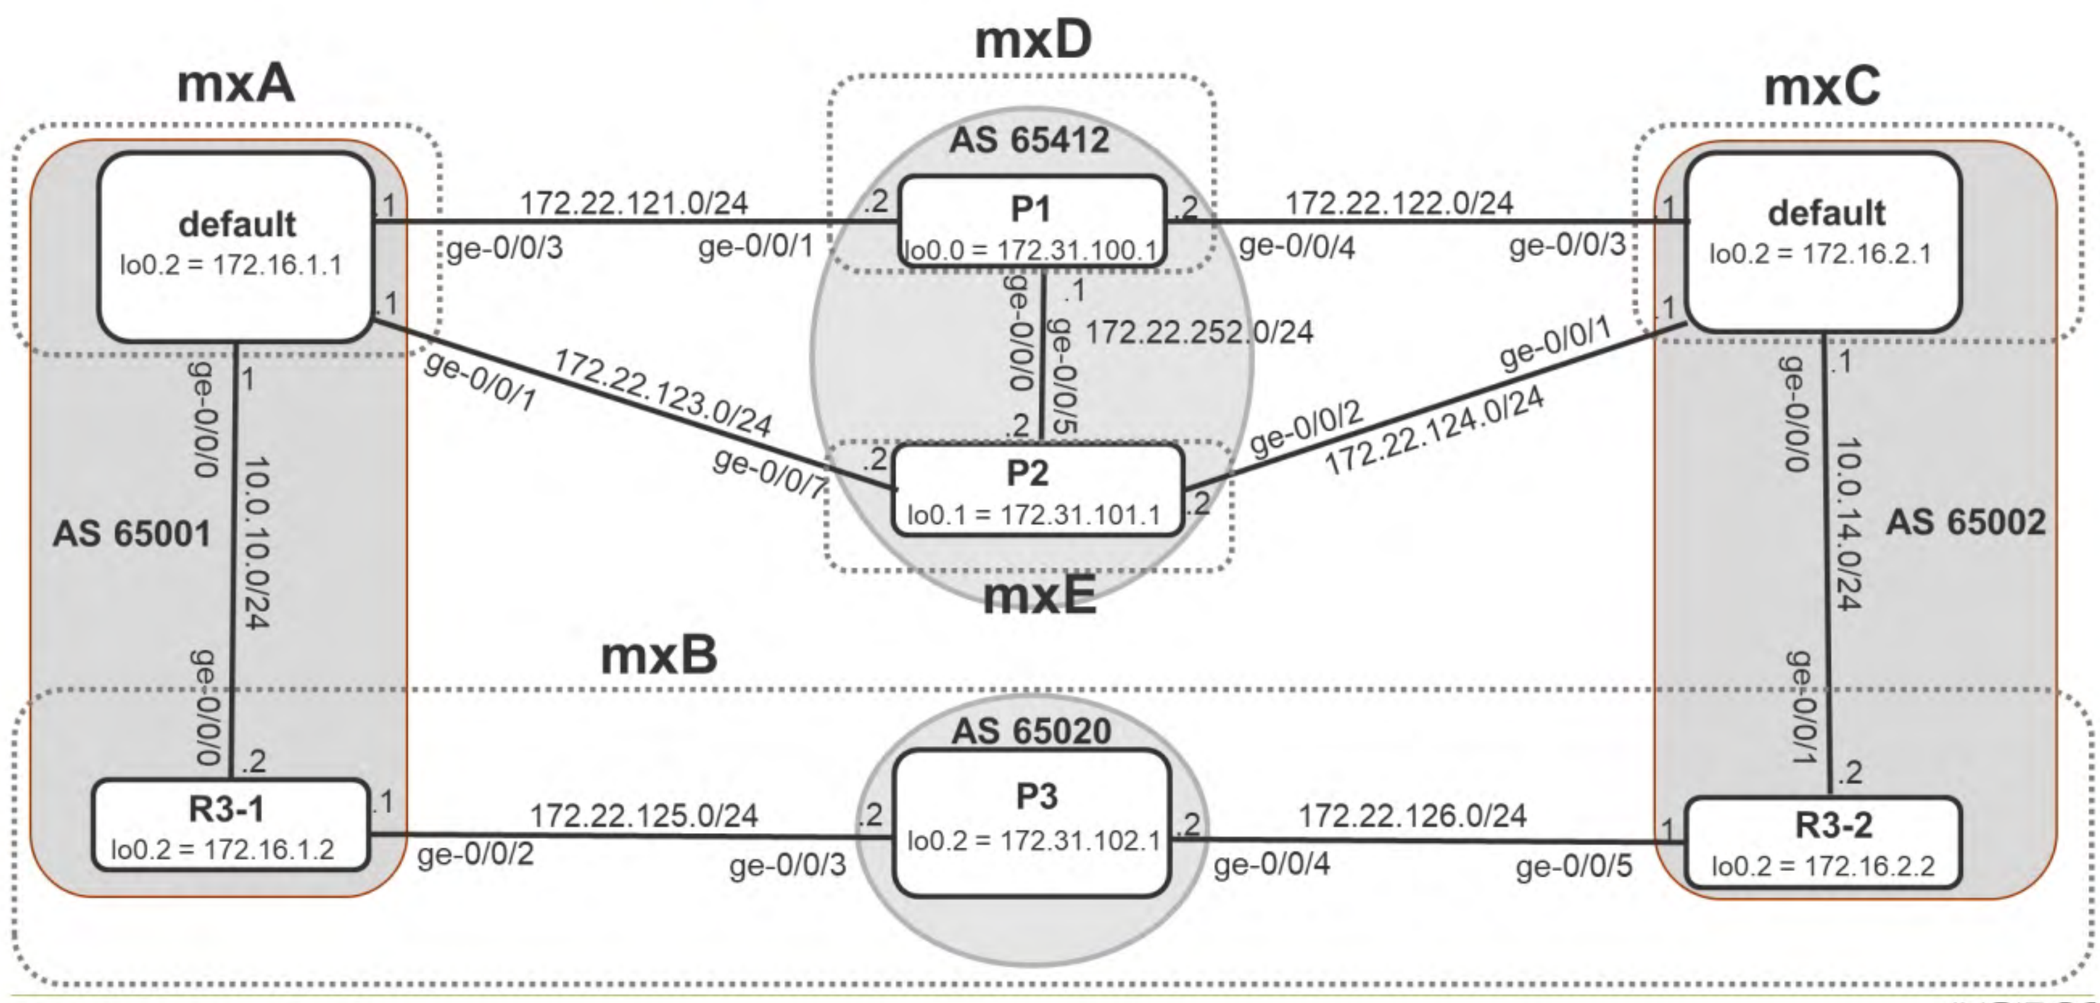

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab

BGP Attributes: Part 2

Overview

In this lab, you use the lab diagram titled “Lab Network Diagram: BGP Attributes - Part 2” to influence traffic leaving your autonomous system. The Local-Preference attribute is used in this lab to define a preferred exit point out for your AS for routes being received from AS 65412. In addition, you use communities to tag the routes being received from the P1, P2, and P3 routers.

By completing this lab you perform the following tasks:

- Load the starting configuration.
- Influence routing using the Origin attribute.
- Influence routing using the MED attribute.
- Use communities to tag routes.
- Influence routing by matching specific communities.

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.

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/lab11-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 26 19:38:00 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/lab11-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 BGP has been enabled.

```
lab@mxB> show configuration
## Last commit: 2019-11-26 19:33:31 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;
                }
            }
            bgp {
                group ibgp {
                    type internal;
                    local-address 172.16.1.2;
                    export [ redistribute-statics next-hop-self ];
                    neighbor 172.16.1.1;
                }
                group P3 {
                    type external;
                    multihop;
                    local-address 172.16.1.2;
                    export export-aggregate;
                    peer-as 65020;
                    neighbor 172.31.102.1;
                }
            }
        }
        policy-options {
            policy-statement export-aggregate {
                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;
                }
            }
            policy-statement next-hop-self {
                term 1 {
                    from {
                        protocol bgp;
                        route-type external;
                    }
                    then {
                        next-hop self;
                    }
                }
            }
            policy-statement redistribute-statics {
                term 1 {
                    from protocol static;
                    then accept;
                }
            }
        }
        routing-options {
```

```
        static {
            route 172.16.1.128/26 reject;
            route 172.16.1.192/26 reject;
            route 172.31.102.1/32 {
                next-hop 172.22.125.2;
                no-readvertise;
            }
        }
        autonomous-system 65001;
        aggregate {
            route 172.16.1.0/24;
        }
    }
}
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 next-hop-self ];
                neighbor 172.16.2.1;
            }
            group P3 {
                type external;
                multihop;
```

```
                    local-address 172.16.2.2;
                    export export-aggregate;
                    neighbor 172.31.102.1 {
                        export export-p3;
                        peer-as 65020;
                    }
                }
            }
        }
        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 next-hop-self {
                term 1 {
                    from {
                        protocol bgp;
                        route-type external;
                    }
                    then {
                        next-hop self;
                    }
                }
            }
            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: Based on the configuration, which routers (logical systems) will mxB represent in this lab?

Answer: The mxB device will, using logical systems, represent the P3, R3-1, and R3-2 routers.

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 any 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 configured for an IBGP session with the mxA router and an EBGP session with the P3 router.

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: What routes will the R3-2 router be advertising to its EBGP peer?

_____

Answer: The R3-2 router is configured with an export policy called *export-aggregate* that is applied to its EBGP peer. This policy causes the R3-2 router to advertise the configured aggregate route of 172.16.2.0/24.

Step 1.3

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
Threading mode: BGP I/O
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      10          6          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2            65001        841        814       0       0     6:04:05 Establ
  inet.0: 1/5/5/0
172.16.2.2            65002        934        926       0       0     6:52:30 Establ
  inet.0: 5/5/5/0
```

Question: What is the state of the P3 router's BGP neighbor relationships?

_____

Answer: The P3 router should have an established EBGP session with both the R3-1 and R3-2 routers.

Step 1.4

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/lab11-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 26 19:44:21 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/lab11-start.config
load complete

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

lab@mxC>
```

Step 1.5

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: 2019-11-26 19:59:44 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 import-P1 {
        term 1 {
            from {
                protocol bgp;
                as-path partner-as;
            }
            then accept;
        }
        term 2 {
            then reject;
        }
    }
    policy-statement next-hop-self {
        term 1 {
            from {
                protocol bgp;
                route-type external;
            }
            then {
                next-hop self;
            }
        }
    }
    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;
        }
    }
    as-path partner-as ".* 65001";
    as-path internal-as "()";
}
```

```
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 next-hop-self ];
            neighbor 172.16.2.2;
        }
        group P1-P2 {
            type external;
            export export-aggregate;
            peer-as 65412;
            neighbor 172.22.122.2 {
                import import-P1;
            }
            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: What routes will the mxC router be advertising to its EBGP peers?

_____

_____

Answer: The mxC router is configured with an export policy called *export-aggregate* that is applied to its EBGP peers. This policy causes the mxC router to advertise the configured aggregate route of 172.16.2.0/24.

Step 1.6

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 1.7

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
                    17         11          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.2.2            65002        914        916       0       0     6:45:23 Establ
  inet.0: 6/7/7/0
172.22.122.2          65412          0          0       0       0        2:04 Connect
172.22.124.2          65412          0          0       0       0        2:04 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. **Because of the state of the lab, you may see an established session with both the P1 and P2 routers. If you do, there is no problem.**

Step 1.8

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/lab11-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 Nov 26 16:19:46 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/lab11-start.config
load complete

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

lab@mxE>
```

Step 1.9

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: 2019-11-26 20:03:37 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 1.10

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
                      16          6          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        7:56 Connect
172.22.124.1          65002        748        755       0       1     5:33:49 Establ
  inet.0: 6/6/6/0
172.31.100.1          65412          0          0       0       0        7:56 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. **Depending on the state of the topology, you may see the BGP sessions with the P1 and mxA router in the established state. If so, this is not a problem.**

Step 1.11

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/lab11-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 26 15:59:57 2019 from 172.25.11.254
--- JUNOS 19.3R1.8 Kernel 64-bit  JNPR-11.0-20190701.269d466_buil
lab@mxD>
```

Step 1.12

Enter configuration mode and load the device's starting configuration by issuing the **`load override`** ***`ajspr/lab11-start.config`*** command. After the configuration has been loaded, commit the changes and exit to operational mode.

```
lab@mxD> configure
Entering configuration mode

[edit]
lab@mxD# load override ajspr/lab11-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: 2019-11-26 20:04:55 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
                      16          6          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        1:21 Connect
172.22.122.1          65002        739        747       0       1     5:29:39 Establ
  inet.0: 6/6/6/0
172.31.101.1          65412        755        752       0       2     5:29:39 Establ
  inet.0: 0/6/6/0
```

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. **Depending on the state of the lab, you may see that the BGP session with the mxA router is in the established state. If so, this is not a problem.**

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/lab11-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 26 13:36:36 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/lab11-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, OSPF, and BGP configuration.

```
lab@mxA> show configuration
## Last commit: 2019-11-26 19:34:59 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;
                }
            }
        }
    }
    policy-options {
        policy-statement export-aggregate {
            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;
            }
        }
        policy-statement next-hop-self {
            term 1 {
                from {
                    protocol bgp;
                    route-type external;
                }
                then {
                    next-hop self;
                }
            }
        }
        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 172.16.1.0/26 reject;
        route 172.16.1.64/26 reject;
    }
    forwarding-table {
        export pfe-load-balance;
    }
    autonomous-system 65001;
    aggregate {
        route 172.16.1.0/24;
    }
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-0/0/0.0;
        }
    }
    bgp {
        group ibgp {
            type internal;
            local-address 172.16.1.1;
            export [ redistribute-statics next-hop-self ];
            neighbor 172.16.1.2;
        }
        group P1-P2 {
            type external;
            export export-aggregate;
            peer-as 65412;
            multipath;
            neighbor 172.22.121.2;
            neighbor 172.22.123.2;
        }
    }
}
```

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: Which routers are mxA configured to peer with using BGP?

_____

_____

Answer: The mxA router is configured with an IBGP session with the R3-1 router and an EBGP session with the P1 and P2 routers.

Question: What routes will the mxA router be advertising to its EBGP peers?

_____

_____

Answer: The mxA router is configured with an export policy called *export-aggregate* that is applied to its EBGP peers. This policy causes the mxA router to advertise the configured aggregate route of 172.16.1.0/24.

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
Threading mode: BGP I/O
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      17         16          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2            65001        720        711       0       2     5:09:45 Establ
  inet.0: 6/7/7/0
172.22.121.2          65412        696        710       0       3     5:09:45 Establ
  inet.0: 5/5/5/0
172.22.123.2          65412        696        709       0       3     5:09:45 Establ
  inet.0: 5/5/5/0
```

Question: What is the status of the mxA router's BGP neighbor relationships?

_____

_____

Answer: The mxA router has established BGP sessions with the R3-1, P1, and P2 routers.

Part 2: Modifying the Origin Attribute

In this lab part, you modify the origin attribute. The BGP origin attribute is a well known mandatory attribute used in the route selection processes. It has three possible values, `IGP (I), EGP (E) or incomplete (?)`. In the route selection process IGP is preferred over EGP and EGP is preferred over incomplete. Using import policy to change the BGP origin, routes can be influenced to prefer a specific peering session. Because the attribute is a transitive attribute, it can also be used in an export policy to influence traffic into the AS.

Step 2.1

From the SSH session with the mxA router, you will need access configuration mode, change to the `protocols bgp` hierarchy. Display the configuration and then remove the `multipath` command from the BGP group *`P1-P2`* in order to demonstrate that the BGP origin attribute can influence traffic. Doing so allows the router to use the BGP route selection process on the duplicate routes being received from P1 and P2 (they are both advertising routes in the 30.30/22 range). Once `multipath` is removed from the configuration, the router will choose only a single path towards routes received from AS 65412. Delete the `multipath` option. Commit the change when completed.

```
lab@mxA> configure
Entering configuration mode

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

```
[edit protocols bgp]
lab@mxA# show
group ibgp {
    type internal;
    local-address 172.16.1.1;
    export [ redistribute-statics next-hop-self ];
    neighbor 172.16.1.2;
}
group P1-P2 {
    type external;
    export export-aggregate;
    peer-as 65412;
    multipath;
    neighbor 172.22.121.2;
    neighbor 172.22.123.2;
}

[edit protocols bgp]
lab@mxA# delete group P1-P2 multipath

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

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

Step 2.2

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 display the routes being received from the P1 and the P2 peers.

```
[edit protocols bgp]
lab@mxA# run show route receive-protocol bgp 172.22.121.2

inet.0: 25 destinations, 31 routes (25 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: 25 destinations, 31 routes (25 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: How many routes are active from the P1 peer? How many routes are active from the P2 peer?

_____

_____

Answer: The asterisk (*) indicates active routes being received from a BGP peer. In this example, the P2 peer has five active routes and the P1 peer has no active routes. Your results might vary.

Step 2.3

Use the **run show route 172.16.2.0/24 detail** command to display the summary route being advertised by AS 65002.

```
[edit protocols bgp]
lab@mxA# run show route 172.16.2.0/24 detail

inet.0: 25 destinations, 31 routes (25 active, 0 holddown, 0 hidden)
172.16.2.0/24 (3 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 582
                Address: 0xc0f1bcc
                Next-hop reference count: 6
                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: 2:02
                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 65002 I
                Aggregator: 65002 172.16.2.1
                Accepted
                Localpref: 100
                Router ID: 172.31.100.1
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 577
                Address: 0xc0f1c94
                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: 1:05:18
                Validation State: unverified
```

```
                Task: BGP_65412.172.22.123.2
                AS path: 65412 65002 I
                Aggregator: 65002 172.16.2.1
                Accepted
                Localpref: 100
                Router ID: 172.31.101.1
        BGP     Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xc0f19d8
                Next-hop reference count: 13
                Source: 172.16.1.2
                Next hop type: Router, Next hop index: 592
                Next hop: 10.0.10.2 via ge-0/0/0.0, selected
                Session Id: 0x185
                Protocol next hop: 172.16.1.2
                Indirect next hop: 0xc285884 1048574 INH Session ID: 0x18d
                State: <Int Ext>
                Inactive reason: Interior > Exterior > Exterior via Interior
                Local AS: 65001 Peer AS: 65001
                Age: 13:29      Metric2: 1
                Validation State: unverified
                Task: BGP_65001.172.16.1.2
                AS path: 65020 65002 I
                Aggregator: 65002 172.16.2.2
                Accepted
                Localpref: 100
                Router ID: 172.16.1.2
```

Question: How many BGP peers are advertising the summary route?

_____

_____

Answer: Three BGP peers are advertising the summary route in to the mxA router, P1, P2, and R3-1.

Question: Which peer is being preferred to activate the route?

_____

_____

Answer: EBGP routes are always preferred over IBGP routes so the preferred peer is P1 or P2. In the previous output, BGP has chosen the P1 router for the active path. Your results might show P2 as the active path.

Question: What is the Origin of each of the three received routes?

_____

_____

Answer: All three routes have an Origin value of IGP. In Junos, this value appears in the AS Path details.

Step 2.4

Navigate to the `[edit policy-options policy-statement P1-P2-import]` hierarchy. Write the import policy to change the `origin` to `egp` on the 172.16.2.0/24 summary route being received from P1 and P2.

```
[edit protocols bgp]
lab@mxA# top edit policy-options policy-statement P1-P2-import

[edit policy-options policy-statement P1-P2-import]
lab@mxA# set term 1 from protocol bgp

[edit policy-options policy-statement P1-P2-import]
lab@mxA# set term 1 from route-filter 172.16.2.0/24 exact

[edit policy-options policy-statement P1-P2-import]
lab@mxA# set term 1 then origin egp

[edit policy-options policy-statement P1-P2-import]
lab@mxA# show
term 1 {
    from {
        protocol bgp;
        route-filter 172.16.2.0/24 exact;
    }
    then origin egp;
}

[edit policy-options policy-statement P1-P2-import]
lab@mxA#
```

Step 2.5

Navigate to the `[edit protocols bgp]` hierarchy and apply the *`P1-P2-import`* policy as an import policy to the *`P1-P2`* group. Commit the changes when completed.

```
[edit policy-options policy-statement P1-P2-import]
lab@mxA# top edit protocols bgp

[edit protocols bgp]
lab@mxA# set group P1-P2 import P1-P2-import

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

Step 2.6

Use the **run show route 172.16.2.0/24 detail** command to display your partner's summary route in the default routing instance.

```
[edit protocols bgp]
lab@mxA# run show route 172.16.2.0/24 detail

inet.0: 25 destinations, 31 routes (25 active, 0 holddown, 0 hidden)
172.16.2.0/24 (3 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39d6f0
                Next-hop reference count: 14
                Source: 172.16.1.2
                Next hop type: Router, Next hop index: 590
                Next hop: 10.0.10.2 via ge-0/0/0.0, selected
                Session Id: 0x14a
                Protocol next hop: 172.16.1.2
                Indirect next hop: 0xb2b2fd0 1048575 INH Session ID: 0x14b
                State: <Active Int Ext>
                Local AS: 65001 Peer AS: 65001
                Age: 41:08      Metric2: 1
                Validation State: unverified
                Task: BGP_65001.172.16.1.2
                Announcement bits (3): 0-KRT 6-BGP_RT_Background 7-Resolve tree 1
                AS path: 65020 65002 I
                Aggregator: 65002 172.16.2.2
                Accepted
                Localpref: 100
                Router ID: 172.16.1.2
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 589
                Address: 0xb39d450
                Next-hop reference count: 5
                Source: 172.22.121.2
                Next hop: 172.22.121.2 via ge-0/0/3.0, selected
                Session Id: 0x149
                State: <Ext>
                Inactive reason: Origin
                Local AS: 65001 Peer AS: 65412
                Age: 37
                Validation State: unverified
                Task: BGP_65412.172.22.121.2
                AS path: 65412 65002 E
                Aggregator: 65002 172.16.2.1
                Accepted
                Localpref: 100
                Router ID: 172.31.100.1
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 584
                Address: 0xb39d630
                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: 0x148
```

```
State: <NotBest Ext>
Inactive reason: Not Best in its group - Router ID
Local AS: 65001 Peer AS: 65412
Age: 37
Validation State: unverified
Task: BGP_65412.172.22.123.2
AS path: 65412 65002 E
Aggregator: 65002 172.16.2.1
Accepted
Localpref: 100
Router ID: 172.31.101.1
```

Question: Which peer is now being preferred to activate the route? What do you notice about the `Inactive reason` for the P1/P2 routes?

_____

_____

Answer: The R3-1 router is now the preferred BGP peer for the summary route. The P1/P2 routes are set as inactive because of the origin attribute.

Step 2.7

Remove the import policy from the *P1-P2* group. Commit the change and return to operational mode.

```
[edit protocols bgp]
lab@mxA# delete group P1-P2 import

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

lab@mxA>
```

Step 2.8

Use the operational mode command **`show route 172.16.2.0/24`** command to verify that the route again prefers one of the P1 or P2 peers in the default routing instance.

```
lab@mxA> show route 172.16.2.0/24

inet.0: 25 destinations, 31 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.0/24      *[BGP/170] 00:00:46, localpref 100
                      AS path: 65412 65002 I, validation-state: unverified
                    > to 172.22.121.2 via ge-0/0/3.0
                    [BGP/170] 00:00:46, localpref 100
                      AS path: 65412 65002 I, validation-state: unverified
                    > to 172.22.123.2 via ge-0/0/1.0
                    [BGP/170] 00:46:47, localpref 100, from 172.16.1.2
                      AS path: 65020 65002 I, validation-state: unverified
```

```
> to 10.0.10.2 via ge-0/0/0.0
```

Question: Did the preferred path return to the P1 or P2 peer?

_____

_____

Answer: Yes, in the above capture, the origin is now the same on all three routes and, in this case, the Best in Group path is through P1.

Part 3: Configuring the MED Attribute

In this lab part, you configure the MED attribute. The MED is an optional nontransitive attribute used to influence traffic coming into your AS. It is a route metric assigned to a BGP route and advertised to a remote peer to influence the remote peer's route selection process. In this part, you use a MED to influence AS 65412 to always use P2 to route traffic to your summary address.

Step 3.1

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 P1 logical system by issuing the **`set cli logical-system P1`** command.

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

lab@mxD:P1>
```

Step 3.2

Issue the **`show route 172.16.1.0/24`** command to display the path used to forward traffic destined your summary route from the perspective of the P1 router.

```
lab@mxD:P1> show route 172.16.1.0/24

inet.0: 19 destinations, 29 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24       *[BGP/170] 13:13:33, localpref 100
                       AS path: 65001 I, validation-state: unverified
                     > to 172.22.121.1 via ge-0/0/1.0
                     [BGP/170] 13:13:37, localpref 100, from 172.31.101.1
                       AS path: 65001 I, validation-state: unverified
                     > to 172.22.252.2 via ge-0/0/0.0
```

Question: What is the preferred forwarding path of your summary route from the perspective of the P1 router?

_____

_____

Answer: From the P1 router, the forwarding path for your summary route is across its directly connected interface into AS 65001.

Step 3.3

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 P2 logical system by issuing the **`set cli logical-system P2`** command.

```
lab@mxE> set cli logical-system P2
Logical system: P2

lab@mxE:P2>
```

Step 3.4

Issue the **`show route 172.16.1.0/24`** command to display the path used to forward traffic destined your summary route from the perspective of the P2 router.

```
lab@mxE:P2> show route 172.16.1.0/24

inet.0: 19 destinations, 29 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24      *[BGP/170] 13:16:16, localpref 100
                      AS path: 65001 I, validation-state: unverified
                    > to 172.22.123.1 via ge-0/0/7.0
                    [BGP/170] 13:16:11, localpref 100, from 172.31.100.1
                      AS path: 65001 I, validation-state: unverified
                    > to 172.22.252.1 via ge-0/0/5.0
```

Question: What is the preferred forwarding path of your summary route from the perspective of the P2 router?

_____

_____

Answer: From the P2 router, the forwarding path for your summary route is across its directly connected interface into AS 65001.

Step 3.5

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 bgp group P1-P2]` hierarchy. Set the `metric-out` value to *`10`* on the P1 neighbor. Commit the configuration when completed.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit protocols bgp group P1-P2

[edit protocols bgp group P1-P2]
lab@mxA# set neighbor 172.22.121.2 metric-out 10

[edit protocols bgp group P1-P2]
lab@mxA# commit
commit complete

[edit protocols bgp group P1-P2]
lab@mxA#
```

Step 3.6

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 172.16.1.0/24`** command to display the path used to forward traffic destined your summary route from the perspective of the P1 router.

```
lab@mxD:P1> show route 172.16.1.0/24

inet.0: 19 destinations, 29 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24       *[BGP/170] 13:26:56, localpref 100, from 172.31.101.1
                       AS path: 65001 I, validation-state: unverified
                     > to 172.22.252.2 via ge-0/0/0.0
                     [BGP/170] 00:08:38, MED 10, localpref 100
                       AS path: 65001 I, validation-state: unverified
                     > to 172.22.121.1 via ge-0/0/1.0
```

Question: What is the preferred forwarding path of your summary route from the perspective of the P1 router?

_____

Answer: From the P1 router, the preferred path for the summary route is now through the P2 router. **This may take a minute or two to change.**

Question: What additional information is present in the routing information learned from the route received directly from mxA?

_____

_____

Answer: You should see a `MED 10` value in the route information learned by P1 directly from mxA.

Step 3.7

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

From the SSH session with the mxE router, issue the **`show route 172.16.1.0/24`** command to display the path used to forward traffic destined your summary route from the perspective of the P2 router.

```
lab@mxE:P2> show route 172.16.1.0/24

inet.0: 19 destinations, 28 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24      *[BGP/170] 13:29:08, localpref 100
                      AS path: 65001 I, validation-state: unverified
                    > to 172.22.123.1 via ge-0/0/7.0
```

Question: What is the preferred forwarding path of your summary route from the perspective of the P2 router?

_____

Answer: From the P2 router, the preferred path for the summary route is still over the ge-0/0/7.0 interface.

Step 3.8

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]` hierarchy. Copy the ***`export-aggregate`*** policy to a new policy named ***`export-p2`***. Display the new policy with the **`show`** command.

```
[edit protocols bgp group P1-P2]
lab@mxA# top edit policy-options

[edit policy-options]
lab@mxA# copy policy-statement export-aggregate to policy-statement export-p2

[edit policy-options]
lab@mxA# show policy-statement export-p2
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]
lab@mxA#
```

Step 3.9

Navigate to the `[edit policy-options policy-statement export-p2]` hierarchy. Set the `metric` to *20* in `term 1` before accepting the summary route.

```
[edit policy-options]
lab@mxA# edit policy-statement export-p2

[edit policy-options policy-statement export-p2]
lab@mxA# set term 1 then metric 20

[edit policy-options policy-statement export-p2]
lab@mxA#
```

Step 3.10

Navigate to the `[edit protocols bgp group P1-P2]` hierarchy and apply the *export-p2* policy as an export policy under the P2 neighbor statement. Commit the changes when completed.

```
[edit policy-options policy-statement export-p2]
lab@mxA# top edit protocols bgp group P1-P2

[edit protocols bgp group P1-P2]
lab@mxA# set neighbor 172.22.123.2 export export-p2

[edit protocols bgp group P1-P2]
lab@mxA# commit
commit complete

[edit protocols bgp group P1-P2]
lab@mxA#
```

Step 3.11

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 172.16.1.0/24`** command to display the path used to forward traffic destined your summary route from the perspective of the P1 router.

```
lab@mxD:P1> show route 172.16.1.0/24

inet.0: 19 destinations, 28 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24      *[BGP/170] 00:18:13, MED 10, localpref 100
                      AS path: 65001 I, validation-state: unverified
                    > to 172.22.121.1 via ge-0/0/1.0
```

Question: What is the preferred forwarding path of your summary route from the perspective of the P1 router?

_____

_____

Answer: From the P1 router, the preferred path for the summary route is now through ge-0/0/1.0.

Step 3.12

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

From the SSH session with the mxE router, issue the **`show route 172.16.1.0/24`** command to display the path used to forward traffic destined your summary route from the perspective of the P2 router.

```
lab@mxE:P2> show route 172.16.1.0/24

inet.0: 19 destinations, 29 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24       *[BGP/170] 00:02:12, MED 10, localpref 100, from 172.31.100.1
                       AS path: 65001 I, validation-state: unverified
                     > to 172.22.252.1 via ge-0/0/5.0
                      [BGP/170] 00:02:12, MED 20, localpref 100
                       AS path: 65001 I, validation-state: unverified
                     > to 172.22.123.1 via ge-0/0/7.0
```

Question: What is the preferred forwarding path of your summary route from the perspective of the P2 router?

_____

_____

Answer: From the P2 router, the preferred path for the summary route is now over the ge-0/0/5.0 interface.

Question: What additional information is present in the routing information learned from the route received directly from mxA?

_____

_____

Question: You should see a `MED 20` value in the route information learned by P2 directly from mxA.

Part 4: Configuring BGP Communities

In this lab part, you configure BGP communities. The community attribute is an optional transitive attribute. An individual BGP process does not have to understand the community attribute but it must advertise it to all established peers. The community attribute is a 4-octet value in the format FFFF:FFFF. The first two octets represent an autonomous system number and the second two octets represent a locally defined value.

Communities can be used simply to provide an administrative tag value to associate specific routes with specific BGP peers. It can also be used to trigger specific actions with an import routing policy. A route's community value can cause specific routes to be accepted, rejected or modified.

Step 4.1

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

From the SSH session with the mxA router, remove the metric from the P1 neighbor and the export policy from the P2 neighbor. Commit the changes when completed.

```
[edit protocols bgp group P1-P2]
lab@mxA# delete neighbor 172.22.121.2 metric-out

[edit protocols bgp group P1-P2]
lab@mxA# delete neighbor 172.22.123.2 export

[edit protocols bgp group P1-P2]
lab@mxA# show
type external;
export export-aggregate;
peer-as 65412;
neighbor 172.22.121.2;
neighbor 172.22.123.2;

[edit protocols bgp group P1-P2]
lab@mxA# commit
commit complete

[edit protocols bgp group P1-P2]
lab@mxA#
```

Step 4.2

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 policy-options]` hierarchy. Define three communities as shown in the following table. The community name should be the same as the member value.

Name	Value
65002:100	65002:100
65002:110	65002:110
65002:120	65002:120

```
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# edit policy-options

[edit policy-options]
lab@mxC# set community 65002:100 members 65002:100

[edit policy-options]
lab@mxC# set community 65002:110 members 65002:110

[edit policy-options]
lab@mxC# set community 65002:120 members 65002:120

[edit policy-options]
lab@mxC#
```

Step 4.3

Create two routing policies named *export-p1* and *export-p2* by using the **copy** command to copy the *export-aggregate* policy into each of the new policies.

```
[edit policy-options]
lab@mxC# copy policy-statement export-aggregate to policy-statement export-p1

[edit policy-options]
lab@mxC# copy policy-statement export-aggregate to policy-statement export-p2

[edit policy-options]
lab@mxC#
```

Step 4.4

Navigate to the [edit policy-options policy-statement *export-p1*] hierarchy and use the **show** command to display the policy. In term *1*, set the community to the *65002:100* community configured previously.

```
[edit policy-options]
lab@mxC# edit policy-statement export-p1

[edit policy-options policy-statement export-p1]
lab@mxC# show
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;
}

[edit policy-options policy-statement export-p1]
lab@mxC# set term 1 then community set 65002:100

[edit policy-options policy-statement export-p1]
lab@mxC#
```

Step 4.5

Navigate to the [edit policy-options policy-statement *export-p2*] hierarchy and use the **show** command to display the policy. In term *1*, set the community with the *65002:120* community configured previously.

```
[edit policy-options policy-statement export-p1]
lab@mxC# up 1 edit policy-statement export-p2

[edit policy-options policy-statement export-p2]
lab@mxC# show
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;
}

[edit policy-options policy-statement export-p2]
lab@mxC# set term 1 then community set 65002:120

[edit policy-options policy-statement export-p2]
lab@mxC#
```

Step 4.6

Navigate to the `[edit protocols bgp group P1-P2]` hierarchy. Apply the *export-p1* policy as an export policy under the P1 neighbor hierarchy and the *export-p2* policy as an export policy under the P2 neighbor hierarchy. Commit the configuration when finished.

```
[edit policy-options policy-statement export-p2]
lab@mxC# top edit protocols bgp group P1-P2

[edit protocols bgp group P1-P2]
lab@mxC# set neighbor 172.22.122.2 export export-p1

[edit protocols bgp group P1-P2]
lab@mxC# set neighbor 172.22.124.2 export export-p2

[edit protocols bgp group P1-P2]
lab@mxC# commit
commit complete

[edit protocols bgp group P1-P2]
lab@mxC#
```

Step 4.7

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 R3-2 logical system by issuing the **`set cli logical-system R3-2`** command.

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

lab@mxB:R3-2>
```

Step 4.8

Enter configuration mode and navigate to the `[edit policy-options]` hierarchy. Define three communities as shown in the following table. The community name should be the same as the member value.

Name	Value
65002:100	65002:100
65002:110	65002:110
65002:120	65002:120

```
lab@mxB:R3-2> configure
Entering configuration mode

[edit]
lab@mxB:R3-2# edit policy-options

[edit policy-options]
lab@mxB:R3-2# set community 65002:100 members 65002:100

[edit policy-options]
lab@mxB:R3-2# set community 65002:110 members 65002:110

[edit policy-options]
lab@mxB:R3-2# set community 65002:120 members 65002:120

[edit policy-options]
lab@mxB:R3-2#
```

Step 4.9

Create a routing policy named *export-p3* by using the **copy** command to copy the *export-aggregate* policy into the new policy.

```
[edit policy-options]
lab@mxB:R3-2# copy policy-statement export-aggregate to policy-statement export-p3

[edit policy-options]
lab@mxB:R3-2#
```

Step 4.10

Navigate to the `[edit policy-options policy-statement export-p3]` hierarchy and use the **show** command to display the policy. In term *1*, set the `community` to the *65002:110* community configured previously.

```
[edit policy-options]
lab@mxB:R3-2# edit policy-statement export-p3

[edit policy-options policy-statement export-p3]
lab@mxB:R3-2# show
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;
}

[edit policy-options policy-statement export-p3]
lab@mxB:R3-2# set term 1 then community set 65002:110

[edit policy-options policy-statement export-p3]
lab@mxB:R3-2#
```

Step 4.11

Navigate to the `[edit protocols bgp group P3]` hierarchy. Apply the *`export-p3`* policy as an export policy under the P3 neighbor hierarchy. Commit the configuration and exit to operational mode when completed.

```
[edit policy-options policy-statement export-p3]
lab@mxB:R3-2# top edit protocols bgp group P3

[edit protocols bgp group P3]
lab@mxB:R3-2# set neighbor 172.31.102.1 export export-p3

[edit protocols bgp group P3]
lab@mxB:R3-2# commit and-quit
commit complete
Exiting configuration mode

lab@mxB:R3-2>
```

Step 4.12

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 route 172.16.2.0/24 detail`** command to display detailed information about the AS 65002 summary route.

```
[edit protocols bgp group P1-P2]
lab@mxA# run show route 172.16.2.0/24 detail

inet.0: 25 destinations, 31 routes (25 active, 0 holddown, 0 hidden)
172.16.2.0/24 (3 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 577
                Address: 0xc0f1c94
                Next-hop reference count: 10
                Source: 172.22.123.2
                Next hop: 172.22.123.2 via ge-0/0/1.0, selected
                Session Id: 0x188
                State: <Active Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 2:13
                Validation State: unverified
```

```
        Task: BGP_65412.172.22.123.2
        Announcement bits (3): 0-KRT 2-BGP_RT_Background 5-Resolve tree 5
        AS path: 65412 65002 I
        Aggregator: 65002 172.16.2.1
        Communities: 65002:120
        Accepted
        Localpref: 100
        Router ID: 172.31.101.1
BGP     Preference: 170/-101
        Next hop type: Router, Next hop index: 0
        Address: 0xc0f1dc0
        Next-hop reference count: 5
        Source: 172.22.121.2
        Next hop: 172.22.121.2 via ge-0/0/3.0, selected
        Session Id: 0x0
        State: <NotBest Ext>
        Inactive reason: Not Best in its group - Active preferred
        Local AS: 65001 Peer AS: 65412
        Age: 2:15
        Validation State: unverified
        Task: BGP_65412.172.22.121.2
        AS path: 65412 65002 I
        Aggregator: 65002 172.16.2.1
        Communities: 65002:100
        Accepted
        Localpref: 100
        Router ID: 172.31.100.1
BGP     Preference: 170/-101
        Next hop type: Indirect, Next hop index: 0
        Address: 0xc0f19d8
        Next-hop reference count: 13
        Source: 172.16.1.2
        Next hop type: Router, Next hop index: 592
        Next hop: 10.0.10.2 via ge-0/0/0.0, selected
        Session Id: 0x185
        Protocol next hop: 172.16.1.2
        Indirect next hop: 0xc285884 1048574 INH Session ID: 0x18d
        State: <Int Ext Changed>
        Inactive reason: Interior > Exterior > Exterior via Interior
        Local AS: 65001 Peer AS: 65001
        Age: 14          Metric2: 1
        Validation State: unverified
        Task: BGP_65001.172.16.1.2
        AS path: 65020 65002 I
        Aggregator: 65002 172.16.2.2
        Communities: 65002:110
        Accepted
        Localpref: 100
        Router ID: 172.16.1.2
```

Question: Are the community values visible in the received BGP routes?

_____

Answer: In this example, the community values are visible. If you do not see any communities in the BGP routes, check you configuration on the mxC and R3-2 router.

Step 4.13

Navigate to the `[edit policy-options]` hierarchy. Define three communities as shown in the following table. The community name should be the same as the member value.

Name	Value
65002:100	65002:100
65002:110	65002:110
65002:120	65002:120

```
[edit protocols bgp group P1-P2]
lab@mxA# top edit policy-options

[edit policy-options]
lab@mxA# set community 65002:100 members 65002:100

[edit policy-options]
lab@mxA# set community 65002:110 members 65002:110

[edit policy-options]
lab@mxA# set community 65002:120 members 65002:120

[edit policy-options]
lab@mxA#
```

Step 4.14

Navigate to the `[edit policy-options policy-statement import-communities]` hierarchy. Create a three term policy. Each term should set the `local-preference` value on the routes received from the P1 and P2 routers based on the community received. Using the community values received from AS 65002, set the `local-preference` to the same value as the administrative value (last two octets) of the community. Use the `show` command to display the policy.

```
[edit policy-options]
lab@mxA# edit policy-statement import-communities

[edit policy-options policy-statement import-communities]
lab@mxA# set term 1 from community 65002:100

[edit policy-options policy-statement import-communities]
lab@mxA# set term 1 then local-preference 100

[edit policy-options policy-statement import-communities]
lab@mxA# set term 1 then accept
```

```
[edit policy-options policy-statement import-communities]
lab@mxA# set term 2 from community 65002:110

[edit policy-options policy-statement import-communities]
lab@mxA# set term 2 then local-preference 110

[edit policy-options policy-statement import-communities]
lab@mxA# set term 2 then accept

[edit policy-options policy-statement import-communities]
lab@mxA# set term 3 from community 65002:120

[edit policy-options policy-statement import-communities]
lab@mxA# set term 3 then local-preference 120

[edit policy-options policy-statement import-communities]
lab@mxA# set term 3 then accept

[edit policy-options policy-statement import-communities]
lab@mxA# show
term 1 {
    from community 65002:100;
    then {
        local-preference 100;
        accept;
    }
}
term 2 {
    from community 65002:110;
    then {
        local-preference 110;
        accept;
    }
}
term 3 {
    from community 65002:120;
    then {
        local-preference 120;
        accept;
    }
}

[edit policy-options policy-statement import-communities]
lab@mxA#
```

Step 4.15

Navigate to the `[edit protocols bgp group P1-P2]` hierarchy. Configure and apply the *`import-communities`* policy to the *P1-P2* BGP group. Commit the configuration when finished.

```
[edit policy-options policy-statement import-communities]
lab@mxA# top edit protocols bgp group P1-P2

[edit protocols bgp group P1-P2]
```

```
lab@mxA# show
type external;
export export-aggregate;
peer-as 65412;
neighbor 172.22.121.2;
neighbor 172.22.123.2;

[edit protocols bgp group P1-P2]
lab@mxA# set import import-communities

[edit protocols bgp group P1-P2]
lab@mxA# commit
commit complete

[edit protocols bgp group P1-P2]
lab@mxA#
```

Step 4.16

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:R3-2> set cli logical-system R3-1
Logical system: R3-1

lab@mxB:R3-1>
```

Step 4.17

Enter configuration mode and navigate to the `[edit policy-options]` hierarchy. Define three communities as shown in the following table. The community name should be the same as the member value.

Name	Value
65002:100	65002:100
65002:110	65002:110
65002:120	65002:120

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

[edit]
lab@mxB:R3-1# edit policy-options

[edit policy-options]
lab@mxB:R3-1# set community 65002:100 members 65002:100

[edit policy-options]
lab@mxB:R3-1# set community 65002:110 members 65002:110

[edit policy-options]
lab@mxB:R3-1# set community 65002:120 members 65002:120
```

```
[edit policy-options]
lab@mxB:R3-1#
```

Step 4.18

Navigate to the `[edit policy-options policy-statement import-communities]` hierarchy. Create a three term policy. Each term should set the `local-preference` value on the routes received from the P3 router based on the community received. Using the community values received from AS 65002, set the `local-preference` to the same value as the administrative value (last two octets) of the community. Use the `show` command to display the policy.

```
[edit policy-options]
lab@mxB:R3-1# edit policy-statement import-communities

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# set term 1 from community 65002:100

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# set term 1 then local-preference 100

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# set term 1 then accept

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# set term 2 from community 65002:110

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# set term 2 then local-preference 110

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# set term 2 then accept

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# set term 3 from community 65002:120

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# set term 3 then local-preference 120

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# set term 3 then accept

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# show
term 1 {
    from community 65002:100;
    then {
        local-preference 100;
        accept;
    }
}
term 2 {
    from community 65002:110;
    then {
        local-preference 110;
        accept;
```

```
    }
}
term 3 {
    from community 65002:120;
    then {
        local-preference 120;
        accept;
    }
}

[edit policy-options policy-statement import-communities]
lab@mxB:R3-1#
```

Step 4.19

Navigate to the `[edit protocols bgp group P3]` hierarchy. Add the *`import-communities`* policy to the BGP P3 group. Commit the configuration and return to operational mode.

```
[edit policy-options policy-statement import-communities]
lab@mxB:R3-1# top edit protocols bgp group P3

[edit protocols bgp group P3]
lab@mxB:R3-1# show
type external;
multihop;
local-address 172.16.1.2;
export export-aggregate;
peer-as 65020;
neighbor 172.31.102.1;

[edit protocols bgp group P3]
lab@mxB:R3-1# set import import-communities

[edit protocols bgp group P3]
lab@mxB:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxB:R3-1>
```

Step 4.20

Use the **`show route 172.16.2.0/24 detail`** command to display detail information about the AS 65002 summary route.

```
lab@mxB:R3-1> show route 172.16.2.0/24 detail

inet.0: 23 destinations, 24 routes (23 active, 0 holddown, 0 hidden)
172.16.2.0/24 (2 entries, 1 announced)
        *BGP    Preference: 170/-121
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39cbb0
                Next-hop reference count: 14
                Source: 172.16.1.1
                Next hop type: Router, Next hop index: 765
                Next hop: 10.0.10.1 via ge-0/0/0.0, selected
```

```
        Session Id: 0x191
        Protocol next hop: 172.16.1.1
        Indirect next hop: 0xb2b2ca0 1048575 INH Session ID: 0x194
        State: <Active Int Ext>
        Local AS: 65001 Peer AS: 65001
        Age: 2:48:56    Metric2: 1
        Validation State: unverified
        Task: BGP_65001.172.16.1.1
        Announcement bits (3): 3-KRT 5-BGP_RT_Background 6-Resolve tree 1
        AS path: 65412 65002 I
        Aggregator: 65002 172.16.2.1
        Communities: 65002:120
        Accepted
        Localpref: 120
        Router ID: 172.16.1.1
BGP     Preference: 170/-111
        Next hop type: Indirect, Next hop index: 0
        Address: 0xb39cd90
        Next-hop reference count: 9
        Source: 172.31.102.1
        Next hop type: Router, Next hop index: 743
        Next hop: 172.22.125.2 via ge-0/0/2.0, selected
        Session Id: 0x18e
        Protocol next hop: 172.31.102.1
        Indirect next hop: 0xb2b31f0 1048576 INH Session ID: 0x195
        State: <Ext>
        Inactive reason: Local Preference
        Local AS: 65001 Peer AS: 65020
        Age: 3:43       Metric2: 0
        Validation State: unverified
        Task: BGP_65020.172.31.102.1
        AS path: 65020 65002 I
        Aggregator: 65002 172.16.2.2
        Communities: 65002:110
        Accepted
        Localpref: 110
        Router ID: 172.31.102.1
```

Question: Why is the route from the mxA router the active route in the R3-1 router's routing table?

_____

_____

Answer: The inactive reason on the route received from the P3 router is Local Preference. Because the local preference of the route received from the mxA router is higher than the local preference set by the import policy based on the community value from the P3 router, the route from the mxA router is active.

Step 4.21

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 route 172.16.2.0/24 detail`** command to display detail information about the AS 65002 summary route.

```
[edit protocols bgp group P1-P2]
lab@mxA# run show route 172.16.2.0/24 detail

inet.0: 26 destinations, 31 routes (26 active, 0 holddown, 0 hidden)
172.16.2.0/24 (2 entries, 1 announced)
        *BGP    Preference: 170/-121
                Next hop type: Router, Next hop index: 582
                Address: 0xb39d630
                Next-hop reference count: 6
                Source: 172.22.123.2
                Next hop: 172.22.123.2 via ge-0/0/1.0, selected
                Session Id: 0x150
                State: <Active Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 2:46:53
                Validation State: unverified
                Task: BGP_65412.172.22.123.2
                Announcement bits (3): 0-KRT 5-BGP_RT_Background 6-Resolve tree 1
                AS path: 65412 65002 I
                Aggregator: 65002 172.16.2.1
                Communities: 65002:120
                Accepted
                Localpref: 120
                Router ID: 172.31.101.1
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 591
                Address: 0xb39d6f0
                Next-hop reference count: 9
                Source: 172.22.121.2
                Next hop: 172.22.121.2 via ge-0/0/3.0, selected
                Session Id: 0x14c
                State: <Ext>
                Inactive reason: Local Preference
                Local AS: 65001 Peer AS: 65412
                Age: 2:46:53
                Validation State: unverified
                Task: BGP_65412.172.22.121.2
                AS path: 65412 65002 I
                Aggregator: 65002 172.16.2.1
                Communities: 65002:100
                Accepted
                Localpref: 100
                Router ID: 172.31.100.1
```

Question: Which of the routes received from the P1 and P2 routers is active? Why?

_____

_____

Answer: The route from the P2 router is the active route. The route from the P1 router has an inactive reason of Local Preference. The import policy has set the local preference based on the community value and the local preference on the P2 route is higher than the local preference of the P1 route.

Step 4.22

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.

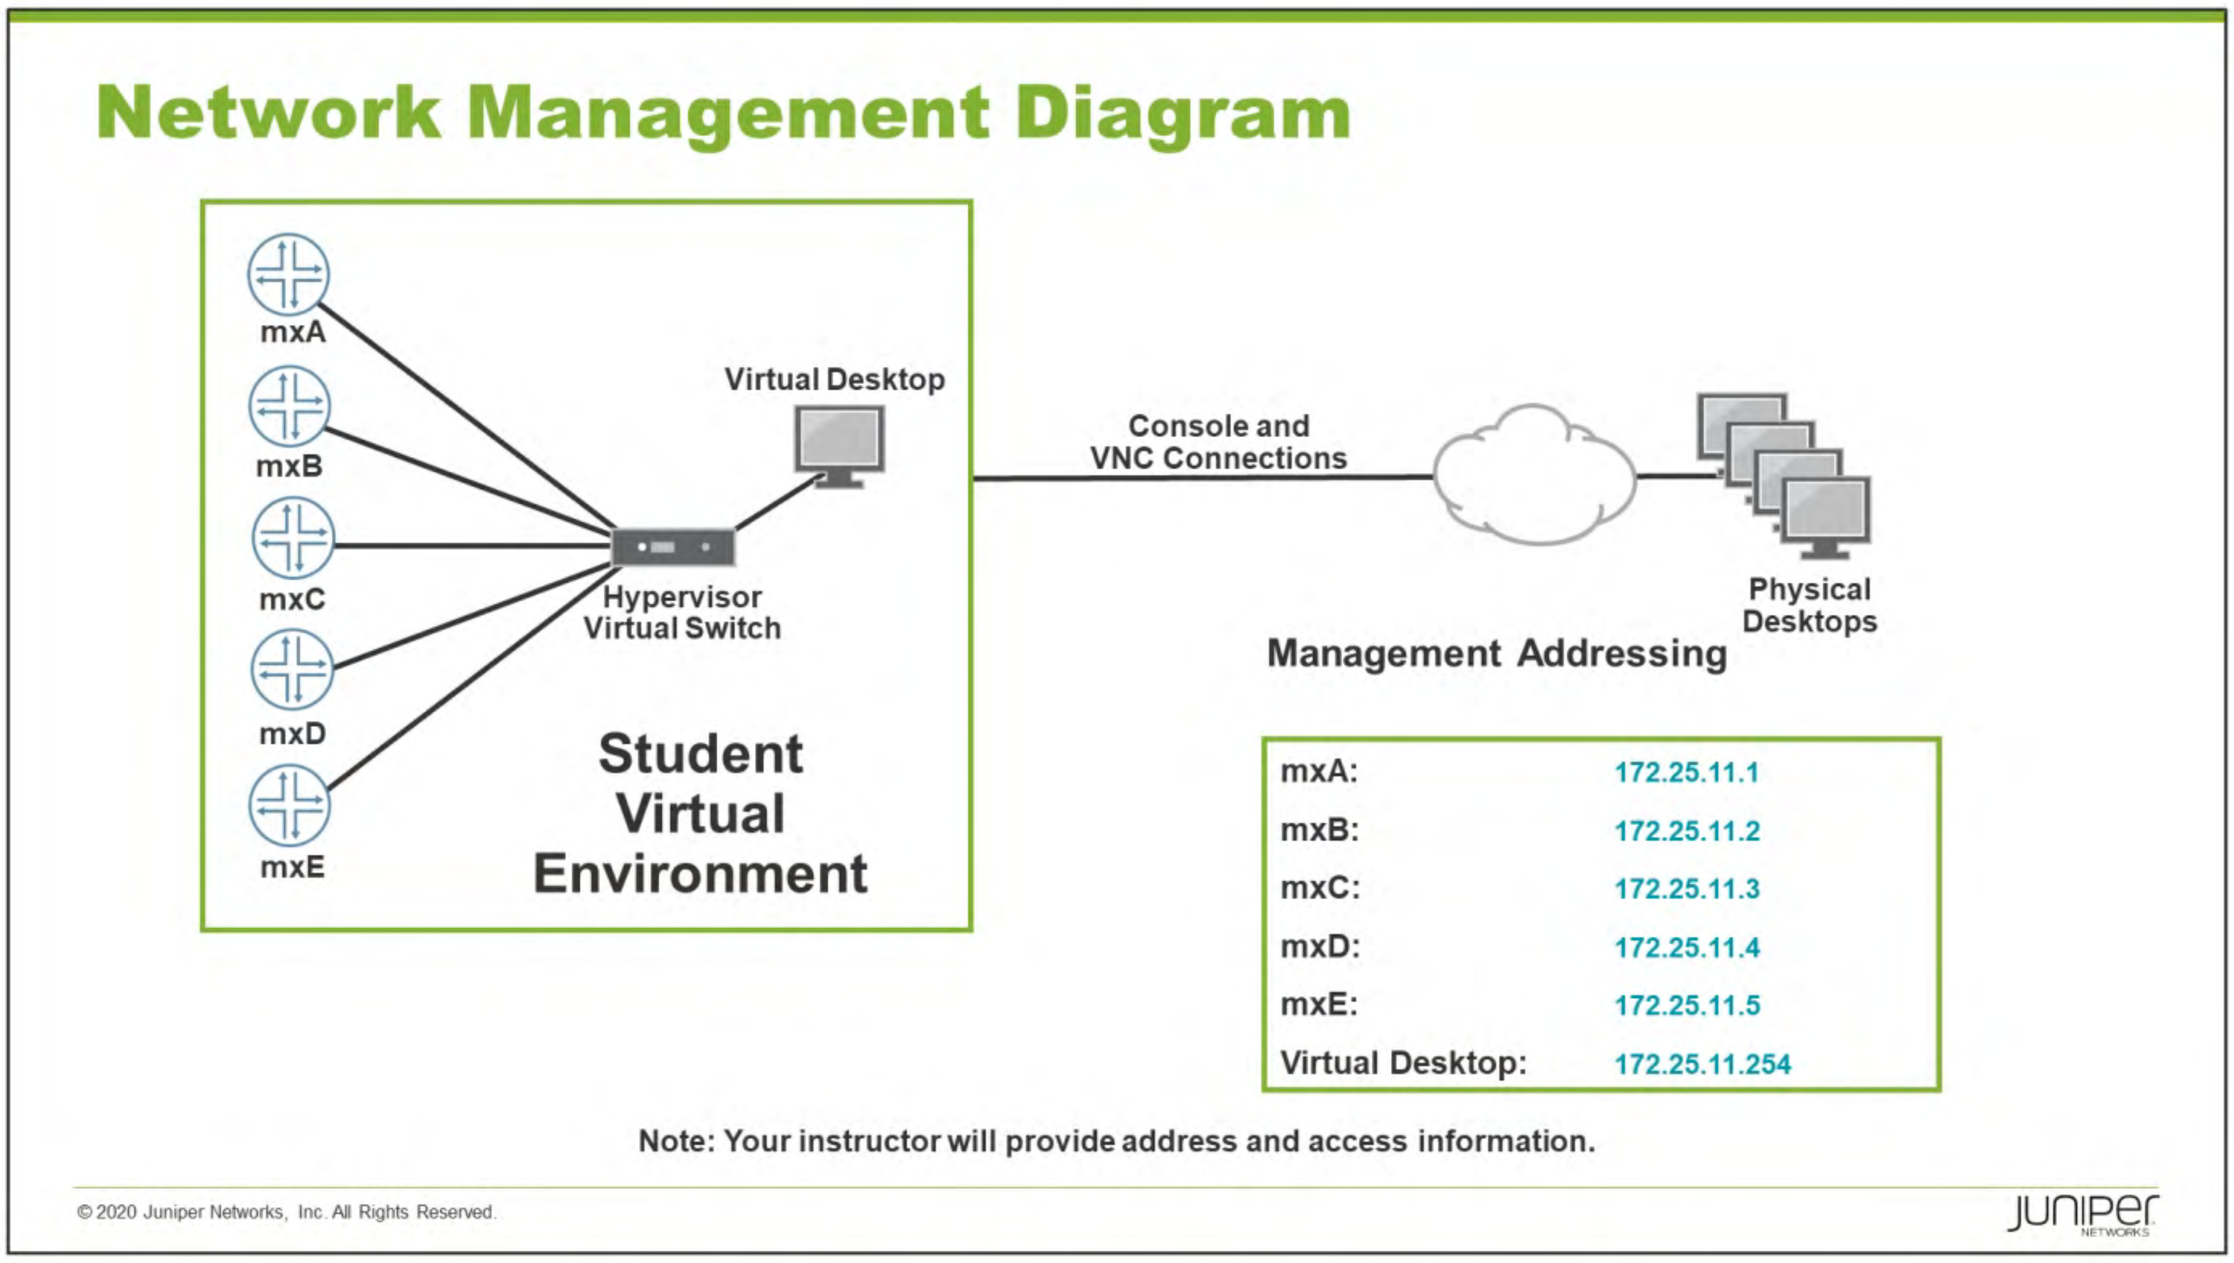
Network Management Diagram
mxA
mxB
mxC
mxD
mxE
Virtual Desktop
Hypervisor
Virtual Switch
Student
Virtual
Environment
Console and
VNC Connections
Physical
Desktops
Management Addressing
mxA: 172.25.11.1
mxB: 172.25.11.2
mxC: 172.25.11.3
mxD: 172.25.11.4
mxE: 172.25.11.5
Virtual Desktop: 172.25.11.254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

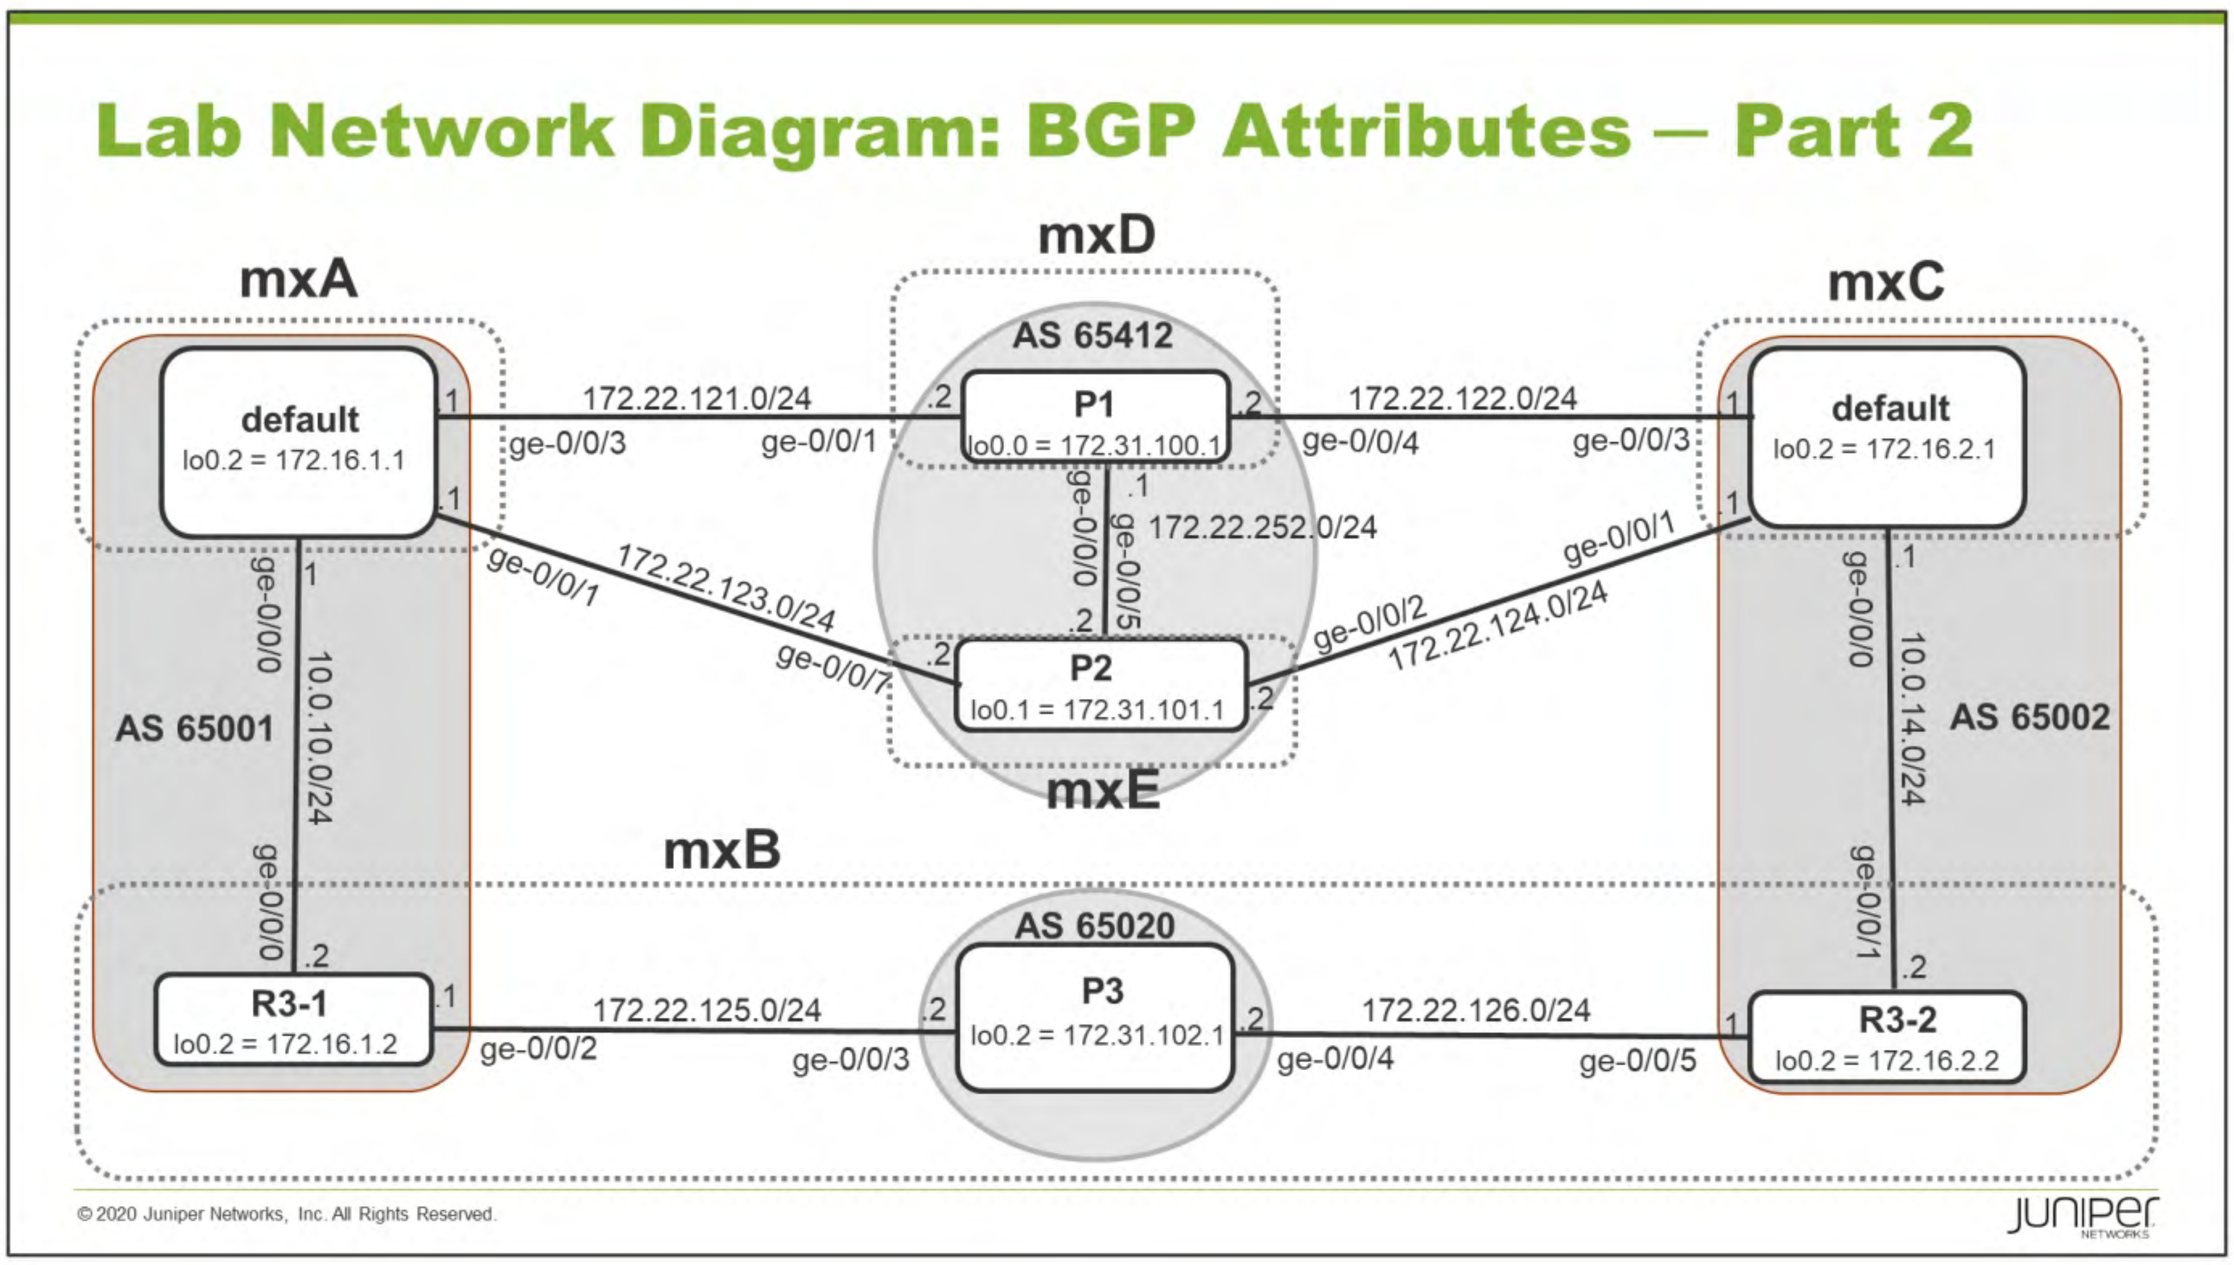
Lab Network Diagram: BGP Attributes – Part 2
mxA
mxD
mxC
mxB
mxE
AS 65412
AS 65001
AS 65002
AS 65020
default
lo0.2 = 172.16.1.1
P1
lo0.0 = 172.31.100.1
default
lo0.2 = 172.16.2.1
P2
lo0.1 = 172.31.101.1
P3
lo0.2 = 172.31.102.1
R3-1
lo0.2 = 172.16.1.2
R3-2
lo0.2 = 172.16.2.2
172.22.121.0/24
172.22.122.0/24
172.22.252.0/24
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
10.0.14.0/24
ge-0/0/3
ge-0/0/1
ge-0/0/4
ge-0/0/3
ge-0/0/0
ge-0/0/5
ge-0/0/1
ge-0/0/1
ge-0/0/7
ge-0/0/2
ge-0/0/0
ge-0/0/0
ge-0/0/1
ge-0/0/2
ge-0/0/3
ge-0/0/4
ge-0/0/5
© 2020 Juniper Networks, Inc. All Rights Reserved.
juniper NETWORKS

Lab

Scaling BGP

Overview

In this lab, you use the lab diagrams titled “Lab Network Diagram: Scaling BGP -- Parts 1-2” and “Lab Network Diagram: Scaling BGP -- Part 3” to configure route reflectors and confederations. Within a local autonomous system topology, the IBGP peers should be fully meshed to prevent routing loops from forming. A fully meshed network inherently has scalability issues which includes the explicit configuration of all IBGP peers with the addition of a new router. Two methods can alleviate the full mesh scaling issue and still ensure a loop-free BGP topology. Route reflection and confederations provide a loop detection mechanism within IBGP without the need for a full mesh of peering sessions and also allow IBGP routes to be readvertised to other IBGP peers.

By completing this lab, you perform the following tasks:

- Load the extended topology.
- Configure route reflection.
- Examine the reflected routes.
- Configure confederations.

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, OSPF neighbor relationships have formed, 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/lab12-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/lab12-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: 2019-12-27 21:44:54 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;
                    }
                }
            }
            ge-0/0/7 {
                unit 0 {
```

```
                    family inet {
                        address 172.22.126.2/24;
                    }
                }
            }
            lo0 {
                unit 0 {
                    family inet {
                        address 172.31.102.1/32;
                        address 20.20.1.1/32;
                        address 40.40.1.1/32;
                    }
                }
            }
        }
        protocols {
            bgp {
                group as65001 {
                    type external;
                    multihop;
                    local-address 172.31.102.1;
                    export ajspr-bgp-export-p3;
                    neighbor 172.16.1.4 {
                        peer-as 65001;
                    }
                }
                group as65002 {
                    type external;
                    multihop;
                    local-address 172.31.102.1;
                    export ajspr-bgp-export-p3;
                    neighbor 172.16.2.4 {
                        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;
    }
}
}
}
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: Based on the configuration, which routers (logical systems) will mxB represent in this lab?

_____

Answer: The mxB device will, using logical systems, represent the P3 router.

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 R5-1 and R5-2 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 any static routes in the 40.40/22 range.

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/lab12-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/lab12-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 IS-IS has been enabled (all interfaces at level 2).

```
lab@mxE> show configuration
## Last commit: 2019-12-27 19:04:24 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 {
            bgp {
                group as65001 {
                    type external;
                    export ajspr-bgp-export;
                    neighbor 172.22.121.1 {
                        peer-as 65001;
                    }
                }
                group as65002 {
                    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;
                }
            }
        }
```

```
    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;
    }
}
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 {
        ospf {
            area 0.0.0.0 {
                interface ge-0/0/6.0;
                interface ge-0/0/7.0;
                interface lo0.1;
            }
        }
        bgp {
```

```
                group P3 {
                    type external;
                    multihop;
                    local-address 172.16.1.4;
                    export export-aggregate;
                    neighbor 172.31.102.1 {
                        peer-as 65020;
                    }
                }
            }
        }
        policy-options {
            policy-statement export-aggregate {
                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;
                }
            }
        }
        routing-options {
            static {
                route 172.31.102.1/32 next-hop 172.22.125.2;
            }
            autonomous-system 65001;
            aggregate {
                route 172.16.1.0/24;
            }
        }
    }
    R5-2 {
        interfaces {
            ge-0/0/1 {
                unit 0 {
                    family inet {
                        address 172.22.126.1/24;
                    }
                }
            }
            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 {
        ospf {
            area 0.0.0.0 {
                interface lo0.2;
                interface ge-0/0/2.0;
                interface ge-0/0/3.0;
            }
        }
        bgp {
            group P3 {
                type external;
                multihop;
                local-address 172.16.2.4;
                export export-aggregate;
                neighbor 172.31.102.1 {
                    peer-as 65020;
                }
            }
            group ibgp {
                type internal;
                local-address 172.16.2.4;
                export next-hop-self;
                neighbor 172.16.2.1;
                neighbor 172.16.2.2;
                neighbor 172.16.2.3;
            }
        }
    }
    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 next-hop-self {
                then {
                    next-hop self;
                }
            }
        }
        routing-options {
            static {
                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 172.25.11.5/24;
            }
        }
    }
}
}
```

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: Which routers are P1 configured to peer with using BGP?

_____

_____

Answer: The P1 router is configured to have an EBGP peering session with the R2-1 and R2-2 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* that is applied to its EBGP peers. This policy causes the P1 router to advertise any static routes in the 30.30/22 range.

Question: Which routers are R5-1 configured to peer with using BGP?

_____

_____

Answer: The R5-1 router is configured for an EBGP session with the P3 router. You will configure the R5-1 router's IBGP sessions in the next parts of this lab.

Question: Which routers are R5-2 configured to peer with using BGP?

_____

_____

Answer: The R5-2 router is configured for an EBGP session with the P3 router and an IBGP session with the R2-2, R3-2, and R4-2 router.

Question: Based on the IBGP configuration of the R5-2 router, how are BGP routing loops being prevented inside AS 65002?

_____

Answer: The routers inside AS 65002 (including the R5-2 router) are configured for a full IBGP mesh. A full IBGP mesh allows all routers in the AS to receive all BGP routes but also prevents routing loops because a router will never readvertise an IBGP-learned route to an IBGP neighbor.

Question: What kind of scaling issue does a full IBGP mesh cause? Use AS 65002 as an example.

_____

_____

Answer: A full IBGP mesh is not a scalable solution in a large environment. For example, if you were to add one more router to AS 65002, the new router as well as R2-2, R3-2, R4-2, and R5-2 would each need to be configured in order to add the new router to the IBGP mesh. Imagine an IBGP mesh of 20 routers, 30 routers, or 100 routers...adding a new router to the mesh would cause you to have to make configuration changes on possibly 101 routers!

Question: What routes will the R5-2 router be advertising to its EBGP peer?

_____

_____

Answer: The R5-2 router is configured with an export policy called *`export-aggregate`* that is applied to its EBGP peer. This policy causes the R5-2 router to advertise the configured aggregate route of 172.16.2.0/24.

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

Issue the **`show bgp summary logical-system R5-2`** command to view the state of the R5-2 router’s EBGP session.

```
lab@mxE> show bgp summary logical-system R5-2
Groups: 2 Peers: 4 Down peers: 3
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       5          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.2.1            65002          0          0       0       0        6:16 Active
172.16.2.2            65002          0          0       0       0        6:16 Active
172.16.2.3            65002          0          0       0       0        6:16 Active
172.31.102.1          65020          8          6       0       0        1:33 Establ
  inet.0: 5/5/5/0
```

Question: What is the state of the R5-2 router’s EBGP neighbor relationship with the P3 router?

_____

Answer: The R5-2 router should have an established EBGP session with the P3 router.

Step 1.8

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/lab12-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/lab12-start.config
load complete
```

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

lab@mxC>
```

Step 1.9

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 commit: 2019-12-06 15:06:19 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;
        }
    }
}
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 {
        ospf {
            area 0.0.0.0 {
                interface ge-0/0/3.0;
                interface ge-0/0/1.0;
                interface lo0.1;
            }
        }
        bgp {
            group ibgp {
                type internal;
                local-address 172.16.2.2;
                neighbor 172.16.2.1;
                neighbor 172.16.2.3;
                neighbor 172.16.2.4;
            }
        }
    }
    routing-options {
        autonomous-system 65002;
    }
}
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 {
                ospf {
                    area 0.0.0.0 {
                        interface ge-0/0/2.0;
                        interface ge-0/0/4.0;
                        interface lo0.2;
                    }
                }
                bgp {
                    group ibgp {
                        type internal;
                        local-address 172.16.2.3;
                        neighbor 172.16.2.1;
                        neighbor 172.16.2.2;
                        neighbor 172.16.2.4;
                    }
                }
            }
            routing-options {
                autonomous-system 65002;
            }
        }
    }
    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.

Question: Based on the configuration, which routers (logical systems) will mxC represent in this lab?

_____

Answer: The mxC device will, using logical systems, represent the R3-2 and R4-2 routers.

Step 1.10

Issue the **`show ospf neighbor logical-system R3-2`** command to ensure the R3-2 router has established an OSPF neighbor relationship with R5-2.

```
lab@mxC> show ospf neighbor logical-system R3-2
Address          Interface              State     ID               Pri  Dead
10.0.16.2        ge-0/0/1.0             Full      172.16.2.4       128    34
```

Question: Has the R3-2 router formed an OSPF neighbor relationship with R5-2?

_____

Answer: Yes, an OSPF neighbor relationship exists between the two routers.

Step 1.11

Issue the **`show ospf neighbor logical-system R4-2`** command to ensure the R4-2 router has established an OSPF neighbor relationship with R5-2.

```
lab@mxC> show ospf neighbor logical-system R4-2
Address          Interface              State     ID               Pri  Dead
10.0.17.2        ge-0/0/2.0             Full      172.16.2.4       128    35
```

Question: Has the R4-2 router formed an OSPF neighbor relationship with R5-2?

_____

Answer: Yes, an OSPF neighbor relationship exists between the two routers.

Step 1.12

Issue the **`show bgp summary logical-system R4-2`** command to view the state of the R4-2 router's IBGP sessions.

```
lab@mxC> show bgp summary logical-system R4-2
Groups: 1 Peers: 3 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       5          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.2.1            65002          0          0       0       0        4:58 Active
172.16.2.2            65002          7          5       0       0        2:12 Establ
  inet.0: 0/0/0/0
172.16.2.4            65002          9          5       0       0        2:10 Establ
  inet.0: 4/5/5/0 0/0/0/0
```

Question: What is the state of the R4-2 router's IBGP neighbor relationships?

_____

Answer: The R4-2 router should have an established IBGP session with the R3-2 and R5-2 routers. The IBGP session with the R2-2 router should be in the active or connect state because the R2-2 router has not been configured for BGP yet.

Step 1.13

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/lab12-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/lab12-start.config
load complete

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

lab@mxD>
```

Step 1.14

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 15:08:17 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 {
            ospf {
                area 0.0.0.0 {
                    interface lo0.1;
                    interface ge-0/0/1.0;
                    interface ge-0/0/2.0;
                }
            }
            bgp {
                group P1 {
                    type external;
                    export export-aggregate;
                    neighbor 172.22.121.2 {
                        peer-as 65412;
                    }
                }
            }
        }
        policy-options {
            policy-statement export-aggregate {
                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;
                }
```

```
            }
        }
        routing-options {
            autonomous-system 65001;
            aggregate {
                route 172.16.1.0/24;
            }
        }
    }
    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 {
            ospf {
                area 0.0.0.0 {
                    interface lo0.2;
                    interface ge-0/0/4.0;
                    interface ge-0/0/5.0;
                }
            }
            bgp {
```

```
                group P1 {
                    type external;
                    export export-aggregate;
                    neighbor 172.22.122.2 {
                        peer-as 65412;
                    }
                }
                group ibgp {
                    type internal;
                    local-address 172.16.2.1;
                    export next-hop-self;
                    neighbor 172.16.2.2;
                    neighbor 172.16.2.3;
                    neighbor 172.16.2.4;
                }
            }
        }
        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 next-hop-self {
                from neighbor 172.22.122.2;
                then {
                    next-hop self;
                }
            }
        }
        routing-options {
            autonomous-system 65002;
            aggregate {
                route 172.16.2.0/24;
            }
        }
    }
}
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 R2-1 and R2-2 routers.

Question: Which routers are R2-1 configured to peer with using BGP?

_____

_____

Answer: The R2-1 router is configured for an EBGP session with the P1 router. You will configure the R2-1 router's IBGP sessions in the next parts of this lab.

Question: Which routers are R2-2 configured to peer with using BGP?

_____

_____

Answer: The R2-2 router is configured for an EBGP session with the P1 router and an IBGP session with the R5-2, R3-2, and R4-2 router.

Step 1.15

Issue the **`show ospf neighbor logical-system`** ***`R2-2`*** command to ensure the R2-2 router has established an OSPF neighbor relationship with the R3-2 and R4-2 routers.

```
lab@mxD> show ospf neighbor logical-system R2-2
Address          Interface              State     ID               Pri  Dead
10.0.14.2        ge-0/0/4.0             Full      172.16.2.2       128    34
10.0.15.2        ge-0/0/5.0             Full      172.16.2.3       128    35
```

Question: Has the R2-2 router formed an OSPF neighbor relationship with the R3-2 and R4-2 routers?

_____

Answer: Yes, an OSPF neighbor relationship exists between the other routers.

Step 1.16

Issue the **`show bgp summary logical-system R2-2`** command to view the state of the R4-2 router's IBGP sessions.

```
lab@mxD> show bgp summary logical-system R2-2
Threading mode: BGP I/O
Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      10          9          0          0          0          0
Peer                     AS      InPkt     OutPkt     OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.2.2            65002          8          9       0       0        2:57 Establ
  inet.0: 0/0/0/0
172.16.2.3            65002          8         10       0       0        2:59 Establ
  inet.0: 0/0/0/0
172.16.2.4            65002         11         10       0       0        2:57 Establ
  inet.0: 4/5/5/0
172.22.122.2          65412         13         12       0       0        3:44 Establ
  inet.0: 5/5/5/0
```

Question: What is the state of the R2-2 router's BGP neighbor relationships?

_____

Answer: The R2-2 router should have an established IBGP session with the R3-2, R4-2, and R5-2 routers and an established EBGP session with the P1 router.

Step 1.17

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/lab12-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/lab12-start.config
```

```
load complete

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

lab@mxA>
```

Step 1.18

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 15:09:10 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 {
                ospf {
                    area 0.0.0.0 {
                        interface lo0.1;
                        interface ge-0/0/3.0;
                        interface ge-0/0/2.0;
                    }
                }
            }
            routing-options {
                autonomous-system 65001;
            }
        }
        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 {
        ospf {
            area 0.0.0.0 {
                interface lo0.2;
                interface ge-0/0/1.0;
                interface ge-0/0/4.0;
            }
        }
    }
    routing-options {
        autonomous-system 65001;
    }
}
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.

Question: Based on the configuration, which routers (logical systems) will mxA represent in this lab?

_____

_____

Answer: The mxC device will, using logical systems, represent the R3-1 and R4-1 routers.

Step 1.19

Issue the **show ospf neighbor logical-system *R3-1*** command to ensure the R3-1 router has established an OSPF neighbor relationship with R2-1 and R5-1.

```
lab@mxA> show ospf neighbor logical-system R3-1
Address          Interface              State     ID               Pri  Dead
10.0.12.2        ge-0/0/2.0             Full      172.16.1.4       128    36
10.0.10.1        ge-0/0/3.0             Full      172.16.1.1       128    34
```

Question: Has the R3-1 router formed an OSPF neighbor relationship with R2-1 and R5-1?

_____

_____

Answer: Yes, an OSPF neighbor relationship exists between the routers.

Step 1.20

Issue the **show ospf neighbor logical-system *R4-1*** command to ensure the R4-1 router has established an OSPF neighbor relationship with R2-1 and R5-1.

```
lab@mxA> show ospf neighbor logical-system R4-1
Address          Interface              State     ID               Pri  Dead
10.0.13.2        ge-0/0/1.0             Full      172.16.1.4       128    38
10.0.11.1        ge-0/0/4.0             Full      172.16.1.1       128    37
```

Question: Has the R4-1 router formed an OSPF neighbor relationship with R2-1 and R5-1?

_____

_____

Answer: Yes, an OSPF neighbor relationship exists between the routers.

Step 1.21

Issue the **show bgp summary logical-system *R3-1*** command and the the **show bgp summary logical-system *R4-1*** command to view the state of the R3-1 and R4-1 routers' IBGP sessions.

```
lab@mxA> show bgp summary logical-system R3-1
BGP is not running
```

```
lab@mxA> show bgp summary logical-system R4-1
BGP is not running

lab@mxA>
```

Question: What is the state of the R3-1 and R4-1 routers' IBGP neighbor relationships? Why?

_____

_____

Answer: The two routers do not have any BGP neighbor relationships because BGP has not been configured on the two routers yet. You will configured the IBGP sessions for the R3-1 and R4-1 routers later in this lab.

Part 2: Configuring Route Reflection

In this lab part, you configure BGP route reflectors. A route reflector utilizes two new BGP attributes. These attributes are never advertised outside the local autonomous system and are used internally for IBGP loop detection. The cluster-list is the first new BGP attribute and operates like the AS path attribute. It contains a list of 32-bit cluster IDs for each cluster a particular route has transited. If a route reflector detects its cluster ID in the cluster-list, it is considered a loop and the BGP update is dropped. The second attribute is the originator ID, which defines the router that first advertised the route to the route reflector. The route reflector uses the originator ID as a second check against routing loops.

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 bgp group rr-cluster]` hierarchy. Configure the R2-1 router as a route reflector for the internal BGP network, and configure the R3-1, R4-2, and R5-1 loopback addresses as neighbors. Use the R2-1 loopback address as the cluster ID and the local-address within the route reflector group. Commit the configuration when finished.

```
lab@mxD:R2-1> configure
Entering configuration mode

[edit]
lab@mxD:R2-1# edit protocols bgp group rr-cluster

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# set type internal
```

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# set local-address 172.16.1.1

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# set cluster 172.16.1.1

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# set neighbor 172.16.1.2

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# set neighbor 172.16.1.3

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# set neighbor 172.16.1.4

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# commit
commit complete

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1#
```

Step 2.3

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 2.4

Enter configuration mode and navigate to the `[edit protocols bgp group ibgp]` hierarchy and configure the R2-1 router's loopback address as an IBGP neighbor. Use the R3-1 router's loopback address as the `local-address` for the internal BGP network. Commit the configuration and exit to operational mode when finished.

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

[edit]
lab@mxA:R3-1# edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxA:R3-1# set type internal

[edit protocols bgp group ibgp]
lab@mxA:R3-1# set local-address 172.16.1.2

[edit protocols bgp group ibgp]
lab@mxA:R3-1# set neighbor 172.16.1.1

[edit protocols bgp group ibgp]
```

```
lab@mxA:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R3-1>
```

Step 2.5

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 2.6

Enter configuration mode and navigate to the `[edit protocols bgp group ibgp]` hierarchy and configure the R2-1 router's loopback address as an IBGP neighbor. Use the R4-1 loopback address as the `local-address` for the internal BGP group. Commit the configuration and exit to operational mode when finished.

```
lab@mxA:R4-1> configure
Entering configuration mode

[edit]
lab@mxA:R4-1# edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxA:R4-1# set type internal

[edit protocols bgp group ibgp]
lab@mxA:R4-1# set local-address 172.16.1.3

[edit protocols bgp group ibgp]
lab@mxA:R4-1# set neighbor 172.16.1.1

[edit protocols bgp group ibgp]
lab@mxA:R4-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R4-1>
```

Step 2.7

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 2.8

Enter configuration mode and navigate to the `[edit protocols bgp group ibgp]` hierarchy and configure the R2-1 router's loopback address as an IBGP neighbor. Use the R5-1 loopback address as the `local-address` for the internal BGP group. Commit the configuration when completed.

```
lab@mxE:R5-1> configure
Entering configuration mode

[edit]
lab@mxE:R5-1# edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxE:R5-1# set type internal

[edit protocols bgp group ibgp]
lab@mxE:R5-1# set local-address 172.16.1.4

[edit protocols bgp group ibgp]
lab@mxE:R5-1# set neighbor 172.16.1.1

[edit protocols bgp group ibgp]
lab@mxE:R5-1# commit
commit complete

[edit protocols bgp group ibgp]
lab@mxE:R5-1#
```

Step 2.9

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

From the SSH session with the mxD router, issue the command **`run show bgp summary`** to ensure that all of the IBGP peer sessions are established.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# run show bgp summary
Threading mode: BGP I/O
Groups: 2 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      10          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2            65001          2          4       0       0          24 Establ
  inet.0: 0/0/0/0
172.16.1.3            65001          2          4       0       0          20 Establ
  inet.0: 0/0/0/0
172.16.1.4            65001          4          4       0       0          16 Establ
  inet.0: 0/5/5/0
172.22.121.2          65412         28         27       0       0       11:02 Establ
  inet.0: 5/5/5/0
```

Question: Are all of the IBGP peers established on the route reflector?

_____

_____

Answer: All of the IBGP and EBGP peering sessions are established. If any of your sessions are not established, contact your instructor.

Step 2.10

Use the **run show route hidden** command to check for any unusable routes in the default routing table.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# run show route hidden

inet.0: 23 destinations, 24 routes (19 active, 0 holddown, 5 hidden)
+ = Active Route, - = Last Active, * = Both

40.40.0.0/24           [BGP/170] 00:01:53, localpref 100, from 172.16.1.4
                         AS path: 65020 I, validation-state: unverified
                         Unusable
40.40.1.0/24           [BGP/170] 00:01:53, localpref 100, from 172.16.1.4
                         AS path: 65020 I, validation-state: unverified
                         Unusable
40.40.2.0/24           [BGP/170] 00:01:53, localpref 100, from 172.16.1.4
                         AS path: 65020 I, validation-state: unverified
                         Unusable
40.40.3.0/24           [BGP/170] 00:01:53, localpref 100, from 172.16.1.4
                         AS path: 65020 I, validation-state: unverified
                         Unusable
172.16.2.0/24          [BGP/170] 00:01:53, localpref 100, from 172.16.1.4
                         AS path: 65020 65002 I, validation-state: unverified
                         Unusable

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: Are any unusable routes in the default routing table?

_____

_____

Answer: Yes, five unusable routes are present.

Step 2.11

Use the **run show route 40.40/24 hidden detail** command to display detailed route information and a possible cause for the route to be marked unusable.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# run show route 40.40/24 hidden detail
```

```
inet.0: 23 destinations, 24 routes (19 active, 0 holddown, 5 hidden)
40.40.0.0/24 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Next hop type: Unusable, Next hop index: 0
                Address: 0xc0d82d0
                Next-hop reference count: 5
                State: <Hidden Int Ext Changed>
                Local AS: 65001 Peer AS: 65001
                Age: 1:16
                Validation State: unverified
                Task: BGP_65001.172.16.1.4
                AS path: 65020 I
                Accepted
                Localpref: 100
                Router ID: 172.16.1.4
```

Question: Does the `detail` switch in the `show route` command display any evidence as to why the route is unusable?

_____

_____

Answer: No information in the detailed output helps determine the reason for the route's unusable status.

Step 2.12

Use the **`run show route 40.40/24 hidden extensive`** to display the extensive information and a possible cause for the route to be marked unusable.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# run show route 40.40/24 hidden extensive


inet.0: 23 destinations, 24 routes (19 active, 0 holddown, 5 hidden)
40.40.0.0/24 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Next hop type: Unusable, Next hop index: 0
                Address: 0xc0d82d0
                Next-hop reference count: 5
                State: <Hidden Int Ext Changed>
                Local AS: 65001 Peer AS: 65001
                Age: 2:20
                Validation State: unverified
                Task: BGP_65001.172.16.1.4
                AS path: 65020 I
                Accepted
                Localpref: 100
                Router ID: 172.16.1.4
                Indirect next hops: 1
                        Protocol next hop: 172.31.102.1
                        Indirect next hop: 0x0 - INH Session ID: 0x0
```

Question: Does the extensive output help resolve the problem with the unusable routes?

_____

_____

Answer: The protocol next-hop attribute in the BGP update cannot be resolved in the routing table. A next-hop self policy must be applied as an IBGP import policy on the R5-1 router to fix this problem on the R2-1 router.

Step 2.13

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 policy-options policy-statement next-hop-self]` hierarchy and create a policy that modifies the next-hop attribute to the local loopback address.

```
[edit protocols bgp group ibgp]
lab@mxE:R5-1# top edit policy-options policy-statement next-hop-self

[edit policy-options policy-statement next-hop-self]
lab@mxE:R5-1# set term 1 then next-hop self

[edit policy-options policy-statement next-hop-self]
lab@mxE:R5-1#
```

Step 2.14

Navigate to the `[edit protocols bgp group ibgp]` hierarchy and apply the *next-hop-self* policy as an export policy in the group. Commit the configuration when finished.

```
[edit policy-options policy-statement next-hop-self]
lab@mxE:R5-1# top edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxE:R5-1# set export next-hop-self

[edit protocols bgp group ibgp]
lab@mxE:R5-1# commit
commit complete

[edit protocols bgp group ibgp]
lab@mxE:R5-1#
```

Step 2.15

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 hidden`** command to determine if there are any hidden routes.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# run show route hidden

inet.0: 23 destinations, 24 routes (23 active, 0 holddown, 0 hidden)

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: Do any hidden routes exist in the R2-1 router's routing table?

_____

_____

Answer: No, there are no hidden routes being learned by the R2-1 router. By applying the *next-hop-self* policy on the R5-1 router, the hidden route problem was fixed on the R2-1 router.

Step 2.16

Issue the **run traceroute 40.40.1.1 source 172.16.1.1** command to make sure that the P3 router is reachable with no routing loops from the R2-1 router.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# run traceroute 40.40.1.1 source 172.16.1.1
traceroute to 40.40.1.1 (40.40.1.1) from 172.16.1.1, 30 hops max, 52 byte packets
 1  10.0.10.2 (10.0.10.2)  1.715 ms  1.642 ms  1.375 ms
 2  10.0.12.2 (10.0.12.2)  1.978 ms  1.824 ms  1.831 ms
 3  40.40.1.1 (40.40.1.1)  3.287 ms  2.655 ms  2.847 ms
```

Question: Do any hidden routes exist in the R2-1 router's routing table?

_____

_____

Answer: No, there are no hidden routes being learned by the R2-1 router. By applying the *next-hop-self* policy on the R5-1 router, the hidden route problem was fixed on the R2-1 router.

Step 2.17

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

From the SSH session with the mxE router, issue the **show route hidden** command to see if there are any hidden routes learned by the R5-1 router.

```
[edit protocols bgp group ibgp]
lab@mxE:R5-1# run show route hidden

inet.0: 24 destinations, 25 routes (20 active, 0 holddown, 5 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
30.30.0.0/24          [BGP/170] 00:15:57, localpref 100, from 172.16.1.1
                        AS path: 65412 I, validation-state: unverified
                        Unusable
30.30.1.0/24          [BGP/170] 00:15:57, localpref 100, from 172.16.1.1
                        AS path: 65412 I, validation-state: unverified
                        Unusable
30.30.2.0/24          [BGP/170] 00:15:57, localpref 100, from 172.16.1.1
                        AS path: 65412 I, validation-state: unverified
                        Unusable
30.30.3.0/24          [BGP/170] 00:15:57, localpref 100, from 172.16.1.1
                        AS path: 65412 I, validation-state: unverified
                        Unusable
172.16.2.0/24         [BGP/170] 00:15:57, localpref 100, from 172.16.1.1
                        AS path: 65412 65002 I, validation-state: unverified
                        Unusable

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: Are the any hidden routes learned by the R5-1 router?

_____

_____

Answer: Yes, there are five hidden routes being learned by the R5-1 router.

Question: How can a hidden route problem on the R5-1 router be solved?

_____

_____

Answer: Hidden routes are usually IBGP learned routes in which the BGP next hop cannot be resolved in the routing table. To solve this problem, a next-hop self policy should be applied on the remote IBGP peers, the R2-1 router in this case.

Step 2.18

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 policy-options policy-statement next-hop-self]` hierarchy and create a policy that modifies the next-hop attribute to the local loopback address.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# top edit policy-options policy-statement next-hop-self

[edit policy-options policy-statement next-hop-self]
lab@mxD:R2-1# set term 1 then next-hop self

[edit policy-options policy-statement next-hop-self]
lab@mxD:R2-1#
```

Step 2.19

Navigate to the `[edit protocols bgp group rr-cluster]` hierarchy and apply the *`next-hop-self`* policy as an export policy in the group. Commit the configuration when completed.

```
[edit policy-options policy-statement next-hop-self]
lab@mxD:R2-1# top edit protocols bgp group rr-cluster

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# set export next-hop-self

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# commit
commit complete

[edit protocols bgp group rr-cluster]
lab@mxD:R2-1#
```

Step 2.20

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

From the SSH session with the mxE router, issue the **`show route hidden`** command to see if there are any hidden routes learned by the R5-1 router.

```
[edit protocols bgp group ibgp]
lab@mxE:R5-1# run show route hidden

inet.0: 24 destinations, 25 routes (24 active, 0 holddown, 0 hidden)

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: Are the any hidden routes learned by the R5-1 router?

_____

_____

Answer: No, there are no more hidden routes being learned by the R5-1 router.

Step 2.21

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 40.40/22`** command to display the P3 routes in the default routing table.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# run show route 40.40/22

inet.0: 23 destinations, 24 routes (23 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

40.40.0.0/24       *[BGP/170] 00:17:29, localpref 100, from 172.16.1.4
                      AS path: 65020 I, validation-state: unverified
```

```
                    > to 10.0.10.2 via ge-0/0/1.0
                      to 10.0.11.2 via ge-0/0/2.0
40.40.1.0/24       *[BGP/170] 00:17:29, localpref 100, from 172.16.1.4
                      AS path: 65020 I, validation-state: unverified
                    > to 10.0.10.2 via ge-0/0/1.0
                      to 10.0.11.2 via ge-0/0/2.0
40.40.2.0/24       *[BGP/170] 00:17:29, localpref 100, from 172.16.1.4
                      AS path: 65020 I, validation-state: unverified
                    > to 10.0.10.2 via ge-0/0/1.0
                      to 10.0.11.2 via ge-0/0/2.0
40.40.3.0/24       *[BGP/170] 00:17:29, localpref 100, from 172.16.1.4
                      AS path: 65020 I, validation-state: unverified
                    > 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 P3 routes active in the default routing table?

_____

_____

Answer: Yes, the P3 routes are now active in the default routing table.

Step 2.22

Use the **`run traceroute 40.40.1.1 source 172.16.1.1 ttl 10`** command to verify connectivity to the 40.40.1.1 address on the P3 router. Be sure to source the traceroute from the R2-1 router's loopback address.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# run traceroute 40.40.1.1 source 172.16.1.1 ttl 10
traceroute to 40.40.1.1 (40.40.1.1) from 172.16.1.1, 10 hops max, 52 byte packets
 1  10.0.10.2 (10.0.10.2)  1.434 ms  1.348 ms  1.135 ms
 2  10.0.10.1 (10.0.10.1)  1.157 ms  1.228 ms  1.151 ms
 3  10.0.10.2 (10.0.10.2)  1.639 ms  1.826 ms  2.050 ms
 4  10.0.10.1 (10.0.10.1)  1.845 ms  1.739 ms  1.776 ms
 5  10.0.10.2 (10.0.10.2)  2.258 ms  2.412 ms  2.387 ms
 6  10.0.10.1 (10.0.10.1)  2.711 ms  2.583 ms  2.432 ms
 7  10.0.10.2 (10.0.10.2)  2.950 ms  2.842 ms  2.792 ms
 8  10.0.10.1 (10.0.10.1)  2.720 ms  3.110 ms  3.093 ms
 9  10.0.10.2 (10.0.10.2)  3.716 ms  3.313 ms  3.371 ms
10  10.0.10.1 (10.0.10.1)  3.464 ms  3.764 ms  3.825 ms
```

Question: Does a problem exist in the path to the P3 router?

_____

_____

Answer: A routing loop has been formed between the R2-1 router and the R3-1 and R4-1 routers. The loop is caused by the *next-hop-self* export policy on the route reflector. When configuring export policies in the route reflector, the match condition must be very specific to only change attributes from the EBGP peers. Right now the *next-hop-self* policy is changing the next hop of the reflected routes being advertised by R5-1.

Step 2.23

Navigate to the [edit policy-options policy-statement *next-hop-self*] hierarchy. Modify the policy to change only the next-hop attribute if the BGP routes are external routes. Use the match condition route-type to accomplish this task. Commit the change when completed.

```
[edit protocols bgp group rr-cluster]
lab@mxD:R2-1# top edit policy-options policy-statement next-hop-self

[edit policy-options policy-statement next-hop-self]
lab@mxD:R2-1# set term 1 from protocol bgp

[edit policy-options policy-statement next-hop-self]
lab@mxD:R2-1# set term 1 from route-type external

[edit policy-options policy-statement next-hop-self]
lab@mxD:R2-1# commit
commit complete

[edit policy-options policy-statement next-hop-self]
lab@mxD:R2-1#
```

Step 2.24

Use the **run traceroute 40.40.1.1 source 172.16.1.1 ttl 10** command to verify connectivity to the 40.40.1.1 address on the P3 router. Be sure to source the traceroute from your default router's loopback address.

```
[edit policy-options policy-statement next-hop-self]
lab@mxD:R2-1# run traceroute 40.40.1.1 source 172.16.1.1 ttl 10
traceroute to 40.40.1.1 (40.40.1.1) from 172.16.1.1, 10 hops max, 52 byte packets
 1  10.0.10.2 (10.0.10.2)  1.609 ms  1.007 ms  1.048 ms
 2  10.0.12.2 (10.0.12.2)  1.685 ms  1.690 ms  1.561 ms
 3  40.40.1.1 (40.40.1.1)  2.890 ms  2.685 ms  2.610 ms
```

Question: Is the traceroute reaching the P3 router?

_____

_____

Answer: Yes. The *next-hop-self* policy change is only changing EBGP routes and not the routes being reflected by the route reflector. Your traceroute path might vary.

Step 2.25

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

From the SSH session with the mxA router, use the **show route 40.40/24 detail** command to display the P3 route in the R4-1 router's routing table.

```
lab@mxA:R4-1> show route 40.40/24 detail

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
40.40.0.0/24 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xc0d8a3c
                Next-hop reference count: 8
                Source: 172.16.1.1
                Next hop type: Router, Next hop index: 707
                Next hop: 10.0.13.2 via ge-0/0/1.0, selected
                Session Id: 0x168
                Protocol next hop: 172.16.1.4
                Indirect next hop: 0xc2dfb04 1048575 INH Session ID: 0x16f
                State: <Active Int Ext>
                Local AS: 65001 Peer AS: 65001
                Age: 30         Metric2: 1
                Validation State: unverified
                Task: BGP_65001.172.16.1.1
                Announcement bits (2): 2-KRT 4-Resolve tree 1
                AS path: 65020 I  (Originator)
                Cluster list:  172.16.1.1
                Originator ID: 172.16.1.4
                Accepted
                Localpref: 100
                Router ID: 172.16.1.1
```

Question: The P3 route is being reflected to the R3-1 and R4-1 routers. What is the value of the cluster list and the originator ID?

_____

_____

Answer: In the example output, the cluster list is 172.16.1.1, which is the cluster ID configured in the BGP cluster group. The originator ID is 172.16.1.4, which is the peer address of the R4-1 router.

Step 2.26

Use the **show route 30.30/24 detail** command to display the P1 route in the R4-1 router's routing table.

```
lab@mxA:R4-1> show route 30.30/24 detail


inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
30.30.0.0/24 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xc0d8bcc
                Next-hop reference count: 10
                Source: 172.16.1.1
                Next hop type: Router, Next hop index: 706
                Next hop: 10.0.11.1 via ge-0/0/4.0, selected
                Session Id: 0x169
                Protocol next hop: 172.16.1.1
                Indirect next hop: 0xc2dfc84 1048577 INH Session ID: 0x16d
                State: <Active Int Ext>
                Local AS: 65001 Peer AS: 65001
                Age: 2:43       Metric2: 1
                Validation State: unverified
                Task: BGP_65001.172.16.1.1
                Announcement bits (2): 2-KRT 4-Resolve tree 1
                AS path: 65412 I
                Accepted
                Localpref: 100
                Router ID: 172.16.1.1
```

Question: Is this route being advertised or reflected to the R4 router?

_____

_____

Answer: Because this route does not have a cluster list or an originator ID, it is being advertised to the R4-1 router and not reflected.

Part 3: Configuring Confederations

In this lab part, you configure a scalable IBGP network using confederations. BGP confederations allows you to break the network up into smaller pieces called a sub-AS or a member-AS network. Each sub-AS is assigned a unique AS number and operates as an independent internal network that must follow the IBGP rules, requiring a full mesh or route reflection topology. Connectivity between the sub-AS networks uses a modified form of EBGP named confederation BGP (CBGP). CBGP peers prepend the sub-AS number to the AS path attribute as routing updates are exchanged, which is used to prevent routing loops while the routes are advertised throughout the IBGP network. Remember to switch to the "Network Diagram: Scaling BGP Lab, Part 3" diagram.

Step 3.1

Enter configuration mode and navigate to the `[edit protocols bgp]` hierarchy. Delete the *ibgp* peer group. Commit the configuration and exit to operational mode when finished.

```
lab@mxA:R4-1> configure
Entering configuration mode

[edit]
lab@mxA:R4-1# edit protocols bgp

[edit protocols bgp]
lab@mxA:R4-1# delete group ibgp

[edit protocols bgp]
lab@mxA:R4-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R4-1>
```

Step 3.2

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 3.3

Enter configuration mode and navigate to the `[edit protocols bgp]` hierarchy. Delete the *ibgp* peer group. Commit the configuration and exit to operational mode when finished.

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

[edit]
lab@mxA:R3-1# edit protocols bgp

[edit protocols bgp]
lab@mxA:R3-1# delete group ibgp

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

lab@mxA:R3-1>
```

Step 3.4

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 protocols bgp]` hierarchy. Delete the *ibgp* peer group. Commit the configuration when finished.

```
[edit protocols bgp group ibgp]
```

```
lab@mxE:R5-1# up

[edit protocols bgp]
lab@mxE:R5-1# delete group ibgp

[edit protocols bgp]
lab@mxE:R5-1# commit
commit complete

[edit protocols bgp]
lab@mxE:R5-1#
```

Step 3.5

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 protocols bgp]` hierarchy. Delete the *`rr-cluster`* peer group. Commit the configuration when finished.

```
[edit policy-options policy-statement next-hop-self]
lab@mxD:R2-1# top edit protocols bgp

[edit protocols bgp]
lab@mxD:R2-1# delete group rr-cluster

[edit protocols bgp]
lab@mxD:R2-1# commit
commit complete

[edit protocols bgp]
lab@mxD:R2-1#
```

Step 3.6

Navigate to the `[edit routing-options]` hierarchy. Set the autonomous system number, confederation global autonomous system number, and the two member AS numbers as described in the following table. Refer to the lab diagram as necessary.

Router	Autonomous System	Confederation Number	Members
R2-1, R3-1	65101	65001	65101, 65201
R4-1, R5-1	65201	65001	65101, 65201

```
[edit protocols bgp]
lab@mxD:R2-1# top edit routing-options

[edit routing-options]
lab@mxD:R2-1# set autonomous-system 65101

[edit routing-options]
lab@mxD:R2-1# set confederation 65001 members 65101
```

```
[edit routing-options]
lab@mxD:R2-1# set confederation 65001 members 65201

[edit routing-options]
lab@mxD:R2-1#
```

Step 3.7

Navigate to the `[edit protocols bgp group ibgp]` hierarchy and configure an internal peer session to the R3-1 router using the loopback addresses. Also apply the *`next-hop-self`* policy as an export group policy. Commit the configuration when finished.

```
[edit routing-options]
lab@mxD:R2-1# top edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxD:R2-1# set type internal

[edit protocols bgp group ibgp]
lab@mxD:R2-1# set local-address 172.16.1.1

[edit protocols bgp group ibgp]
lab@mxD:R2-1# set neighbor 172.16.1.2

[edit protocols bgp group ibgp]
lab@mxD:R2-1# set export next-hop-self

[edit protocols bgp group ibgp]
lab@mxD:R2-1# commit
commit complete

[edit protocols bgp group ibgp]
lab@mxD:R2-1#
```

Step 3.8

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

From the SSH session with the mxA router, enter configuration mode and navigate to the `[edit routing-options]` hierarchy. Set the autonomous system number, the confederation global autonomous system number, and the two member AS numbers using the values from the table shown earlier. Refer to the lab diagram as necessary.

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

[edit]
lab@mxA:R3-1# edit routing-options

[edit routing-options]
lab@mxA:R3-1# set autonomous-system 65101

[edit routing-options]
lab@mxA:R3-1# set confederation 65001 members 65101

[edit routing-options]
lab@mxA:R3-1# set confederation 65001 members 65201
```

```
[edit routing-options]
lab@mxA:R3-1#
```

Step 3.9

Navigate to the `[edit protocols bgp group ibgp]` hierarchy and configure an internal peer session to the default router using loopback addresses. Commit the configuration and exit to operational mode when completed.

```
[edit routing-options]
lab@mxA:R3-1# top edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxA:R3-1# set type internal

[edit protocols bgp group ibgp]
lab@mxA:R3-1# set local-address 172.16.1.2

[edit protocols bgp group ibgp]
lab@mxA:R3-1# set neighbor 172.16.1.1

[edit protocols bgp group ibgp]
lab@mxA:R3-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R3-1>
```

Step 3.10

Use the **`show bgp summary`** command to display the BGP peer sessions.

```
lab@mxA:R3-1> 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
                       5          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.1            65101          4          2       0       0          6 Establ
  inet.0: 5/5/5/0
```

Question: Is the IBGP peer session between the R2-1 router and the R3-1 router established?

_____

_____

Answer: Yes, the IBGP peer session between the R2-1 router and the R3-1 router is established. If the peer is not established, check the configuration or contact your instructor.

Question: How many routes are being received from the R2-1 router by the R3-1 router?

_____

_____

Answer: The R3-1 router is receiving 5 active routes from the R2-1 router.

Step 3.11

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 3.12

Enter configuration mode and navigate to the `[edit routing-options]` hierarchy. Set the autonomous system number, confederation global autonomous system number, and the two member AS numbers as described in the table shown earlier. Refer to the lab diagram as necessary.

```
lab@mxA:R4-1> configure
Entering configuration mode

[edit]
lab@mxA:R4-1# edit routing-options

[edit routing-options]
lab@mxA:R4-1# set autonomous-system 65201

[edit routing-options]
lab@mxA:R4-1# set confederation 65001 members 65101

[edit routing-options]
lab@mxA:R4-1# set confederation 65001 members 65201

[edit routing-options]
lab@mxA:R4-1#
```

Step 3.13

Navigate to the `[edit protocols bgp group ibgp]` hierarchy and configure an internal peer session to the R5-1 router using loopback addresses. Commit your configuration when finished.

```
[edit routing-options]
lab@mxA:R4-1# top edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxA:R4-1# set type internal
```

```
[edit protocols bgp group ibgp]
lab@mxA:R4-1# set local-address 172.16.1.3

[edit protocols bgp group ibgp]
lab@mxA:R4-1# set neighbor 172.16.1.4

[edit protocols bgp group ibgp]
lab@mxA:R4-1# commit
commit complete

[edit protocols bgp group ibgp]
lab@mxA:R4-1#
```

Step 3.14

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 routing-options]` hierarchy. Set the autonomous system number, the confederation global autonomous system number, and the two member AS numbers using the values from the table shown earlier. Refer to the lab diagram as necessary.

```
[edit protocols bgp]
lab@mxE:R5-1# top edit routing-options

[edit routing-options]
lab@mxE:R5-1# set autonomous-system 65201

[edit routing-options]
lab@mxE:R5-1# set confederation 65001 members 65101

[edit routing-options]
lab@mxE:R5-1# set confederation 65001 members 65201

[edit routing-options]
lab@mxE:R5-1#
```

Step 3.15

Navigate to the `[edit protocols bgp group ibgp]` hierarchy and configure an internal peer session to the R4-1 router using the loopback addresses. Also configure the next-hop-self policy as a group export policy. Commit the configuration when completed.

```
[edit routing-options]
lab@mxE:R5-1# top edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxE:R5-1# set type internal

[edit protocols bgp group ibgp]
lab@mxE:R5-1# set local-address 172.16.1.4

[edit protocols bgp group ibgp]
lab@mxE:R5-1# set neighbor 172.16.1.3

[edit protocols bgp group ibgp]
lab@mxE:R5-1# set export next-hop-self
```

```
[edit protocols bgp group ibgp]
lab@mxE:R5-1# commit
commit complete

[edit protocols bgp group ibgp]
lab@mxE:R5-1#
```

Step 3.16

Use the **`run show bgp summary`** command to verify that the peer session between the R4-1 router and the R5-1 router is established.

```
[edit protocols bgp group ibgp]
lab@mxE:R5-1# 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
                       5          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.3            65201          2          4       0       0          4 Establ
  inet.0: 0/0/0/0
172.31.102.1          65020          6          3       0       0          8 Establ
  inet.0: 5/5/5/0
```

Question: Are the BGP peer sessions established to the R4-1 router and the P3 router?

_____

_____

Answer: Both the EBGP session to P3 and the IBGP session to R4-1 are established. If the peering sessions are not established, check your configuration and consult with your instructor.

Step 3.17

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 R5-1 router is advertising routes to the R4-1 peer.

```
[edit protocols bgp group ibgp]
lab@mxA:R4-1# 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
                       5          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.4            65201          6          2       0       0         31 Establ
  inet.0: 5/5/5/0
```

Question: Is the R4-1 router receiving routes from the R5-1 peer? How many of the routes are active?

_____

_____

Answer: The R4-1 router is receiving five active routes from the R5-1 peer.

Step 3.18

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 protocols bgp group cbgp]` hierarchy. Configure a confederation BGP session between the R2-1 router and the R5-1 router. Because the CBGP session uses loopback addresses to peer, both the **`local-address`** and **`multihop`** commands are required to establish the session. Commit the configuration when you are finished.

```
[edit protocols bgp group ibgp]
lab@mxD:R2-1# top edit protocols bgp group cbgp

[edit protocols bgp group cbgp]
lab@mxD:R2-1# set type external

[edit protocols bgp group cbgp]
lab@mxD:R2-1# set multihop

[edit protocols bgp group cbgp]
lab@mxD:R2-1# set local-address 172.16.1.1

[edit protocols bgp group cbgp]
lab@mxD:R2-1# set neighbor 172.16.1.4

[edit protocols bgp group cbgp]
lab@mxD:R2-1# set peer-as 65201

[edit protocols bgp group cbgp]
lab@mxD:R2-1# commit
commit complete

[edit protocols bgp group cbgp]
lab@mxD:R2-1#
```

Step 3.19

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 protocols bgp group cbgp]` hierarchy and configure the confederation BGP session to the default router. Commit the configuration when you are finished.

```
[edit protocols bgp group ibgp]
lab@mxE:R5-1# top edit protocols bgp group cbgp

[edit protocols bgp group cbgp]
lab@mxE:R5-1# set type external
```

```
[edit protocols bgp group cbgp]
lab@mxE:R5-1# set multihop

[edit protocols bgp group cbgp]
lab@mxE:R5-1# set local-address 172.16.1.4

[edit protocols bgp group cbgp]
lab@mxE:R5-1# set neighbor 172.16.1.1

[edit protocols bgp group cbgp]
lab@mxE:R5-1# set peer-as 65101

[edit protocols bgp group cbgp]
lab@mxE:R5-1# commit
commit complete

[edit protocols bgp group cbgp]
lab@mxE:R5-1#
```

Step 3.20

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 bgp summary`** command to display the bgp peering sessions on the R2-1 router.

```
[edit protocols bgp group cbgp]
lab@mxD:R2-1# run show bgp summary
Threading mode: BGP I/O
Groups: 3 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      10          5          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2            65101         15         17       0       0        5:43 Establ
  inet.0: 0/0/0/0
172.16.1.4            65201          4          4       0       0           9 Establ
  inet.0: 0/5/5/0
172.22.121.2           65412         23         18       0       0        6:55 Establ
  inet.0: 5/5/5/0
```

Question: Does the default router have an established CBGP peer session with the R5-1 router?

_____

_____

Answer: Yes, the CBGP peer session with R5-1 is established. If the peer session is not established, check the configuration and consult with your instructor.

Question: How many routes are being received from the R5-1 router? How many routes are active from R5-1?

_____

_____

Answer: Five routes being are received from the R5-1 router, however none of the routes are active.

Step 3.21

Use the **run show route hidden** command to display any hidden routes.

```
[edit protocols bgp group cbgp]
lab@mxD:R2-1# run show route hidden

inet.0: 23 destinations, 24 routes (19 active, 0 holddown, 5 hidden)
+ = Active Route, - = Last Active, * = Both

40.40.0.0/24          [BGP/170] 00:05:27, localpref 100, from 172.16.1.4
                        AS path: (65201) 65020 I, validation-state: unverified
                        Unusable
40.40.1.0/24          [BGP/170] 00:05:27, localpref 100, from 172.16.1.4
                        AS path: (65201) 65020 I, validation-state: unverified
                        Unusable
40.40.2.0/24          [BGP/170] 00:05:27, localpref 100, from 172.16.1.4
                        AS path: (65201) 65020 I, validation-state: unverified
                        Unusable
40.40.3.0/24          [BGP/170] 00:05:27, localpref 100, from 172.16.1.4
                        AS path: (65201) 65020 I, validation-state: unverified
                        Unusable
172.16.2.0/24         [BGP/170] 00:05:27, localpref 100, from 172.16.1.4
                        AS path: (65201) 65020 65002 I, validation-state: unverified
                        Unusable

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: Are there any hidden routes?

_____

_____

Answer: Yes, there are five received routes that are hidden.

Question: Why are the routes received from the R5-1 router marked as unusable?

_____

_____

Answer: The `next-hop` attribute in the routes being received from the R5-1 router cannot be resolved in the R2-1 router's routing table. CBGP peer sessions do not change any of the BGP attributes except the AS path.

Step 3.22

Configure the *`next-hop-self`* policy as the group export policy. Commit the configuration and return to operational mode.

```
[edit protocols bgp group cbgp]
lab@mxD:R2-1# set export next-hop-self

[edit protocols bgp group cbgp]
lab@mxD:R2-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxD:R2-1>
```

Step 3.23

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

From the SSH session with the mxE router, configure the *`next-hop-self`* policy as the group export policy. Commit the configuration and return to operational mode.

```
[edit protocols bgp group cbgp]
lab@mxE:R5-1# set export next-hop-self

[edit protocols bgp group cbgp]
lab@mxE:R5-1# commit and-quit
commit complete
Exiting configuration mode

lab@mxE:R5-1>
```

Step 3.24

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 hidden`** command to display the routes that are hidden.

```
lab@mxD:R2-1> show route hidden

inet.0: 23 destinations, 24 routes (23 active, 0 holddown, 0 hidden)

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: Are there any hidden routes?

_____

_____

Answer: No, there are no more hidden routes.

Step 3.25

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

```
lab@mxD:R2-1> exit

mxD (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

Network Management Diagram

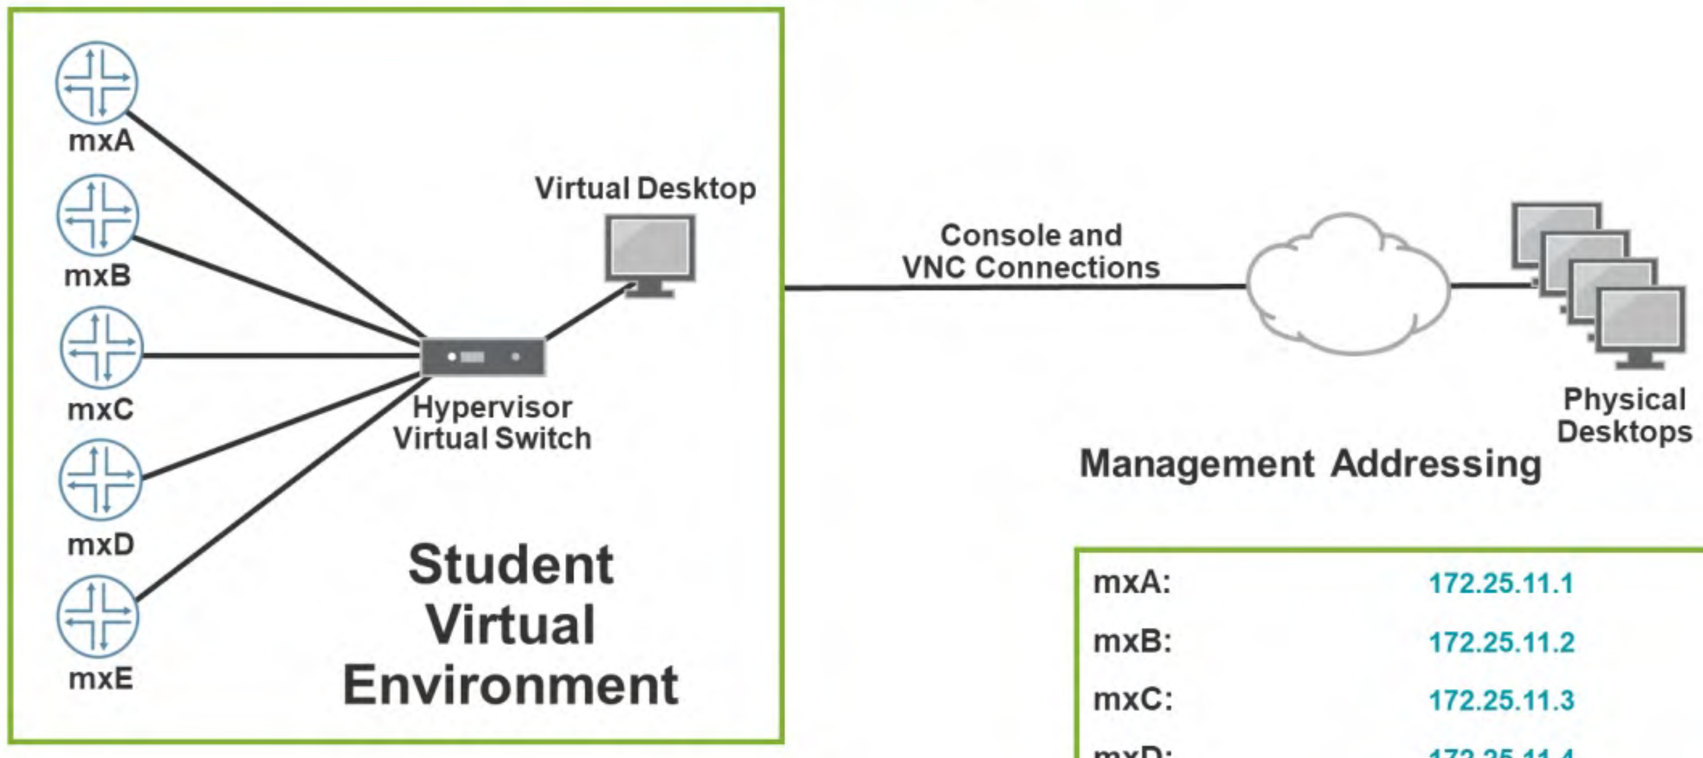

Management Addressing

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: Scaling BGP – Parts 1-2

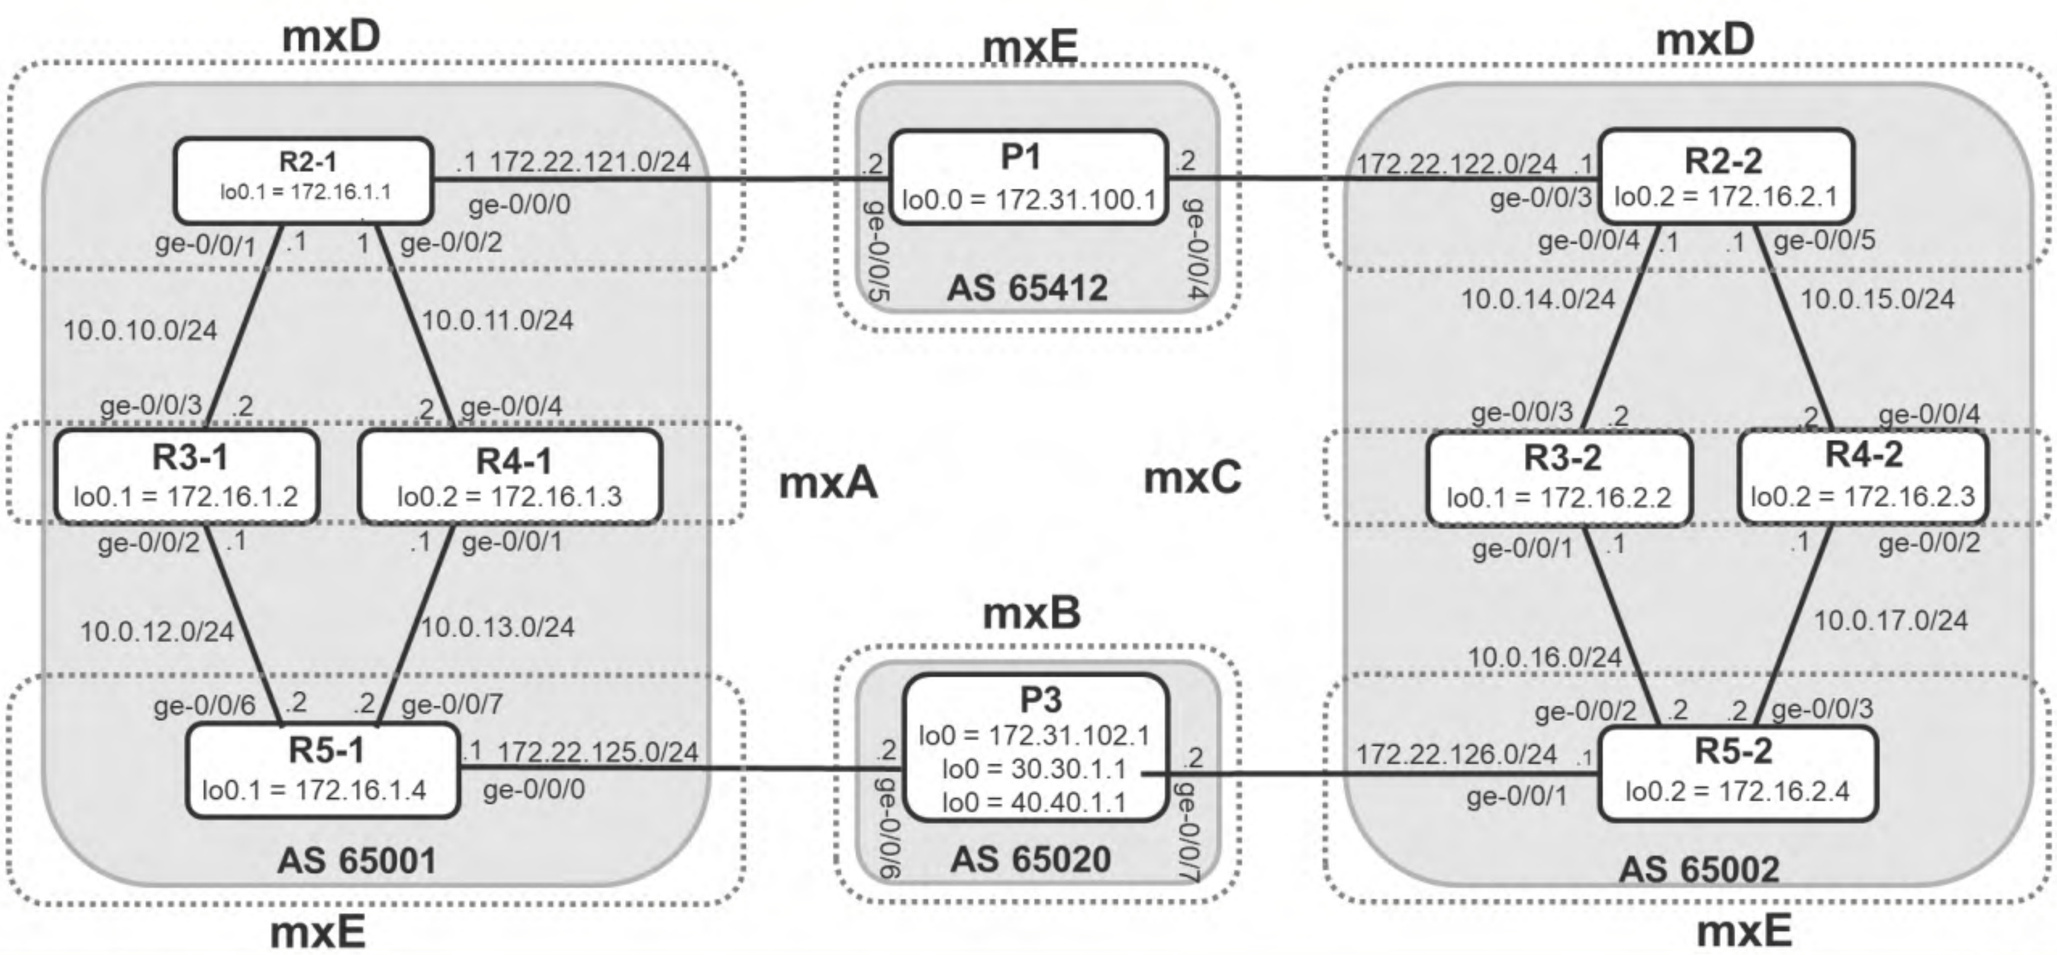

© 2020 Juniper Networks, Inc. All Rights Reserved.

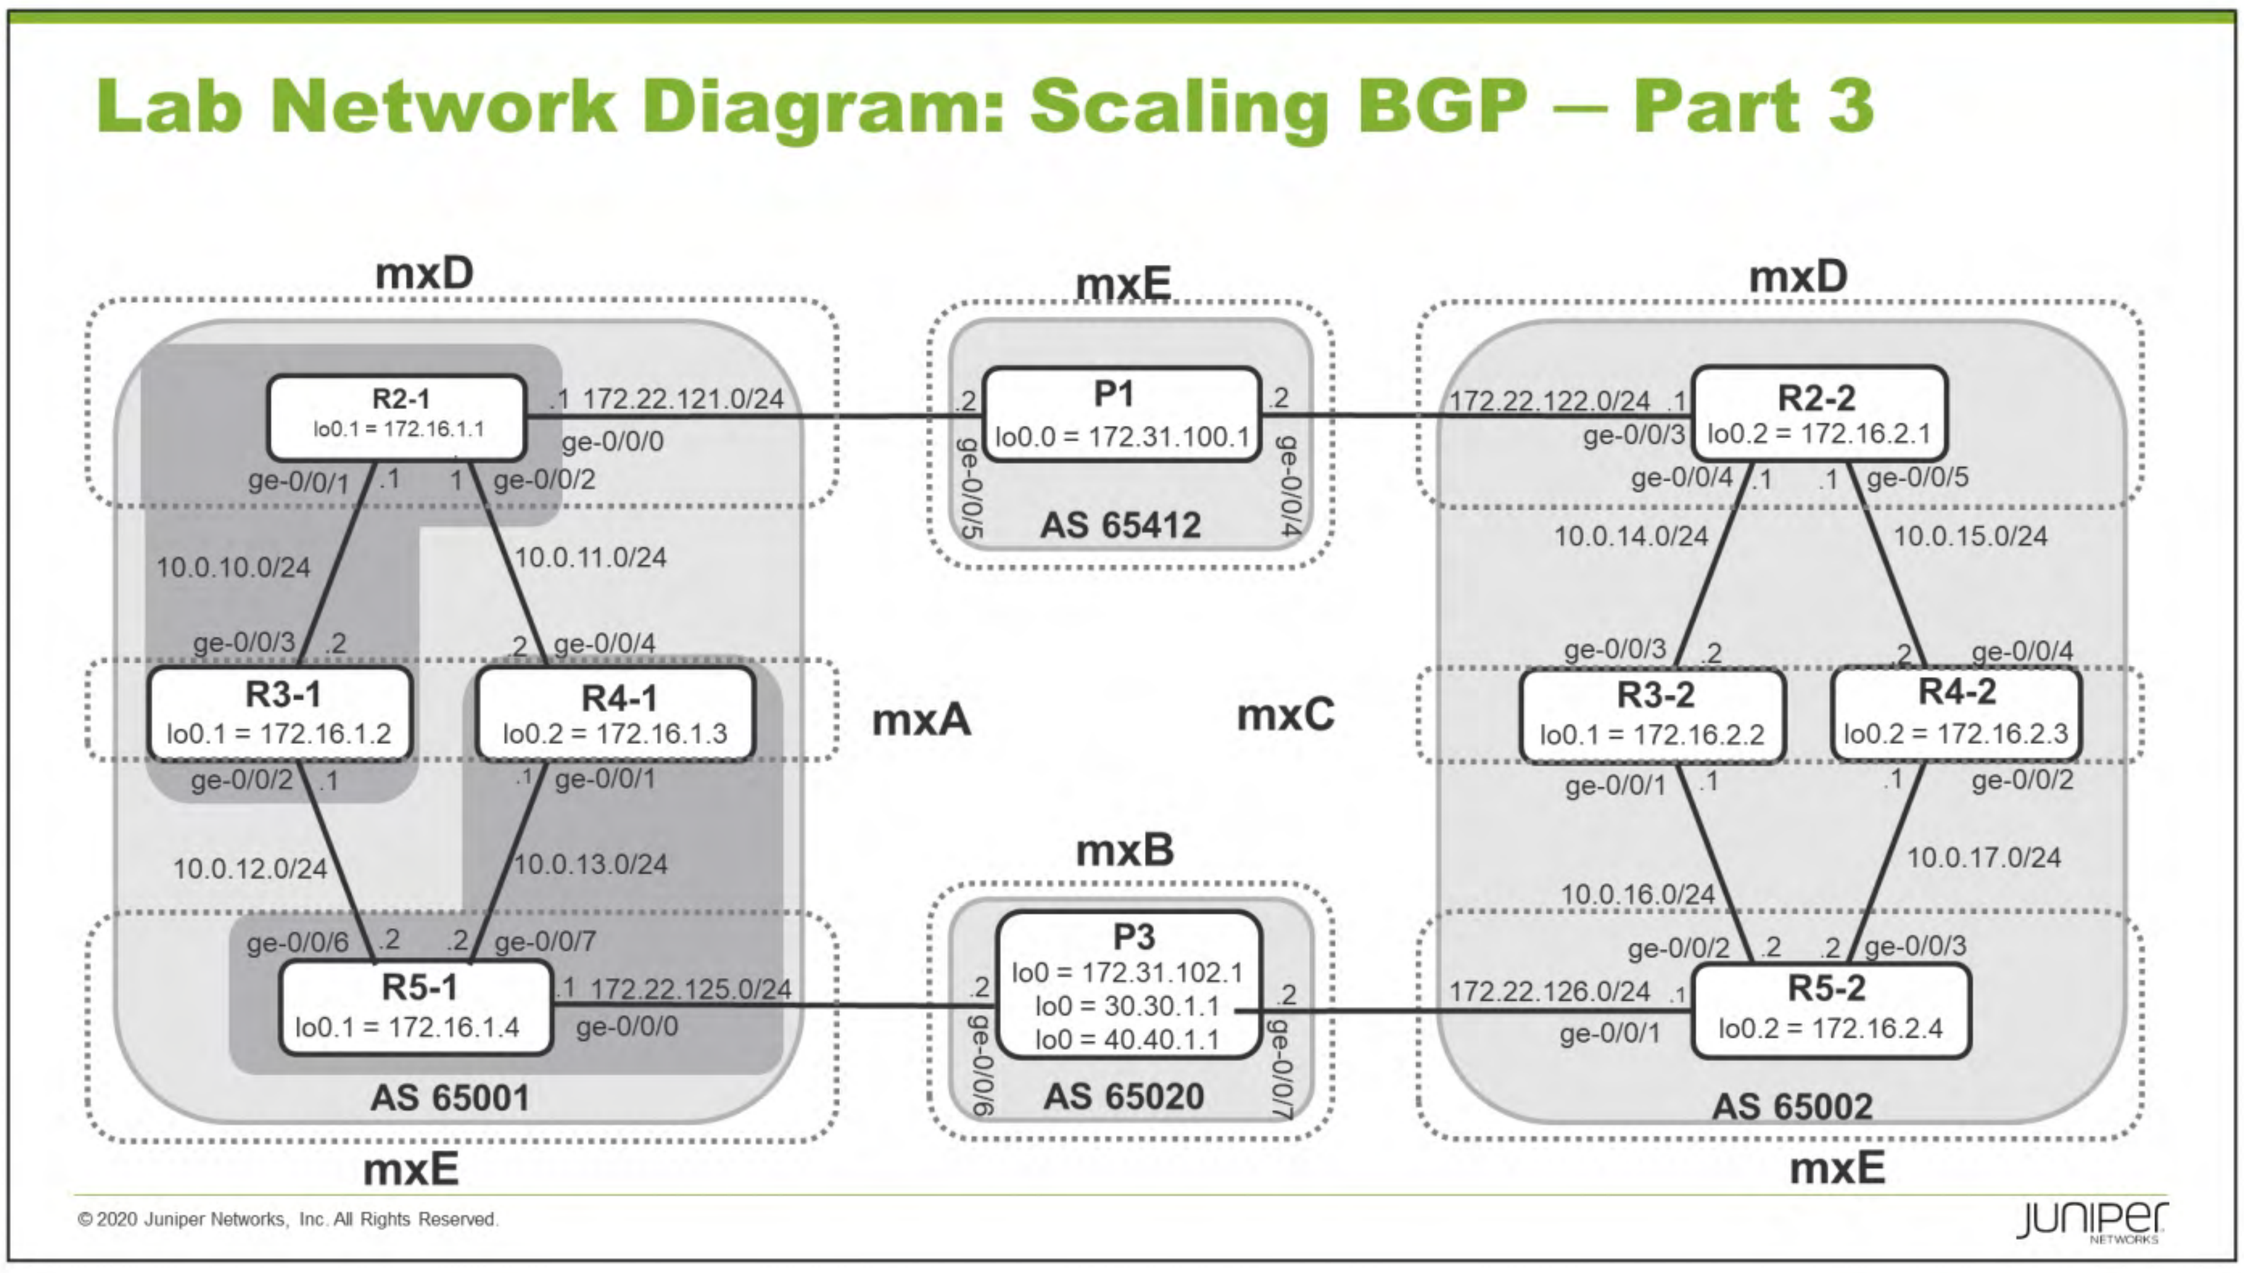
Lab Network Diagram: Scaling BGP – Part 3
mxD
R2-1
lo0.1 = 172.16.1.1
.1 172.22.121.0/24
ge-0/0/0
ge-0/0/1
.1
1
ge-0/0/2
10.0.10.0/24
10.0.11.0/24
ge-0/0/3
.2
2
ge-0/0/4
R3-1
lo0.1 = 172.16.1.2
R4-1
lo0.2 = 172.16.1.3
ge-0/0/2
.1
.1
ge-0/0/1
10.0.12.0/24
10.0.13.0/24
ge-0/0/6
.2
.2
ge-0/0/7
R5-1
lo0.1 = 172.16.1.4
.1 172.22.125.0/24
ge-0/0/0
AS 65001
mxE
mxE
P1
lo0.0 = 172.31.100.1
.2
.2
ge-0/0/5
ge-0/0/4
AS 65412
mxA
mxC
mxB
P3
lo0 = 172.31.102.1
lo0 = 30.30.1.1
lo0 = 40.40.1.1
.2
.2
ge-0/0/6
ge-0/0/7
AS 65020
mxD
172.22.122.0/24 .1
R2-2
lo0.2 = 172.16.2.1
ge-0/0/3
ge-0/0/4
.1
.1
ge-0/0/5
10.0.14.0/24
10.0.15.0/24
ge-0/0/3
.2
2
ge-0/0/4
R3-2
lo0.1 = 172.16.2.2
R4-2
lo0.2 = 172.16.2.3
ge-0/0/1
.1
.1
ge-0/0/2
10.0.16.0/24
10.0.17.0/24
ge-0/0/2
.2
.2
ge-0/0/3
172.22.126.0/24 .1
R5-2
lo0.2 = 172.16.2.4
ge-0/0/1
AS 65002
mxE
© 2020 Juniper Networks, Inc. All Rights Reserved.
JUNIPER NETWORKS

Lab

BGP FlowSpec

Overview

In this lab, you use the lab diagram titled “Lab Network Diagram: BGP FlowSpec” to block DNS traffic that is being used to attack a Web Server located at the Customer site. In the first part of the lab you will add a flow specification to the customer device and send that information to the service provider. The service provider will then forward that information to the mxC-R3 device which creates a filter to block the Attackers traffic based on the flow data. In the next portion of the lab, the service provider will be responsible for adding the flow specification to the RR and have the flow routes distributed to mxA-R1, and mxC-R3 devices to block the DNS attack.

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

- Configure a Customer edge router to forward flow specification data to a service provider edge device
- Configure a service provider network to receive flow specification routes from a customer, verify that the routes are within the assigned customer address range and forward the flow information to all IBGP peers.
- Monitor and verify the edge devices that they are blocking the correct traffic and the attack no longer exists
- Configure a RR in the service provider network to distribute flow specification information to IBGP peers.
- Monitor and verify the edge devices receive the routes and create the required firewall filters.

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.

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 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/lab13-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: Wed Nov 27 17:20:28 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/lab13-start.config
load complete

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

lab@mxA-R1>
```

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/lab13-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 27 17:20:28 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/lab13-start.config
load complete

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

lab@mxB-R2>
```

Step 1.4

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/lab13-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 27 17:21:22 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/lab13-start.config
load complete

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

lab@mxC-R3>
```

Step 1.5

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/lab13-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: Wed Nov 27 17:21:45 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/lab13-start.config
load complete

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

lab@mxE-Attacker>
```

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/lab13-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 27 15:27:40 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/lab13-start.config
load complete

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

lab@mxD-Customer>
```

Part 2: Configuring Customer-based FlowSpec

In this lab part, you configure the Customer device (mxD) to create a flow specification which will be advertised to the service provider network via the mxA-R1 EBGP peering.

Step 2.1

Use the **show bgp summary** command to view the BGP neighbor relationship on the Customer router.

```
lab@mxD-Customer> 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...
172.22.131.10         64512          2          2       0       0          16 Establ
  inet.0: 0/0/0/0
```

Question: How many neighbors does the mxD-Customer router have configured?

_____

_____

Answer: You should see one BGP neighbor that has been configured on the Customer router.

Question: What is the status of the BGP session?

_____

_____

Answer: The BGP session in the `Established` state.

Step 2.2

Use the **`show bgp neighbor 172.22.131.38 | match NLRI`** command to examine the EBGP session that is currently in an `Established` state. Follow this command with the **`show configuration protocols bgp`** command.

```
lab@mxD-Customer> show bgp neighbor 172.22.131.10 | match NLRI
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast inet-flow
  NLRI for this session: inet-unicast
  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

lab@mxD-Customer> show configuration protocols bgp
group SP {
    type external;
    peer-as 64512;
    neighbor 172.22.131.10;
}
```

Question: Which NLRIs are being advertised by the BGP peer?

_____

_____

Answer: The peer is advertising the inet-unicast and inet-flow NLRIs.

Question: What NLRI is being used for this session?

_____

_____

Answer: The NLRI used in the session is inet-unicast.

Question: Why do you think that only the inet-unicast NLRI is being used?

_____

_____

Answer: the family unicast-flow has not been added to the BGP configuration.

Step 2.3

Enter configuration mode and navigate to the `[edit protocols bgp]` hierarchy. Add **`family inet flow`** to the BGP group configuration. When satisfied with the changes, perform a **`commit`**.

```
lab@mxD-Customer> configure
Entering configuration mode

[edit]
lab@mxD-Customer# edit protocols bgp

[edit protocols bgp]
lab@mxD-Customer# set group SP family inet flow

[edit protocols bgp]
lab@mxD-Customer# commit
commit complete

[edit protocols bgp]
lab@mxD-Customer#
```

Step 2.4

Use the **`run show bgp neighbor 172.22.131.10 | match NLRI`** command again to examine the EBGP session.

```
[edit protocols bgp]
lab@mxD-Customer# run show bgp neighbor 172.22.131.10 | match NLRI
  NLRI for restart configured on peer: inet-flow
  NLRI advertised by peer: inet-unicast inet-flow
  NLRI for this session: inet-flow
  NLRI that restart is negotiated for: inet-flow
  NLRI of received end-of-rib markers: inet-flow
  NLRI of all end-of-rib markers sent: inet-flow

[edit protocols bgp]
lab@mxD-Customer#
```

Question: What NLRIs have been selected for this session now?

_____

_____

Answer: Only the inet-flow has been selected for this session.

Question: Why is inet-unicast not included?

_____

_____

Answer: Once an address family is defined, all addresses have to be defined for them to be included it in the session.

Step 2.5

In this step you will add **family inet unicast** back into the configuration. When you are satisfied with the changes, **commit** the changes.

```
[edit protocols bgp]
lab@mxD-Customer# set group SP family inet unicast

[edit protocols bgp]
lab@mxD-Customer# commit
commit complete
```

Step 2.6

Use the **run show bgp neighbor 172.22.131.10 | match NLRI** command again to examine the EBGP session.

```
[edit protocols bgp]
lab@mxD-Customer# run show bgp neighbor 172.22.131.10 | match NLRI
  NLRI for restart configured on peer: inet-unicast inet-flow
  NLRI advertised by peer: inet-unicast inet-flow
  NLRI for this session: inet-unicast inet-flow
  NLRI that restart is negotiated for: inet-unicast inet-flow
  NLRI of received end-of-rib markers: inet-unicast inet-flow
  NLRI of all end-of-rib markers sent: inet-unicast inet-flow

[edit protocols bgp]
lab@mxD-Customer#
```

Question: Are both NLRIs successfully negotiated?

_____

_____

Answer: Yes, both NLRIs have been successfully negotiated. If not, reverify your work, or work with your Instructor.

Step 2.7

Before the flow specification is configured, look at the existing `routing-options` configuration using the **`show routing-options`** command.

```
[edit protocols bgp]
lab@mxD-Customer# top show routing-options
static {
    route 203.0.113.0/24 discard;
    route 0.0.0.0/0 next-hop 172.22.131.10;
}
autonomous-system 65000;
```

Question: What is the purpose of the two static routes?

_____

_____

Answer: The 0/0 route is the default route for Internet access. The 203.0.113.0/24 represents the internal network that belongs to the Customer. Any route within the range of the routes owned by the customer can be advertised within the flow specification.

Step 2.8

In this lab step you will configure BGP FlowSpec to report that address 203.0.113.1 is being attacked by multiple sources using the UDP protocol and port 53. This traffic should be discarded at the edge of the service provider network. Once the flow specification has been created, review the configuration, and apply the changes using the **`commit and-quit`** command.

```
[edit protocols bgp]
lab@mxD-Customer# top edit routing-options flow

[edit routing-options flow]
lab@mxD-Customer# set route DNS match protocol udp

[edit routing-options flow]
lab@mxD-Customer# set route DNS match port 53

[edit routing-options flow]
lab@mxD-Customer# set route DNS match destination 203.0.113.1

[edit routing-options flow]
lab@mxD-Customer# set route DNS then discard

[edit routing-options flow]
lab@mxD-Customer# show
route DNS {
    match {
        protocol udp;
        port 53;
        destination 203.0.113.1/32;
    }
```

```
    then discard;
}

[edit routing-options flow]
lab@mxD-Customer# commit and-quit
commit complete
Exiting configuration mode

lab@mxD-Customer>
```

Step 2.9

Use the **show route table inetflow.0 table extensive** to determine if a FlowSpec route has been created.

```
lab@mxD-Customer> show route table inetflow.0 extensive

inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
203.0.113.1,*,proto=17,port=53/term:1 (1 entry, 1 announced)
TSI:
KRT in dfwd;
Action(s): discard,count
Page 0 idx 0, (group SP type External) Type 1 val 0xd7e7adc (adv_entry)
   Advertised metrics:
     Flags: NoNexthop
     AS path: [65000] I
     Communities: traffic-rate:65000:0
Path 203.0.113.1,*,proto=17,port=53
Vector len 4.  Val: 0
        *Flow   Preference: 5
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xc0f1528
                Next-hop reference count: 27
                Next hop:
                State: <Active SendNhToPFE>
                Local AS: 65000
                Age: 1:18
                Validation State: unverified
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP_RT_Background
                AS path: I
                Communities: traffic-rate:65000:0
```

Question: Did the FlowSpec route display?

_____

_____

Answer: Yes the FlowSpec route did display.

Question: What action(s) has been set for the route?

_____

_____

Answer: The actions have been set to discard and count

Question: Why is the Next hop type set to Fictitious?

_____

_____

Answer: The Next hop type is set to Fictitious because there is no next hop for this type of NLRI

Step 2.10

Use the **show route advertising protocol bgp 172.22.131.10 extensive** to see if the route is being advertised to the service provider network.

```
lab@mxD-Customer> show route advertising-protocol bgp 172.22.131.10 extensive

inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
* 203.0.113.1,*,proto=17,port=53/term:1 (1 entry, 1 announced)
 BGP group SP type External
     Flags: NoNexthop
     AS path: [65000] I
     Communities: traffic-rate:65000:0
```

Question: Is the route being advertised?

_____

Answer: Yes, the route is being advertised.

Step 2.11

Return to the SSH session that you have open to the mxA-R1 router.

From the SSH session with the mxA-R1 router, use the **show route receiving protocol bgp 172.22.131.9 extensive table inetflow.0** to see if the Flow-route is being received from the Customer network.

```
lab@mxA-R1> show route receive-protocol bgp 172.22.131.9 table inetflow.0

inetflow.0: 1 destinations, 1 routes (0 active, 0 holddown, 1 hidden)

lab@mxA-R1>
```

Question: Is the mxA-R1 router receiving the Flow-route?

_____

Answer: Yes, The mxA-R1 router is receiving the Flow-route, but the route is hidden.

Step 2.12

Use the **show route receive-protocol bgp 172.22.131.9 table inetflow.0 extensive hidden** to see why the route is hidden.

```
lab@mxA-R1> show route receive-protocol bgp 172.22.131.9 table inetflow.0
   extensive hidden

inetflow.0: 1 destinations, 1 routes (0 active, 0 holddown, 1 hidden)
  203.0.113.1,*,proto=17,port=53/term:N/A (1 entry, 0 announced)
     Accepted
     Flags: NoNexthop
     AS path: 65000 I
     Communities: traffic-rate:65000:0
     Hidden reason: Flow-route fails validation

lab@mxA-R1>
```

Question: Why is the FlowSpec route hidden?

_____

Answer: It is hidden because the route has failed validation.

Question: What is needed for the route to pass validation?

_____

Answer: There are two ways for the route to pass validation. The Flow-route must be within the range of routes originated and advertised by the Customer to the Service Provider or a custom validation policy must exist on the service provider router to override the default validation process.

Step 2.13

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

From the SSH session with the mxD-Customer router, use the **show route advertising-protocol bgp 172.22.131.9 table inet.0** to see if any routes are being advertised.

```
lab@mxD-Customer> show route advertising-protocol bgp 172.22.131.9 table inet.0

lab@mxD-Customer>
```

Question: Is the Customer site advertising the route that represents its internal network?.

_____

Answer: No, it is not advertising its internal route. This causes the FlowSpec to fail the validation test.

Step 2.14

Configure the Customer router to only advertise their internal network. This will require a policy. Since the goal is to only advertise the one route that belongs to the Customer and not any others, a reject all term will be needed. This reject all term will cause the FlowSpec route to also be discarded. In this policy, make sure the FlowSpec route is also included.

```
lab@mxD-Customer> configure
Entering configuration mode

[edit]
lab@mxD-Customer# edit policy-options policy-statement TO-SP

[edit policy-options policy-statement TO-SP]
lab@mxD-Customer# set term FLOW-to-SP from rib inetflow.0

[edit policy-options policy-statement TO-SP]
lab@mxD-Customer# set term FLOW-to-SP then accept

[edit policy-options policy-statement TO-SP]
lab@mxD-Customer# set term CUSTOMER-ROUTES from route-filter 203.0.113.0/24 exact

[edit policy-options policy-statement TO-SP]
lab@mxD-Customer# set term CUSTOMER-ROUTES then accept

[edit policy-options policy-statement TO-SP]
lab@mxD-Customer# set term REJECT-ALL then reject

[edit policy-options policy-statement TO-SP]
lab@mxD-Customer# top edit protocols bgp group SP

[edit protocols bgp group SP]
lab@mxD-Customer# set export TO-SP

[edit protocols bgp group SP]
lab@mxD-Customer# commit and-quit
commit complete
Exiting configuration mode

lab@mxD-Customer>
```

Step 2.15

Ensure that the Customer site is advertising both routes to the service provider using the **show route advertising route bgp 172.22.131.10.**

```
lab@mxD-Customer> show route advertising-protocol bgp 172.22.131.10

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 203.0.113.0/24          Self                                    I

inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  203.0.113.1,*,proto=17,port=53/term:1
*                                                                 I
```

Question: Are both routes being advertised?

_____

_____

Answer: Yes, both routes are being advertised.

Step 2.16

Return to the SSH session that you have open to the mxA-R1 router.

From the SSH session with the mxA-R1 router, verify that the routes are being received from the Customer and the routes are active by issuing the **`show route receive-protocol bgp 172.22.131.9`**.

```
lab@mxA-R1> show route receive-protocol bgp 172.22.131.9

inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 203.0.113.0/24          172.22.131.9                            65000 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  203.0.113.1,*,proto=17,port=53/term:1
*                                                                  65000 I

lab@mxA-R1>
```

Question: Are both routes being received?

_____

Answer: Yes, both routes are being received.

Question: Is either route hidden?

_____

Answer: None of the routes are hidden.

Step 2.17

Verify that R1 has established a FlowSpec peering with all neighbors by running the **`show bgp summary`** command.

```
lab@mxA-R1> 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
                        1          1          0          0          0          0
inetflow.0
                        1          1          0          0          0          0
Peer                      AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.9          65000        182        178       0       8     1:19:49 Establ
  inet.0: 1/1/1/0
  inetflow.0: 1/1/1/0
192.168.71.2          64512        462        463       0       0     3:27:18 Establ
  inet.0: 0/0/0/0

lab@mxA-R1>
```

Question: Is there a FlowSpec session established with R2 which peers using the address of 192.168.71.2

_____

_____

Answer: No, there is no FlowSpec peering.

Step 2.18

Verify the IBGP configuration on R1 using the **show configuration protocols bgp group ibgp** command. If family inet flow is not part of the group configuration, add the family inet flow option.

```
lab@mxA-R1> show configuration protocols bgp group ibgp
type internal;
local-address 192.168.71.1;
family inet {
    unicast;
}
export nhs;
neighbor 192.168.71.2;

lab@mxA-R1> configure
Entering configuration mode

[edit]
lab@mxA-R1# set protocols bgp group ibgp family inet flow
```

Step 2.19

Configure the **routing-options flow term-order** to use the **standard** option. All edge devices should have this set if they are participating in FlowSpec. Run the **commit and-quit** command when finished.

```
[edit]
lab@mxA-R1# set routing-options flow term-order standard

[edit]
```

```
lab@mxA-R1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA-R1>
```

Step 2.20

Verify they the FlowSpec peering is established using the **show bgp summary** command and that the FlowSpec routes are being sent to the RR using the **show route advertising-protocol bgp 192.168.71.2.**

```
lab@mxA-R1> 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
                       1          1          0          0          0          0
inetflow.0
                       1          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.9          65000        196        193       0       8     1:26:10 Establ
  inet.0: 1/1/1/0
  inetflow.0: 1/1/1/0
192.168.71.2          64512          5          7       0       0        1:09 Establ
  inet.0: 0/0/0/0
  inetflow.0: 0/0/0/0

lab@mxA-R1>


lab@mxA-R1> show route advertising-protocol bgp 192.168.71.2

inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 203.0.113.0/24          Self                         100        65000 I

inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  203.0.113.1,*,proto=17,port=53/term:1
*                                                      100        65000 I
```

Question: Is the FlowSpec session established with R2?

_____

_____

Answer: Yes, the FlowSpec session is established with R2

Question: Is R1 advertising the Customer routes to R2?

_____

_____

Answer: Yes, R1 is advertising the Customer routes to R2.

Step 2.21

Return to the SSH session that you have open to the mxC-R3 router.

From the SSH session with the mxC-R3 router, verify that the FlowSpec route is received by the mxC-R3 router and that it has validated the Flow-route using the **`show route flow validation 203.0.113.0/24`** command. Then ensure that the firewall filter was created using the **`show firewall`** command.

```
lab@mxC-R3> show route flow validation 203.0.113.0/24

inet.0:
203.0.113.0/24
                Active unicast route
                    Dependent flow destinations: 1
                    Origin: 192.168.71.1, Neighbor AS: 65000
203.0.113.1/32
                Flow destination (1 entries, 1 match origin, next-as)
                    Unicast best match: 203.0.113.0/24
                    Flags: Consistent


lab@mxC-R3> show firewall

Filter: __default_bpdu_filter__

Filter: __flowspec_default_inet__
Counters:
Name                                                Bytes              Packets
203.0.113.1,*,proto=17,port=53                          0                    0

lab@mxC-R3> 0                      0
```

Question: Were both routes received at R3?

_____

_____

Answer: Yes, both routes were received at R3.

Question: Was the FlowSpec firewall filter created?

_____

_____

Question: Yes the FlowSpec firewall filter was created.

Step 2.22

Return to the SSH session that you have open to the mxE-Attacker router.

From the SSH session with the mxE-Attacker router, use the JUNOS traceroute command and set the port to 53. This will generate one packet of UDP port 53 traffic before switching over to random ports.

Execute the **`traceroute 203.0.113.1 port 53`** command on the mxE-Attacker router. The first traceroute packet will be dropped and we can verify by looking at the traceroute output and then running the show firewall command on mxC.

```
lab@mxE-Attacker> traceroute 203.0.113.1 port 53
traceroute to 203.0.113.1 (203.0.113.1), 30 hops max, 52 byte packets
 1  * 172.22.131.13 (172.22.131.13)  2.085 ms  1.702 ms
 2  172.22.131.5 (172.22.131.5)  2.396 ms  2.411 ms  2.083 ms
 3  172.22.131.1 (172.22.131.1)  3.309 ms  2.528 ms  2.684 ms
 4  203.0.113.1 (203.0.113.1)  4.238 ms  5.102 ms  3.857 ms
```

Question: Was there a delay and was the first traceroute output an *?

_____

_____

Answer: Yes there was a delay. This produced an * in the first line of the returned output. This was the one port 53 packet sent by traceroute before randomizing the port numbers.

Step 2.23

Return to the SSH session that you have open to the mxC-R3 router.

From the SSH session with the mxC-R3 router, run the **`show firewall`** command and verify that one packet was stopped.

```
lab@mxC-R3> show firewall

Filter: __default_bpdu_filter__

Filter: __flowspec_default_inet__
Counters:
Name                                                Bytes              Packets
203.0.113.1,*,proto=17,port=53                         52                    1

lab@mxC-R3>
```

Question: Was there one packet dropped?

_____

_____

Answer: Yes, one packet was dropped.

Part 3: Setting up the NOC to perform DDoS Attack Remediation

In this lab part, the NOC has been contacted by the Customer which is complaining that they are receiving an excessive amount of large ping packets directed at the Server at 203.0.113.1. They would like the service provider to block ICMP-request that are 100 bytes or larger directed to this server.

In this example the NOC PE routers must be configured to validate routes received from the route reflector, since the Flow-route is not being sent from the originator of the network prefix route. Since these routes are being introduced from inside the AS to PE routers inside the AS the validation checks are meant to ensure a configuration mistake does not end up blocking a large number or prefixes. This is accomplished by creating a policy that only accepts /32 prefixes into the inetflow.0 table.

Step 3.1

Return to the SSH session that you have open to the mxA-R1 router.

From the SSH session with the mxA-R1 router, enter configuration mode and create a policy to accept any /32 Flow-route received from the RR. Then apply this policy under the **family inet flow** heirarchy after the **no-validation** option.

```
lab@mxA-R1> configure
Entering configuration mode

[edit]
lab@mxA-R1# edit policy-options policy-statement Allow-FS-IN

[edit policy-options policy-statement Allow-FS-IN]
lab@mxA-R1# set term FS-RR-IN from rib inetflow.0

[edit policy-options policy-statement Allow-FS-IN]
lab@mxA-R1# set term FS-RR-IN from route-filter 0/0 prefix-length-range /32-/32

[edit policy-options policy-statement Allow-FS-IN]
lab@mxA-R1# set term FS-RR-IN then accept

[edit policy-options policy-statement Allow-FS-IN]
lab@mxA-R1# set term REJECT-ALL then reject

[edit policy-options policy-statement Allow-FS-IN]
lab@mxA-R1# show
term FS-RR-IN {
    from {
        rib inetflow.0;
        route-filter 0.0.0.0/0 prefix-length-range /32-/32;
    }
    then accept;
}
term REJECT-ALL {
    then reject;
}
```

```
[edit policy-options policy-statement Allow-FS=IN]
lab@mxA-R1# top edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxA-R1# set family inet flow no-validate Allow-FS-IN

[edit protocols bgp group ibgp]
lab@mxA-R1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA-R1>
```

Step 3.2

Return to the SSH session that you have open to the mxC-R3 router.

From the SSH session with the mxC-R3 router, repeat the previous configuration that you just completed on mxA-R1.

```
lab@mxC-R3> configure
Entering configuration mode

[edit]
lab@mxC-R3# edit policy-options policy-statement Allow-FS-IN

[edit policy-options policy-statement Allow-FS-IN]
lab@mxC-R3# set term FS-IN-RR from rib inetflow.0

[edit policy-options policy-statement Allow-FS-IN]
lab@mxC-R3# set term FS-IN-RR from route-filter 0/0 prefix-length-range /32-/32

[edit policy-options policy-statement Allow-FS-IN]
lab@mxC-R3# set term FS-IN-RR then accept

[edit policy-options policy-statement Allow-FS-IN]
lab@mxC-R31# set term REJECT-ALL then reject

[edit policy-options policy-statement Allow-FS-IN]
lab@mxC-R3# show
term FS-IN-RR {
    from {
        rib inetflow.0;
        route-filter 0.0.0.0/0 prefix-length-range /32-/32;
    }
    then accept;
}
term REJECT-ALL {
    then reject;
}

[edit policy-options policy-statement Allow-FS-IN]
lab@mxC-R3# top edit protocols bgp group ibgp

[edit protocols bgp group ibgp]
lab@mxC-R3# set family inet flow no-validate Allow-FS-IN
```

```
[edit protocols bgp group ibgp]
lab@mxC-R3# commit and-quit
commit complete
Exiting configuration mode

lab@mxC-R3>
```

Step 3.3

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

From the SSH session with the mxD router, configure the Flow specification to drop ICMP-request packets that exceed 100 bytes. This will required that you set a packet-length. Since the packet size configured excludes layer 2 headers but includes the IPv4 packet header plus the ICMP header, the size will need to be 28 bytes larger than the packet size. Enter configuration mode and use the **routing-options flow** command to set the **destination** to be **203.0.113.1,** the **icmp-type** to be **echo-request** and the packet-length to be ***128-65535*** bytes. Verify your work using the **show** command before committing your work.

```
lab@mxB-R2> configure
Entering configuration mode

[edit]
lab@mxB-R2# edit routing-options flow route ICMP

[edit routing-options flow route ICMP]
lab@mxB-R2# set match packet-length 128-65535

[edit routing-options flow route ICMP]
lab@mxB-R2# set match icmp-type echo-request

[edit routing-options flow route ICMP]
lab@mxB-R2# set match destination 203.0.113.1

[edit routing-options flow route ICMP]
lab@mxB-R2# set then discard

[edit routing-options flow route ICMP]
lab@mxB-R2# show
match {
    packet-length 128-65535;
    destination 203.0.113.1/32;
    icmp-type echo-request;
}
then discard;

[edit routing-options flow route ICMP]
lab@mxB-R2# commit and-quit
commit complete
Exiting configuration mode

lab@mxB-R2>
```

Step 3.4

In this lab step you will verify that the Flow-route has been created and is being forwarded to the IBGP neighbors. View the inetflow.0 table using the **`show route table inetflow.0`** command to verify that the new Flow-route is added. To verify that the routes have been advertised to the neighbors using the command **`show route advertising-protocol bgp 192.168.71.3`**.

```
lab@mxB-R2> show route table inetflow.0

inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

203.0.113.1,*,proto=17,port=53/term:1
                   *[BGP/170] 00:26:46, localpref 100, from 192.168.71.1
                      AS path: 65000 I, validation-state: unverified
                      Fictitious
203.0.113.1,*,icmp-type=8,len>=128&<=65535/term:2
                   *[Flow/5] 00:01:27
                      Fictitious
```

Question: Has the new flow route for ICMP traffic destined for the 203.0.113.1 server been added?

_____

_____

Answer: Yes, the new flow route for ICMP traffic destined for the 203.0.113.1 server has been added.

```
lab@mxB-R2> show route advertising-protocol bgp 192.168.71.3

inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 203.0.113.0/24          192.168.71.1                 100        65000 I

inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  203.0.113.1,*,proto=17,port=53/term:1
*                                                      100        65000 I
  203.0.113.1,*,icmp-type=8,len>=128&<=65535/term:2
*                                                      100        I
```

Question: Is the ICMP Flow-route being advertised to its RR client peer at 192.168.71.3?

_____

_____

Question: Yes the ICMP Flow-route is being advertised to it RR client peer.

Step 3.5

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

From the SSH session with the mxD router, verify that the Flow-route has been received and the Firewall filter has been installed. Use the **`show route receive-protocol bgp 192.168.71.2`** command. The use the **`show firewall`** command to ensure the `__flowspec_default_inet__` firewall filter is present.

```
lab@mxC-R3> show route receive-protocol bgp 192.168.71.2

inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 203.0.113.0/24          192.168.71.1                 100        65000 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  203.0.113.1,*,proto=17,port=53/term:1
*                                                      100        65000 I
  203.0.113.1,*,icmp-type=8,len>=128&<=65535/term:2
*                                                      100        I

lab@mxC-R3>
```

Question: Has the new flow route for ICMP traffic destined for the 203.0.113.1 server been received?

_____

_____

Answer: Yest the new flow route for ICMP traffic destined for the 203.0.113.1 server has been received.

```
lab@mxC-R3> show firewall

Filter: __default_bpdu_filter__

Filter: __flowspec_default_inet__
Counters:
Name                                                Bytes              Packets
203.0.113.1,*,icmp-type=8,len>=128&<=65535                0                    0
203.0.113.1,*,proto=17,port=53                           52                    1

lab@mxC-R3>
```

Question: Has the flowspec filter for ICMP been created?

_____

_____

Answer: Yes, the flowspec filter for ICMP has been created.

Step 3.6

Return to the SSH session that you have open to the mxE-Attacker router.

From the SSH session with the mxE-Attacker router, test the effectiveness of the icmp FlowSpec Filter. Try three different pings. First try a normal ping using the **`ping 203.0.113.1 count 1`** command. This should fail. Next try a ping with a packet size of 99 bytes using the **ping `203.0.113.1 count 1 size 99`.** This should also fail. Finally perform a ping with a packet size of 100 using the **`ping 203.0.113.1 count 1 size 100`** command.

```
lab@mxE-Attacker> ping 203.0.113.1 count 1
PING 203.0.113.1 (203.0.113.1): 56 data bytes
64 bytes from 203.0.113.1: icmp_seq=0 ttl=61 time=4.540 ms

--- 203.0.113.1 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.540/4.540/4.540/0.000 ms

lab@mxE-Attacker> ping 203.0.113.1 count 1 size 99
PING 203.0.113.1 (203.0.113.1): 99 data bytes
107 bytes from 203.0.113.1: icmp_seq=0 ttl=61 time=3.459 ms

--- 203.0.113.1 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 3.459/3.459/3.459/0.000 ms

lab@mxE-Attacker> ping 203.0.113.1 count 1 size 100
PING 203.0.113.1 (203.0.113.1): 199 data bytes

--- 203.0.113.1 ping statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

lab@mxE-Attacker>
```

Question: Did the first two pings succeed and the last one fail?

_____

_____

Answer: Yes the first two pings did succeed and the last ping attempt failed. The failure was caused by the packet size matching the conditions set by the flowspec rule.

Step 3.7

Return to the SSH session that you have open to the mxC-R3 router.

From the SSH session with the mxC-R3 router, Verify that the filter trapped the Attack traffic based on the criteria using the **show firewall** command.

```
lab@mxC-R3> show firewall

Filter: __default_bpdu_filter__

Filter: __flowspec_default_inet__
Counters:
Name                                                Bytes              Packets
203.0.113.1,*,icmp-type=8,len>=128&<=65535            128                    1
203.0.113.1,*,proto=17,port=53                         52                    1

lab@mxC-R3>
```

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

```
lab@mxC-R3> exit


mxC-R3 (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

Network Management Diagram

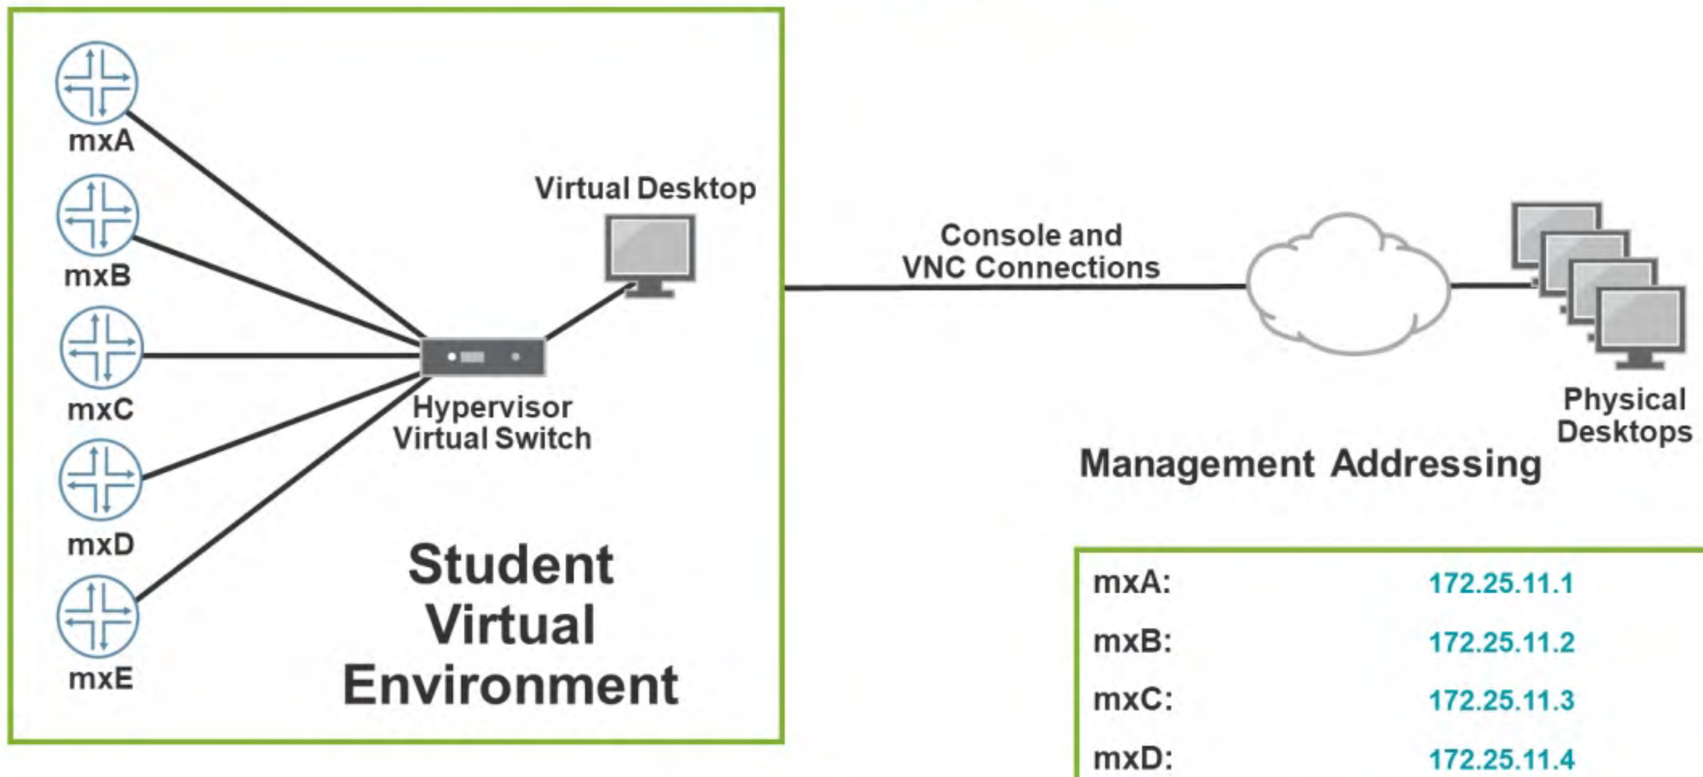

Management Addressing

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: BGP FlowSpec

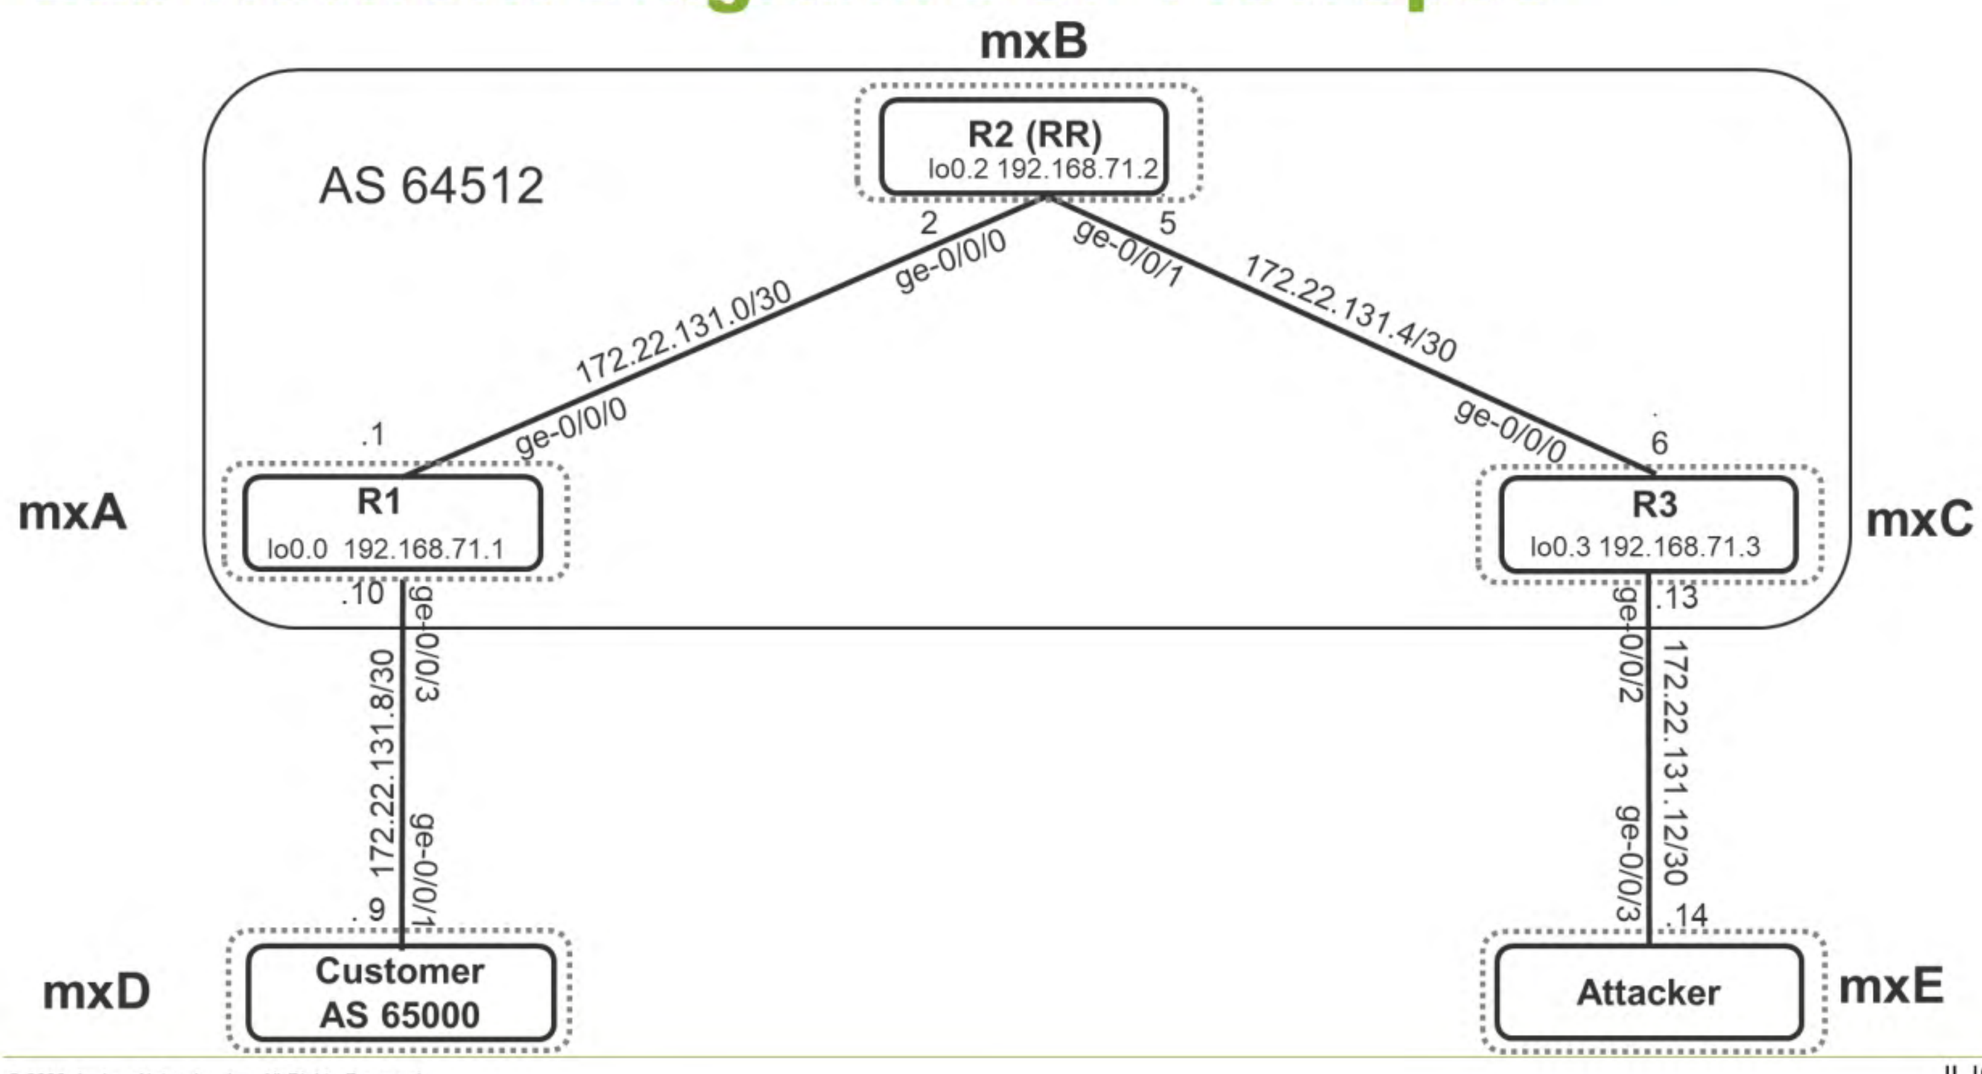

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab

Troubleshooting BGP

Overview

In this lab, you use the lab diagram titled "Lab Network Diagram: BGP Troubleshooting" to troubleshoot BGP neighbor relationships and routing issues. The goal is to get all neighbor relationships up and functioning and ensure that routing is functioning correctly. Ensure that you do not have any hidden routes within your BGP network.

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

- Troubleshoot various BGP 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 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 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/lab14-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 27 17:20:28 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/lab14-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 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/lab14-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: Wed Nov 27 17:20:28 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/lab14-start.config
load complete

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

lab@mxE>
```

Step 1.3

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/lab14-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 27 17:21:22 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/lab14-start.config
load complete

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

lab@mxC>
```

Step 1.4

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/lab14-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 27 17:21:45 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/lab14-start.config
load complete

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

lab@mxB>
```

Step 1.5

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/lab14-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: Wed Nov 27 15:27:40 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/lab14-start.config
load complete

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

lab@mxA>
```

Part 2: Troubleshooting BGP

In this lab part, you will be presented with a number of EBGP and IBGP adjacency issues. You will use various CLI tools to troubleshoot and correct these BGP peering issues.

Step 2.1

Change the perspective of the CLI to that of the R1 logical system by issuing the **`set cli logical-system R1`** command.

```
lab@mxA> set cli logical-system R1
Logical system: R1

lab@mxA:R1>
```

Step 2.2

Use the **`show bgp summary`** command to view the BGP neighborships on the R1 router.

```
lab@mxA:R1> show bgp summary
Threading mode: BGP I/O
Groups: 2 Peers: 3 Down peers: 3
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       0          0          0          0          0          0
inet6.0
                       0          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.37         65230          0          0       0       0          23 Active
192.168.71.2          65501          0          0       0       0          23 Idle
192.168.71.4          65501          0          0       0       0          23 Idle
```

Question: How many neighbors does the R1 router have configured?

_____

_____

Answer: You should see three BGP neighbors that have been configured on the R1 router.

Question: What is the status of the BGP sessions?

_____

_____

Answer: You should see two BGP sessions in `Idle` state and you should see one session in an `Active` state.

Step 2.3

Use the **`show bgp neighbor 172.22.131.37`** command to examine the EBGP session that is currently in an Active state. You should look for a line beginning with 'Error'.

```
lab@mxA:R1> show bgp neighbor 172.22.131.37
Peer: 172.22.131.37 AS 65230   Local: 172.22.131.38 AS 65501
  Group: ebgp-AS65530          Routing-Instance: master
  Forwarding routing-instance: master
  Type: External    State: Active         Flags: <>
  Last State: Idle          Last Event: Start
  Last Error: Open Message Error
```

```
Options: <Preference AddressFamily PeerAS Refresh>
Options: <GracefulShutdownRcv>
Address families configured: inet-unicast inet6-unicast
Holdtime: 90 Preference: 170
Graceful Shutdown Receiver local-preference: 0
Number of flaps: 0
Error: 'Open Message Error' Sent: 7 Recv: 0
```

Question: Why is the neighborship showing in an `Active` state?

_____

_____

Answer: The peer is down due to an `Open Message Error`.

Step 2.4

Enter into configuration mode and navigate to the `[edit protocols bgp]` hierarchy and enable traceoptions flagging the **`normal`** and **`open`** conditions. The traceoptions file should be called ***`bgp-trace`***. When you are satisfied with the changes, **`commit`** the configuration changes before moving on.

```
lab@mxA:R1> configure
Entering configuration mode

[edit]
lab@mxA:R1# edit protocols bgp

[edit protocols bgp]
lab@mxA:R1# set traceoptions file bgp-trace

[edit protocols bgp]
lab@mxA:R1# set traceoptions flag normal

[edit protocols bgp]
lab@mxA:R1# set traceoptions flag open

[edit protocols bgp]
lab@mxA:R1# commit
commit complete

[edit protocols bgp]
lab@mxA:R1#
```

Step 2.5

Review the contents of the new traceoptions file and identify the problem relating the failure to establish a neighborship with your EBGP peer. You should be looking for messages that relate to the EBGP session that is currently in `Active` state.

```
[edit protocols bgp]
lab@mxA:R1# run show log bgp-trace
Dec  1 00:03:14 trace_on: Tracing to "/var/log/R1/bgp-trace" started...

Dec  1 00:04:00.500082 BGP RECV 172.22.131.37+179 -> 172.22.131.38+52484
Dec  1 00:04:00.500104 BGP RECV message type 1 (Open) length 71
```

```
Dec  1 00:04:00.500459 bgp_process_open:4185: NOTIFICATION sent to 172.22.131.37
   (External AS 65230): code 2 (Open Message Error) subcode 2 (bad peer AS number),
   Reason: peer 172.22.131.37 (External AS 65230) claims 65530, 65230 configured
Dec  1 00:04:00.500473 Notify sent to 172.22.131.37 (External AS 65230), code 2,
   subcode 2
Dec  1 00:04:00.500485 BGP_65230.172.22.131.37: send proc: sending 21 bytes
Dec  1 00:04:00.500502
Dec  1 00:04:00.500502 BGP SEND 172.22.131.38+52484 -> 172.22.131.37+179
Dec  1 00:04:00.500521 BGP SEND message type 3 (Notification) length 21
Dec  1 00:04:00.500533 BGP SEND Notification code 2 (Open Message Error) subcode 2
   (bad peer AS number)
Dec  1 00:04:00.500587 BGP_65230.172.22.131.37: send proc: writev 21/21 bytes, rc
   1
Dec  1 00:04:00.500603 bgp_peer_close_and_restart: peer 172.22.131.37 (External AS
   65230), state is 4 (OpenSent) event RecvOpen, flags=0x0
Dec  1 00:04:00.500609 bgp_peer_close_and_restart: closing peer 172.22.131.37
   (External AS 65230), state is 4 (OpenSent) event RecvOpen
Dec  1 00:04:00.500615 bgp_send_deactivate:2873: 172.22.131.37 (External AS 65230)
   ,flags=0x0: removed from active list
Dec  1 00:04:00.500626 BGP peer 172.22.131.37 (External AS 65230) CLOSE: close
   socket 98
Dec  1 00:04:00.500682 bgp_peer_close_and_restart: 8177: BGP peer 172.22.131.37
   (External AS 65230): state OpenSent task close fd -1
Dec  1 00:04:00.501004 bgp_delete_pending_set_impatient: peer 172.22.131.37
   (External AS 65230) job-exist:no, flags:, need-job:no
Dec  1 00:04:00.501019 bgp_peer_post_close:7813 BGP Peer 172.22.131.37 (External
   AS 65230) CLOSE: post close cleanup - NEED restart, (last_flap NoEvent, state
   Idle, flags )
```

Note

You might have to parse through quite a few messages before finding the messages that relates to the EBGP session. You will also see many messages that relate to other BGP sessions. You can ignore those messages at this time.

Question: Why is the peering session not establishing?

_____

_____

Answer: The reason the session is not establishing is because the peer AS is not configured correctly for this session. the message indicates that 65230 is configured but the peer is looking for 65530.

Step 2.6

Correct the peer AS number that is not configured correctly for the EBGP group by changing the configured peer AS to 65530. After making the correct changes **commit** your configuration.

```
[edit protocols bgp]
lab@mxA:R1# show
traceoptions {
```

```
    file bgp-trace;
    flag normal;
    flag open;
}
group ibgp-RR {
    type internal;
    local-address 192.168.0.4;
    family inet {
        unicast;
    }
    family inet6 {
        unicast;
    }
    export next-hop-self;
    neighbor 192.168.71.2;
    neighbor 192.168.71.4;
}
group ebgp-AS65530 {
    family inet {
        unicast;
    }
    family inet6 {
        unicast;
    }
    peer-as 65230;
    neighbor 172.22.131.37;
}

[edit protocols bgp]
lab@mxA:R1# set group ebgp-AS65530 peer-as 65530

[edit protocols bgp]
lab@mxA:R1# commit
commit complete
```

Step 2.7

Use the **`run show bgp summary`** command to verify the state of the EBGP peering session after making your configuration changes.

```
[edit protocols bgp]
lab@mxA:R1# run show bgp summary
Groups: 2 Peers: 3 Down peers: 2
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      11         11          0          0          0          0
inet6.0
                       1          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.37         65530         11          3       0       0          22 Establ
  inet.0: 11/11/11/0
  inet6.0: 1/1/1/0
192.168.71.2          65501          0          0       0       0     5:41:36 Idle
192.168.71.4          65501          0          0       0       0     5:41:36 Idle
```

Question: What is the current status of the EBGP peering session?

_____

Answer: From the output, you should be able to see that the EBGP session is up and established.

Step 2.8

Begin troubleshooting why the other two sessions are not working. You should start by verifying that you have routes to their loopback addresses using the **`run show route 192.168.71/24`** command.

```
[edit protocols bgp]
lab@mxA:R1# run show route 192.168.71/24

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

192.168.71.1/32    *[Direct/0] 05:42:28
                    > via lo0.0
192.168.71.2/32    *[IS-IS/18] 05:42:13, metric 10
                    > to 172.22.131.2 via ge-0/0/0.0
192.168.71.3/32    *[IS-IS/18] 05:42:04, metric 20
                    > to 172.22.131.2 via ge-0/0/0.0
192.168.71.4/32    *[IS-IS/18] 05:42:03, metric 20
                    > to 172.22.131.2 via ge-0/0/0.0
```

Question: Do you have routes to the peer loopback addresses?

_____

Answer: Yes, you should have routes to the loopback addresses of the R2 and R4 routers.

Question: What is the next step?

_____

_____

Answer: We have verified that we have a route to the remote peers. The next step is to verify what address we are using locally for the peering session.

Step 2.9

Verify the address that you are using locally when trying to establish the peering sessions to the R2 and R4 routers.

```
[edit protocols bgp]
lab@mxA:R1# show group ibgp-RR
type internal;
local-address 192.168.0.4;
family inet {
    unicast;
```

```
}
family inet6 {
    unicast;
}
export next-hop-self;
neighbor 192.168.71.2;
neighbor 192.168.71.4;

[edit protocols bgp]
lab@mxA:R1#
```

Question: Is the local address being used for your IBGP peerings configured correctly?

_____

Answer: No, the local address is not correct and should be configured as your local loopback address.

Step 2.10

Change the local address being used to establish the sessions to your IBGP neighbors. This address should be configured as your local loopback address. When you are satisfied with the configuration changes, **commit** and exit to operational mode.

```
[edit protocols bgp]
lab@mxA:R1# set group ibgp-RR local-address 192.168.71.1

[edit protocols bgp]
lab@mxA:R1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R1>
```

Step 2.11

Use the **show bgp summary** command to verify the recent changes have resolved the peering issues to the R2 and R4 routers.

```
lab@mxA:R1> show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      34         22          0          0          0          0
inet6.0
                       1          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.37         65530         22         20       0       0        5:28 Establ
  inet.0: 11/11/11/0
  inet6.0: 1/1/1/0
192.168.71.2          65501         11         11       0       0          20 Establ
  inet.0: 11/11/11/0
  inet6.0: 0/0/0/0
192.168.71.4          65501          9          9       0       0          16 Establ
  inet.0: 0/12/12/0
```

Question: What is the current status of the neighborships to the R2 and R4 routers?

_____

_____

Answer: Both neighborships should be up and established at this point.

Step 2.12

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 2.13

Using the **`show bgp summary`** command, verify that the R2 router's BGP sessions are up and operational. R2 is a route reflector (along with R4) and should have IBGP sessions to all internal devices.

```
lab@mxB:R2> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      23         22          0          0          0          0
inet6.0
                       2          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
192.168.71.1          65501         19         17       0       0        3:16 Establ
  inet.0: 11/11/11/0
  inet6.0: 1/1/1/0
192.168.71.3          65501        820        824       0       0     6:06:50 Establ
  inet.0: 11/12/12/0
  inet6.0: 0/1/1/0
192.168.71.4          65501        828        828       0       0     6:06:54 Establ
  inet.0: 0/0/0/0
192.168.71.5          65501          0          0       0       0     6:07:38 Active
```

Question: Are all the IBGP sessions up and established?

_____

_____

Answer: All sessions are up and established except for the session to the R5 router. This session is in an `Active` state.

Step 2.14

Verify that you have a route to the loopback address for R5 using the **`show route 192.168.71.5`** command.

```
lab@mxB:R2> show route 192.168.71.5

lab@mxB:R2>
```

Question: Do you have a route to the R5 router's loopback address?

_____

_____

Answer: No, you do not have a route to the R5 router's loopback.

Question: What is the next step?

_____

_____

Answer: The next step is to verify you have the correct IS-IS adjacencies.

Step 2.15

Using the **`show isis adjacency`** command, verify that you have the appropriate IS-IS adjacencies to neighboring devices.

```
lab@mxB:R2> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/0.0            mxA-R1         2  Up                   19  0:50:56:a9:25:1a
ge-0/0/2.0            mxB-R3         2  Up                    8  0:50:56:a9:21:f2
lt-0/0/10.0           mxB-R4         2  Up                    7  0:5:86:4e:5a:1
```

Question: Do you see all your IS-IS neighbors?

_____

_____

Answer: Yes, You should have neighborships with R1, R3 and R4 based on the network topology. These sessions should be up and functioning at this point.

Step 2.16

Use the **`show isis database`** command to verify that you see the R5 router's link-state PDUs in the link-state database.

```
lab@mxB:R2> 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-R1.00-00                    0x22   0x31d9      810 L1 L2
mxB-R2.00-00                    0x22   0xcc5b      610 L1 L2
mxB-R2.02-00                    0x1f   0x64cc      798 L1 L2
mxB-R3.00-00                    0x20   0x9983     1023 L1 L2
mxB-R3.02-00                    0x1e   0x61c7     1087 L1 L2
mxB-R3.03-00                    0x1d   0xcf5d     1087 L1 L2
mxB-R4.00-00                    0x23   0x8f41      881 L1 L2
mxB-R4.02-00                    0x20   0xcb56     1020 L1 L2
mxB-R4.03-00                    0x1d   0xf732      668 L1 L2
mxC-R5.00-00                    0x20   0x2968      741 L1 L2
  10 LSPs
```

Step 2.17

Use the **`show isis database mxC-R5.00-00 extensive | match prefix`** command to determine if R5 is sending its own loopback.

```
lab@mxB:R2> show isis database mxC-R5.00-00 extensive | match prefix
   IP prefix: 172.22.131.16/30                Metric:       10 Internal Up
IP prefix: 172.22.131.16/30, Internal, Metric: default 10, Up
IP extended prefix: 172.22.131.16/30 metric 10 up
```

Question: Do you see the R5 router's loopback address in the IS-IS link-state database?

_____

_____

Answer: No, you should not see the R5 router's loopback at this time. This indicates there is a problem on R5. We will address this problem when we move to R5.

Note

The problem with the peering session from R2 to R5 can not be resolved on R2. To keep some order to our troubleshooting steps we will move on to verifying and troubleshooting R3's BGP sessions and we will address the R2 to R5 neighborship problem when we get to verifying and troubleshooting R5.

Step 2.18

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 2.19

Use the **`show bgp summary`** command to verify the state of the R3 router's BGP sessions.

```
lab@mxB:R3> show bgp summary
Groups: 3 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
```

```
inet6.0
                      2          2          0          0          0          0
inet.0
                     33         22          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.46         65532        855        847       0       0     6:21:20 12/
   12/12/0           0/0/0/0
192.168.71.2          65501        840        835       0       0     6:14:26 Establ
  inet6.0: 1/1/1/0
  inet.0: 10/11/11/0
192.168.71.4          65501        852        847       0       0     6:21:12 Establ
  inet.0: 0/10/10/0
2001:db8:101:3::46       65532        849        841       0       0     6:21:05
   Establ
  inet6.0: 1/1/1/0
```

Question: Are all the BGP neighborships up and established?

_____

_____

Answer: Yes, all neighborships should be up and established on R3. You should have four peering session from R3.

Step 2.20

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 2.21

Use the **`show bgp summary`** command to verify the state of the BGP sessions. Keep in mind that R4 is also a route reflector and should have BGP neighborships with all internal devices.

```
lab@mxB:R4> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                     23         22          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
192.168.71.1          65501         37         36       0       0       12:41 10/
   11/11/0           0/0/0/0
192.168.71.2          65501        851        849       0       0     6:16:23 0/0/
   0/0               0/0/0/0
192.168.71.3          65501        853        856       0       0     6:23:05 12/
   12/12/0           0/0/0/0
192.168.71.5          65501          0          0       0       0     6:34:26 Active
```

Question: Are all the BGP sessions up and established from R4?

_____

_____

Answer: No, The peering session to R5 is in an `Active` state.

Step 2.22

Verify that you have a route to the loopback address for the R5 router using the **`show route 192.168.71.5`** command.

```
lab@mxB:R4> show route 192.168.71.5

lab@mxB:R4>
```

Question: Do you have a route to the R5 router's loopback address?

_____

_____

Answer: No, you do not have a route to the R5 router's loopback.

Question: What is the next step?

_____

_____

Answer: The next step is to verify that you have the correct IS-IS adjacencies.

Step 2.23

Using the **`show isis adjacency`** command, verify that you have the appropriate IS-IS adjacencies to neighboring devices.

```
lab@mxB:R4> show isis adjacency
Interface             System         L State        Hold (secs) SNPA
ge-0/0/1.0            mxC-R5         2  Up                   26  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                   22  0:5:86:4e:5a:0
```

Question: Do you see all your IS-IS neighbors?

_____

_____

Answer: Yes, you should have neighborships with the R2, R3, and R5 routers based on the network topology. These sessions should be up and functioning at this point.

Step 2.24

Use the **`show isis database`** command to verify that you see R5's link-state PDUs in the link-state database.

```
lab@mxB:R4> 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-R1.00-00                    0x23   0x2fda     1076 L1 L2
mxB-R2.00-00                    0x23   0xca5c      850 L1 L2
mxB-R2.02-00                    0x20   0x62cd     1081 L1 L2
mxB-R3.00-00                    0x21   0x9784     1158 L1 L2
mxB-R3.02-00                    0x1e   0x61c7      515 L1 L2
mxB-R3.03-00                    0x1d   0xcf5d      515 L1 L2
mxB-R4.00-00                    0x24   0x8d42     1155 L1 L2
mxB-R4.02-00                    0x21   0xc957     1155 L1 L2
mxB-R4.03-00                    0x1e   0xf533      803 L1 L2
mxC-R5.00-00                    0x21   0x2769      846 L1 L2
  10 LSPs
```

Step 2.25

Use the **`show isis database`** ***`mxC-R5.00-00`*** **`extensive | match prefix`** command to determine if the R5 router is advertising its own loopback address.

```
lab@mxB:R4> show isis database mxC-R5.00-00 extensive | match prefix
   IP prefix: 172.22.131.16/30                Metric:       10 Internal Up
    IP prefix: 172.22.131.16/30, Internal, Metric: default 10, Up
    IP extended prefix: 172.22.131.16/30 metric 10 up
```

Question: Do you see the R5 router's loopback address in the IS-IS link-state database?

_____

_____

Answer: No, you should not see the R5 router's loopback at this time. This indicates there is a problem on R5. We will address this problem when we move to R5.

Note

The problem with the peering session from R4 to R5 can not be resolved on R4. This problem seems to be related to the same issue that R2 had with its peering to R5. Move to R5 now and verify and troubleshoot its BGP neighborships.

Step 2.26

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 R5 logical system by issuing the **`set cli logical-system`** ***`R5`*** command.

```
lab@mxC> set cli logical-system R5
Logical system: R5

lab@mxC:R5>
```

Step 2.27

Use the **show bgp summary** command to verify the state of the R5 router's BGP sessions.

```
lab@mxC:R5> show bgp summary
Groups: 2 Peers: 3 Down peers: 2
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                       9          9          0          0          0          0
inet6.0
                       1          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.42         65532        823        815       0       0     6:04:44 Establ
  inet.0: 9/9/9/0
  inet6.0: 1/1/1/0
192.168.71.2           65501          0          0       0       0     6:04:48 Connect
192.168.71.4           65501          0          0       0       0     6:04:48 Connect
```

Question: Are all the BGP neighborships established?

_____

_____

Answer: No, the BGP sessions to R2 and R4 are showing in a `Connect` state. You should have expected this based on the results from troubleshooting the sessions from R2 and R4.

Question: What does the `Connect` state mean?

_____

_____

Answer: The `Connect` state means that the open messages have been sent and we are waiting on a reply from the neighbor. You know that R2 and R4 do not have a route back to R5's loopback so they are unable to respond to the open messages.

Note

Based on your troubleshooting on R2 and R4, you know that R5 is not sending its loopback address in its IS-IS link-state advertisements.

Step 2.28

Use the **show configuration protocols isis** command to verify the current IS-IS configuration.

```
lab@mxC:R5> show configuration protocols isis
level 1 disable;
```

```
interface ge-0/0/0.0;
inactive: interface lo0.0;

lab@mxC:R5>
```

Question: Do you see any problems with the current IS-IS configuration?

_____

_____

Answer: You should notice that the R5 router's loopback address has been deactivated. This explains why the R5 router is not advertising its loopback address to its neighbors.

Step 2.29

Enter into configuration mode and enable the loopback interface under the IS-IS protocol. Make sure you commit your changes and exit back into operational mode.

```
lab@mxC:R5> configure
Entering configuration mode

[edit]
lab@mxC:R5# activate protocols isis interface lo0.0

[edit]
lab@mxC:R5# commit and-quit
commit complete
Exiting configuration mode

lab@mxC:R5>
```

Step 2.30

Use the **show bgp summary** command to verify your recent changes have resolved the BGP peering issues.

```
lab@mxC:R5> show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      53         22          0          0          0          0
inet6.0
                       2          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.42         65532        830        829       0       0     6:07:48 Establ
  inet.0: 8/9/9/0
  inet6.0: 1/1/1/0
192.168.71.2          65501         18         11       0       0          16 Establ
  inet.0: 14/22/22/0
  inet6.0: 1/1/1/0
192.168.71.4          65501         15         11       0       0          22 Establ
  inet.0: 0/22/22/0
```

Question: What is the current status of the BGP peering sessions on the R5 router?

_____

_____

Answer: All of the R5 router's BGP sessions should be up and established correctly now.

Now that all our BGP sessions are up and functioning the next step is to verify that all routes are being accepted and are not being hidden. We should also verify that there are no missing NLRI families or other reachability problems in our network.

You should start your verification at the route reflectors because they should receive all advertised routes and are responsible for distributing the appropriate routes to the IBGP peers.

Step 2.31

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 2.32

Use the **`show route summary`** command to review the current status of routes on this device.

```
lab@mxB:R2> show route summary
Autonomous system number: 65501
Router ID: 192.168.71.2

inet.0: 35 destinations, 44 routes (35 active, 0 holddown, 8 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      4 routes,      4 active
               Local:      3 routes,      3 active
                 BGP:     31 routes,     22 active
               IS-IS:      6 routes,      6 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active

inet6.0: 16 destinations, 17 routes (15 active, 0 holddown, 2 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      3 routes,      3 active
               Local:      3 routes,      3 active
                 BGP:      3 routes,      1 active
               IS-IS:      7 routes,      7 active
               INET6:      1 routes,      1 active
```

Question: Do you have hidden routes?

_____

Answer: Yes, you should notice that you have both IPv4 and IPv6 hidden routes.

Question: Can you tell based on this output what device advertised these routes to you?

_____

Answer: No, there is not enough information in this output to determine from where the routes are coming.

Step 2.33

Use the **show route hidden** command to get more details about the hidden routes in the R2 router's routing table. You should be able to determine which peer is advertising these routes to the R2 router.

```
lab@mxB:R2> show route hidden

inet.0: 35 destinations, 44 routes (35 active, 0 holddown, 8 hidden)
+ = Active Route, - = Last Active, * = Both

172.21.1.0/24          [BGP/170] 13:35:21, localpref 100, from 192.168.71.5
                         AS path: 65532 I, validation-state: unverified
                         Unusable
172.21.2.0/24          [BGP/170] 13:35:21, localpref 100, from 192.168.71.5
                         AS path: 65532 I, validation-state: unverified
                         Unusable
172.21.3.0/24          [BGP/170] 13:35:21, localpref 100, from 192.168.71.5
                         AS path: 65532 I, validation-state: unverified
                         Unusable
172.21.4.0/24          [BGP/170] 13:35:21, localpref 100, from 192.168.71.5
                         AS path: 65532 6 7 7 I, validation-state: unverified
                         Unusable
172.21.5.0/24          [BGP/170] 13:35:21, localpref 100, from 192.168.71.5
                         AS path: 65532 6 7 31 8 I, validation-state: unverified
                         Unusable
172.21.6.0/24          [BGP/170] 13:35:21, localpref 100, from 192.168.71.5
                         AS path: 65532 6 7 32 9 I, validation-state: unverified
                         Unusable
172.21.7.0/24          [BGP/170] 13:35:21, localpref 100, from 192.168.71.5
                         AS path: 65532 6 7 31 32 I, validation-state: unverified
                         Unusable
172.22.0.0/24          [BGP/170] 13:35:21, localpref 100, from 192.168.71.5
                         AS path: 65532 69 86 73 76 I, validation-state: unverified
                         Unusable

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 16 destinations, 17 routes (15 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
fd69:8:2::/48          [BGP/170] 13:35:21, localpref 100, from 192.168.71.5
                         AS path: 65532 I, validation-state: unverified
                         Unusable
                       [BGP/170] 20:01:37, localpref 100, from 192.168.71.3
                         AS path: 65532 I, validation-state: unverified
                         Unusable
```

Step 2.34

Use the **show route hidden extensive | match protocol** command to identify the protocol next hop for each of the hidden routes. After identifying the protocol next hop, you can determine if you have a route to this next hop using the **show route *address*** command.

```
lab@mxB:R2> show route hidden extensive | match protocol
                     Protocol next hop: 172.22.131.42
                     Protocol next hop: 172.22.131.42
                     Protocol next hop: 172.22.131.42
                     Protocol next hop: 172.22.131.42
                     Protocol next hop: 172.22.131.42
                     Protocol next hop: 172.22.131.42
                     Protocol next hop: 172.22.131.42
                     Protocol next hop: 172.22.131.42
                     Protocol next hop: ::ffff:172.22.131.42
                     Protocol next hop: ::ffff:192.168.71.3

lab@mxB:R2> show route 172.22.131.42

lab@mxB:R2> show route ::ffff:172.22.131.42

lab@mxB:R2> show route ::ffff:192.168.71.3

lab@mxB:R2>
```

Question: Do you have routes to the protocol next hops?

_____

Answer: No, you do not have a route to any of the protocol next hops for the hidden routes.

Question: On which devices are the protocol next hops defined?

_____

_____

_____

_____

Answer: Based on the network diagram, the IPv4 protocol next-hop address is configured on the R5 router's EBGP neighbor's peering interface. The first IPv6 protocol next-hop address is also configured on the R5 router's EBGP neighbor's peering interface. The second IPv6 protocol next-hop address is IPv4-mapped IPv6 address from the R3 router's loopback interface.

Question: Why are the routes from the R5 router showing the EBGP peers interface address as the protocol next hop?

_____

_____

Answer: You see this problem because IBGP does not change the next hop when advertising to internal peers by default. To avoid this problem, you must either create a policy to change the next-hop or include the EBGP facing interface in your AS's IGP while also setting it to **`passive`**.

Question: Why is IPv6 route's protocol next hop unreachable?

_____

_____

Answer: The next hop could be unreachable because it appears that the R3 router is changing the next hop to its IPv4-mapped IPv6 loopback address, but the R3 router does not seem to be distributing this address to its neighbors.

Step 2.35

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

From the SSH session with the mxC router, use the **`show configuration policy-options`** command to review the current policies configured on R5. If there is a policy defined to change the next hop you should use the **`show configuration protocols bgp group`** ***`ibgp-RR`*** command to verify if the policy has been applied to the BGP group. If there is not a policy then you should determine if the EBGP peering interface is included in your IGP as a passive interface.

```
lab@mxC:R5> show configuration policy-options
policy-statement next-hop-self {
    from {
        protocol bgp;
        route-type external;
    }
    then {
        next-hop self;
    }
}

lab@mxC:R5> show configuration protocols bgp group ibgp-RR
type internal;
local-address 192.168.71.5;
family inet {
    unicast;
}
family inet6 {
    unicast;
}
neighbor 192.168.71.2;
neighbor 192.168.71.4;
```

Question: Do you have a policy designed to alter the next hop on external BGP routes?

_____

_____

Answer: Yes, you should have a policy designed to alter the next hop for EBGP routes.

Question: Is the policy applied to your IBGP peering sessions?

_____

_____

Answer: No, the policy has not been applied to the IBGP peering sessions.

Step 2.36

Enter into configuration mode and navigate to the `[edit protocols bgp group ibgp-RR]` hierarchy and apply the policy as an export policy. Commit and exit configuration mode after you are satisfied with your configuration changes.

```
lab@mxC:R5> configure
Entering configuration mode

[edit]
lab@mxC:R5# edit protocols bgp group ibgp-RR

[edit protocols bgp group ibgp-RR]
lab@mxC:R5# set export next-hop-self

[edit protocols bgp group ibgp-RR]
lab@mxC:R5# show
type internal;
local-address 192.168.71.5;
family inet {
    unicast;
}
family inet6 {
    unicast;
}
export next-hop-self;
neighbor 192.168.71.2;
neighbor 192.168.71.4;

[edit protocols bgp group ibgp-RR]
lab@mxC:R5# commit and-quit
commit complete
Exiting configuration mode

lab@mxC:R5>
```

Step 2.37

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

From the SSH session with the mxB router, use the **`show route summary`** command to review the current status of routes on this device and determine if the policy you applied on the R5 router has resolved the problem causing hidden routes from that peer.

```
lab@mxB:R2> show route summary
Autonomous system number: 65501
Router ID: 192.168.71.2

inet.0: 35 destinations, 43 routes (35 active, 0 holddown, 0 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      4 routes,      4 active
               Local:      3 routes,      3 active
                 BGP:     30 routes,     22 active
               IS-IS:      6 routes,      6 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active

inet6.0: 16 destinations, 17 routes (16 active, 0 holddown, 1 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      3 routes,      3 active
               Local:      3 routes,      3 active
                 BGP:      3 routes,      2 active
               IS-IS:      7 routes,      7 active
               INET6:      1 routes,      1 active
```

Question: Do you still have hidden routes?

_____

_____

Answer: Yes, you will still have one hidden route. You should remember that you still have a hidden route that was learned from the R3 router. You should notice that the hidden IPv4 routes are no longer there as well as the one IPv6 route that we are learning from the R5 router.

Step 2.38

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 2.39

Use the **`show configuration protocols bgp group`** ***`ibgp-RR`*** command to review the current BGP configuration to ensure the next-hop-self policy is applied. Then review the loopback interface configuration to verify the IPv4-mapped IPv6 address has been configured.

```
lab@mxB:R3> show configuration protocols bgp group ibgp-RR
type internal;
local-address 192.168.71.3;
family inet {
    unicast;
}
family inet6 {
    unicast;
}
export next-hop-self;
neighbor 192.168.71.2;
neighbor 192.168.71.4;

lab@mxB:R3> show configuration interfaces lo0
unit 3 {
    family inet {
        address 192.168.71.3/32;
    }
    family iso {
        address 49.0001.0000.0000.7803.00;
    }
    family inet6 {
        address 2001:db8:0101::3/128 {
            primary;
        }
        inactive: address ::ffff:192.168.71.3/128;
    }
}
```

Question: Is the next-hop policy applied?

_____

Answer: Yes, the policy is applied.

Question: Is the IPv4-mapped IPv6 loopback address defined on the loopback interface?

_____

Answer: Yes, The IPv4-mapped IPv6 address is defined on R3's loopback interface but this address has been deactivated.

Step 2.40

Enter into configuration mode and activate the IPv4-mapped IPv6 address on the loopback address. Commit and exit to operational mode when you are satisfied with your configuration changes.

```
lab@mxB:R3> configure
Entering configuration mode

[edit]
lab@mxB:R3# activate interfaces lo0 unit 3 family inet6 address
   ::ffff:192.168.71.3/128
```

```
[edit]
lab@mxB:R3# show interfaces lo0
unit 3 {
    family inet {
        address 192.168.71.3/32;
    }
    family iso {
        address 49.0001.0000.0000.7803.00;
    }
    family inet6 {
        address 2001:db8:0101::3/128 {
            primary;
        }
        address ::ffff:192.168.71.3/128;
    }
}

[edit]
lab@mxB:R3# commit and-quit
commit complete
Exiting configuration mode

lab@mxB:R3>
```

Step 2.41

Change the perspective of the CLI to that of the R2 logical system by issuing the **`set cli logical-system R2`** command.

```
lab@mxB:R3> set cli logical-system R2
Logical system: R2

lab@mxB:R2>
```

Step 2.42

Use the **`show route summary`** command to review the current status of routes on this device and determine if activating the loopback IPv4-mapped IPv6 address on the R3 router has resolved the problem causing the hidden route from that peer.

```
lab@mxB:R2> show route summary
Autonomous system number: 65501
Router ID: 192.168.71.2

inet.0: 35 destinations, 43 routes (35 active, 0 holddown, 0 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      4 routes,      4 active
               Local:      3 routes,      3 active
                 BGP:     30 routes,     22 active
               IS-IS:      6 routes,      6 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active
```

```
inet6.0: 17 destinations, 18 routes (17 active, 0 holddown, 0 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      3 routes,      3 active
               Local:      3 routes,      3 active
                 BGP:      3 routes,      2 active
               IS-IS:      8 routes,      8 active
               INET6:      1 routes,      1 active
```

Question: Do you have any hidden routes?

_____

Answer: No, you should not see hidden routes.

Step 2.43

Change the perspective of the CLI to that of the R4 logical system by issuing the **`set cli logical-system R4`** command.

```
lab@mxB:R2> set cli logical-system R4
Logical system: R4

lab@mxB:R4>
```

Step 2.44

Use the **`show route summary`** command to review the current status of routes on this device.

```
lab@mxB:R4> show route summary
Autonomous system number: 65501
Router ID: 192.168.71.4

inet.0: 35 destinations, 43 routes (35 active, 0 holddown, 0 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      4 routes,      4 active
               Local:      3 routes,      3 active
                 BGP:     30 routes,     22 active
               IS-IS:      6 routes,      6 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active

inet6.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      3 routes,      3 active
               Local:      3 routes,      3 active
               IS-IS:      8 routes,      8 active
               INET6:      1 routes,      1 active
```

Question: Do you see any hidden routes?

_____

Answer: No, you should not see any hidden routes. The alterations we made to resolve issues on the R2 router have ensured that we do not have the same problems on the R4 route reflector.

Question: Do you see all the BGP routes you had on the R2 router?

_____

Answer: No, you should notice that there are no IPv6 BGP routes on the R4 router.

Question: What are some possible causes?

_____

_____

_____

Answer: You might not have the IPv6 peer configured in BGP. You might be missing the IPv6 NLRI family if you are using the IPv4 sessions to transport the IPv6 addresses. These are just some possible causes.

Step 2.45

Enter into configuration mode and navigate to the `[edit protocols bgp]` hierarchy and review your BGP peering configurations.

```
lab@mxB:R4> configure
Entering configuration mode

[edit]
lab@mxB:R4# edit protocols bgp

[edit protocols bgp]
lab@mxB:R4# show
group ibgp {
    type internal;
    local-address 192.168.71.4;
    neighbor 192.168.71.2;
}
group ibgp-clients {
    type internal;
    local-address 192.168.71.4;
    cluster 1.1.1.1;
    neighbor 192.168.71.1;
    neighbor 192.168.71.3;
    neighbor 192.168.71.5;
}
```

Question: Do you see any problems?

_____

_____

Answer: You should notice that there are no IPv6 sessions or no IPv6 NLRIs have been enabled for the IBGP sessions.

Step 2.46

Configure the IPv6 family in your ***ibgp*** and **ibgp-clients** BGP groups. This configuration will ensure that your IPv6 routes are transported over your IPv4 BGP sessions. Remember that you also have to explicitly define the `inet` NLRI when defining another family for the BGP sessions. Commit and exit to operational mode when you are satisfied with your configuration changes.

```
[edit protocols bgp]
lab@mxB:R4# set group ibgp family inet unicast

[edit protocols bgp]
lab@mxB:R4# set group ibgp family inet6 unicast

[edit protocols bgp]
lab@mxB:R4# set group ibgp-clients family inet unicast

[edit protocols bgp]
lab@mxB:R4# set group ibgp-clients family inet6 unicast

[edit protocols bgp]
lab@mxB:R4# show
group ibgp {
    type internal;
    local-address 192.168.71.4;
    family inet {
        unicast;
    }
    family inet6 {
        unicast;
    }
    neighbor 192.168.71.2;
}
group ibgp-clients {
    type internal;
    local-address 192.168.71.4;
    family inet {
        unicast;
    }
    family inet6 {
        unicast;
    }
    cluster 1.1.1.1;
    neighbor 192.168.71.1;
    neighbor 192.168.71.3;
    neighbor 192.168.71.5;
}
```

```
[edit protocols bgp]
lab@mxB:R4# commit and-quit
commit complete
Exiting configuration mode

lab@mxB:R4>
```

Step 2.47

Use the **show route summary** command to review the current status of routes on this device and determine if activating the IPv6 NLRI for you BGP sessions has resolved the problem causing the missing IPv6 routes from neighbors.

```
lab@mxB:R4> show route summary
Autonomous system number: 65501
Router ID: 192.168.71.4

inet.0: 35 destinations, 43 routes (35 active, 0 holddown, 0 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      4 routes,      4 active
               Local:      3 routes,      3 active
                 BGP:     30 routes,     22 active
               IS-IS:      6 routes,      6 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active

inet6.0: 17 destinations, 18 routes (17 active, 0 holddown, 0 hidden)
Limit/Threshold: 32768/32768 destinations
              Direct:      3 routes,      3 active
               Local:      3 routes,      3 active
                 BGP:      3 routes,      2 active
               IS-IS:      8 routes,      8 active
               INET6:      1 routes,      1 active
```

Question: Do you have IPv6 BGP routes?

_____

Answer: Yes, you should now have IPv6 BGP routes.

Step 2.48

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

```
lab@mxB:R4> exit


mxB (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

Network Management Diagram

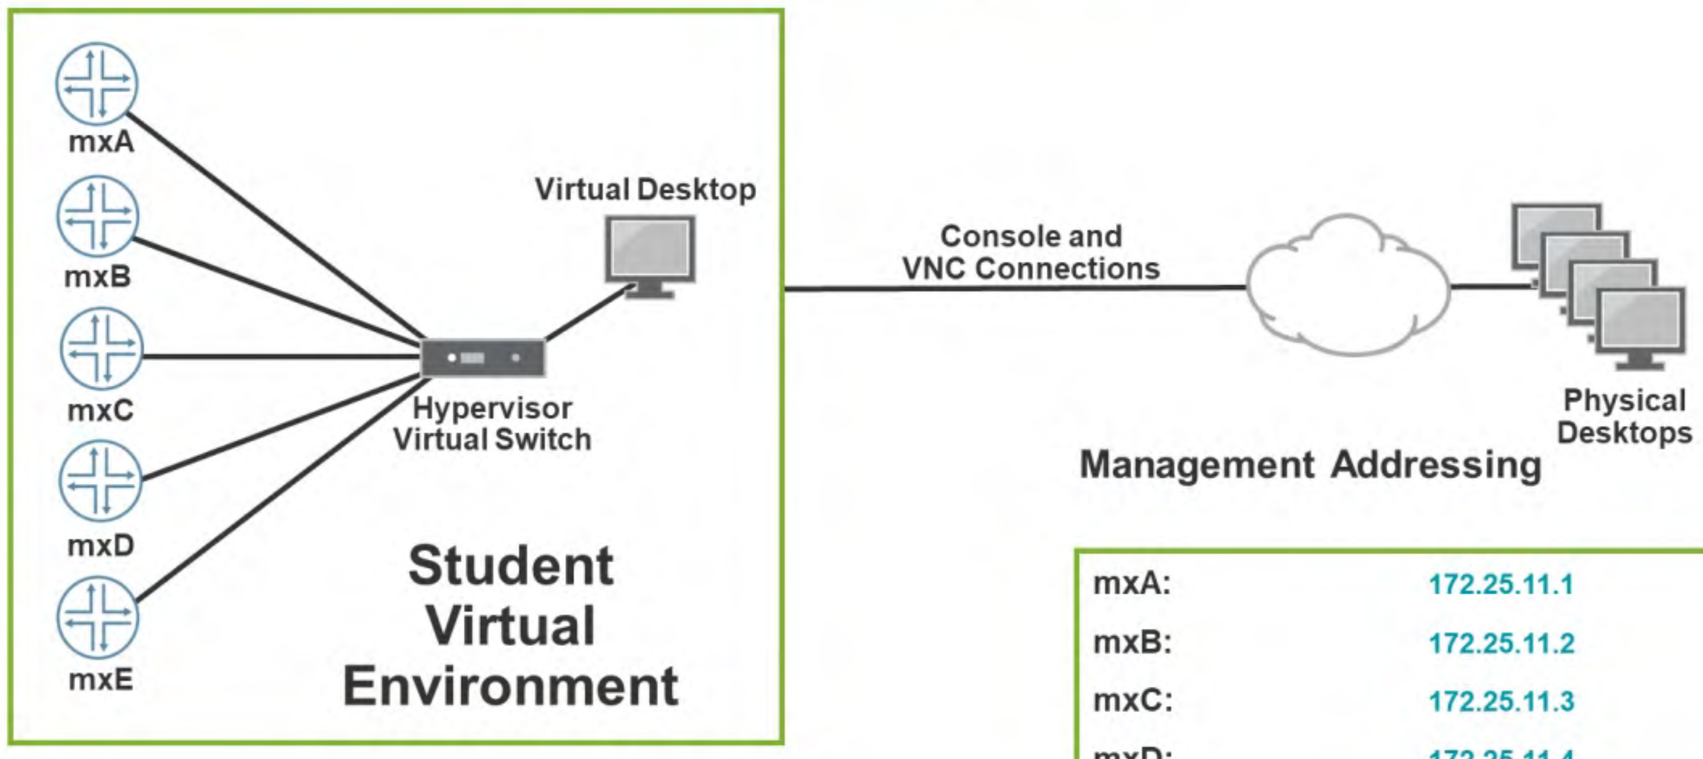

Management Addressing

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: BGP Troubleshooting

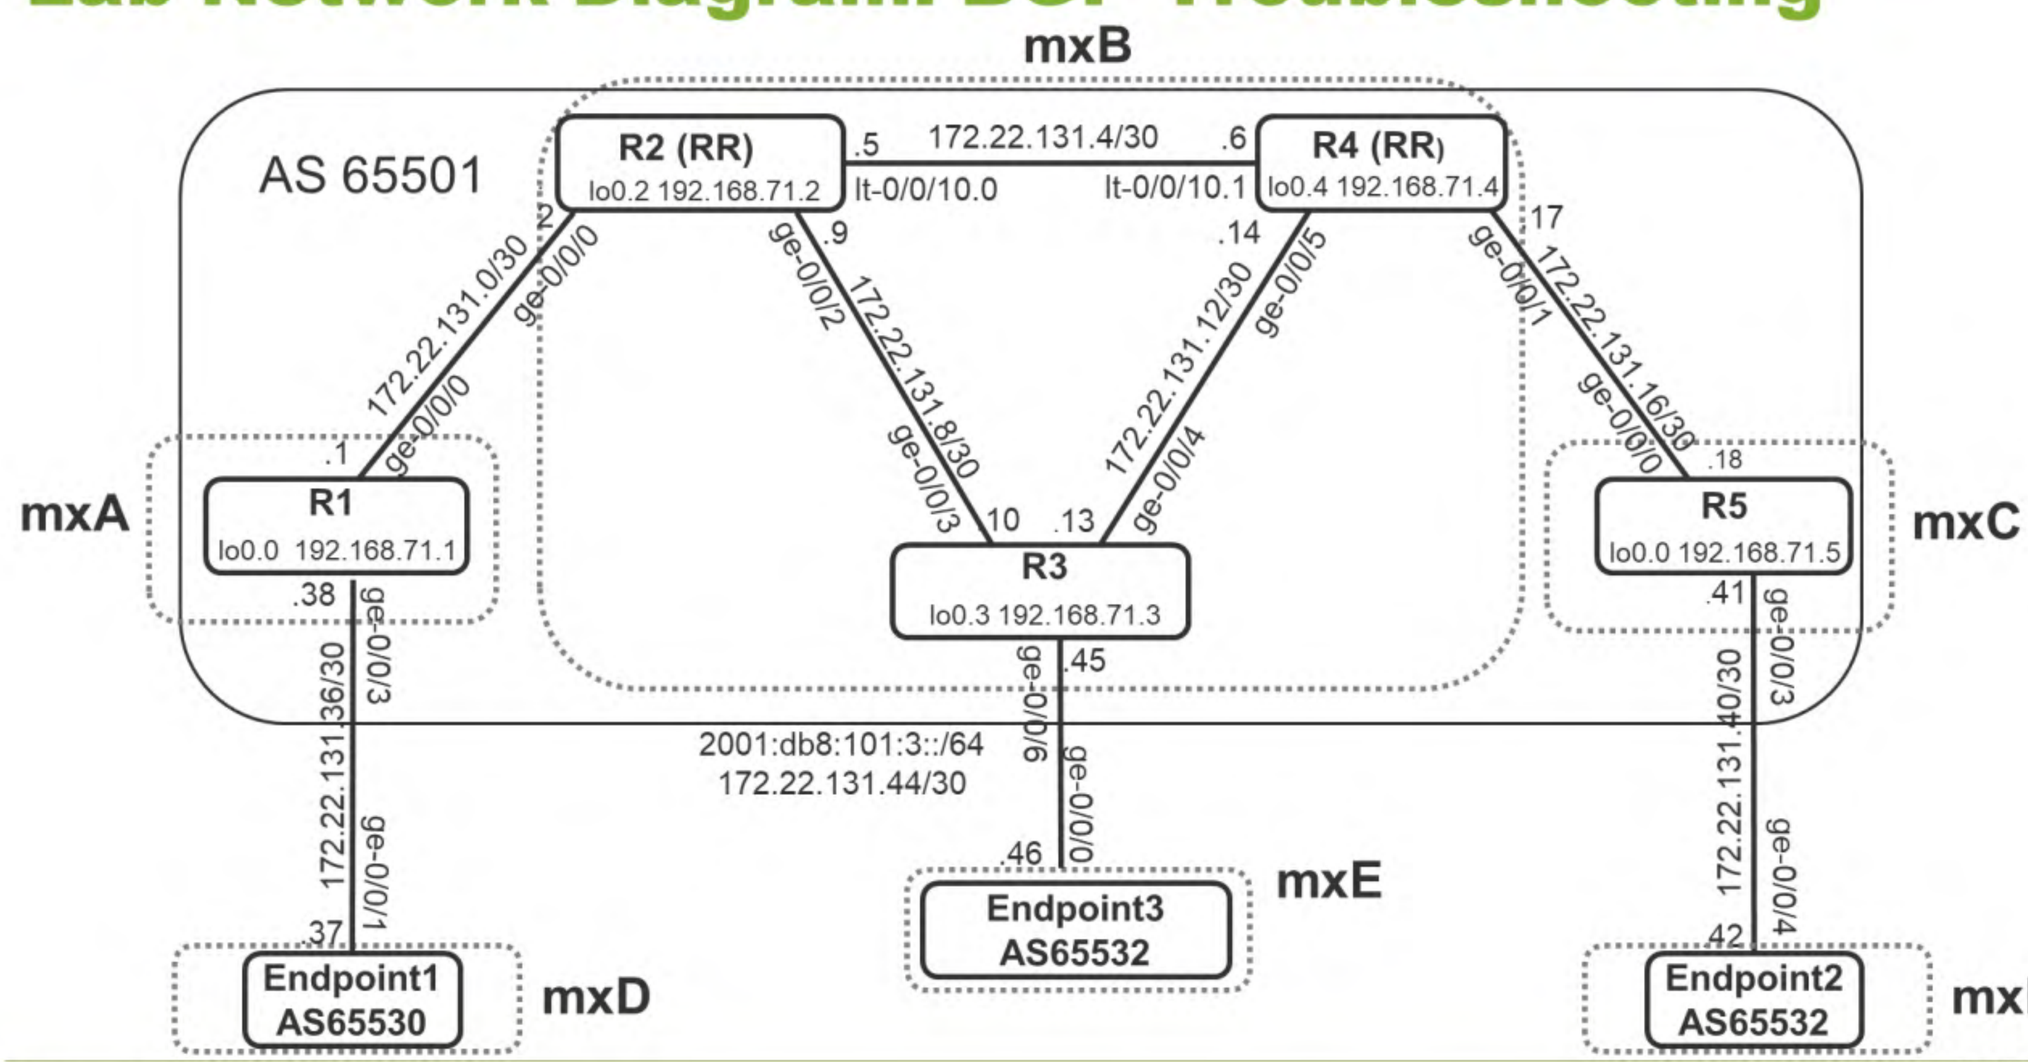

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab

Troubleshooting Routing Policy

Overview

In this lab, you use the lab diagram titled "Lab Network Diagram: Troubleshooting Routing Policy" to verify and troubleshoot routing problems. You will ensure routing policies are correctly configured and implemented. As the network administrator, you must adhere to certain routing requirements.

The local AS honors a set of communities to influence AS_PATH prepending when advertising routes to peers:

- Community 65501:3 is used to prepend the local AS three times before advertising to peers.
- Community 65501:6 is used to prepend the local AS six times before advertising to peers.
- Community 65501:9 is used to prepend the local AS nine times before advertising to peers.

In addition, there are specific requirements for each of the peering points:

- EBGP session to AS65530 on R1
  - Inbound: Routes transiting through in AS73 should be rejected; all other routes should be tagged with community 65501:100
- EBGP session to AS65532 on R3
  - Inbound: no route transiting either AS 31 or 32 should be accepted; accepted routes should be tagged with community 65501:200
  - Outbound: no route tagged with community 65501:100 and more specific than /24 should be sent
- EBGP session to AS65532 on R5
  - Inbound: route originating from AS3 should be rejected; accepted routes should be tagged with community 65530:200

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

- Troubleshoot import and export policies to ensure that you are meeting the defined requirements.
- You will fix any issues that you come across during the troubleshooting steps.

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 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 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/lab15-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 27 17:20:28 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/lab15-start.config
load complete

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

lab@mxD>
```

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/lab15-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: Wed Nov 27 17:20:28 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/lab15-start.config
load complete

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

lab@mxE>
```

Step 1.4

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/lab15-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 27 17:20:28 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/lab15-start.config
load complete

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

lab@mxC>
```

Step 1.5

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/lab15-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 27 17:20:28 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/lab15-start.config
load complete

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

lab@mxB>
```

Step 1.6

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/lab15-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: Wed Nov 27 17:20:28 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/lab15-start.config
load complete

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

lab@mxA>
```

Part 2: Troubleshooting Policies

In this lab part, you will be use CLI commands to discover routing issues. Once discovered, you will use various CLI tools to troubleshoot and correct import and export policies.

Step 2.1

Change the perspective of the CLI to that of the R1 logical system by issuing the **`set cli logical-system R1`** command.

```
lab@mxA> set cli logical-system R1
Logical system: R1

lab@mxA:R1>
```

Step 2.2

Use the **`show bgp summary`** command to view the BGP neighborships and routes being learned on the R1 router.

```
lab@mxA:R1> show bgp summary
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      29          9          0          0          0          0
inet6.0
                       3          1          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.37         65530        278        289       0       0     1:59:42 Establ
  inet.0: 0/11/0/0
  inet6.0: 0/1/0/0
192.168.71.2          65501        300        275       0       0     1:59:09 Establ
  inet.0: 9/9/9/0
  inet6.0: 1/1/1/0
192.168.71.4          65501        127         83       0       1       36:42 Establ
  inet.0: 0/9/9/0
  inet6.0: 0/1/1/0
```

Question: How many prefixes are you receiving from your EBGP peer?

_____

_____

Answer: You should see 11 IPv4 routes and 1 IPv6 route being learned from the R1 router's EBGP peer.

Question: How many of the routes learned from the Endpoint1 EBGP peer are active?

_____

_____

Answer: According to the output, none of the BGP routes learned from the Endpoint1 EBGP peer are active in the routing table.

Step 2.3

Use the **show route receive-protocol bgp 172.22.131.37** to examine the routes learned from the EBGP peer.

```
lab@mxA:R1> show route receive-protocol bgp 172.22.131.37

inet.0: 32 destinations, 42 routes (22 active, 0 holddown, 11 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 20 destinations, 21 routes (19 active, 0 holddown, 1 hidden)
```

Question: Where are the 11 routes that we saw in the previous output?

_____

_____

Answer: The routes are there, they are just marked as hidden.

Step 2.4

Use the **show route receive-protocol bgp 172.22.131.37 hidden detail** command to review the received routes that are being marked as hidden.

```
lab@mxA:R1> show route receive-protocol bgp 172.22.131.37 hidden detail

inet.0: 32 destinations, 42 routes (22 active, 0 holddown, 11 hidden)
  172.20.0.0/24 (1 entry, 0 announced)
     Nexthop: 172.22.131.37
     AS path: 65530 I
     Communities: 65501:3
     Hidden reason: rejected by import policy

  172.20.0.0/25 (1 entry, 0 announced)
     Nexthop: 172.22.131.37
     AS path: 65530 I
     Communities: 65501:3
     Hidden reason: rejected by import policy

  172.20.0.128/25 (1 entry, 0 announced)
     Nexthop: 172.22.131.37
     AS path: 65530 I
     Communities: 65501:3
     Hidden reason: rejected by import policy

  172.20.1.0/24 (1 entry, 0 announced)
     Nexthop: 172.22.131.37
     AS path: 65530 I
     Communities: 65501:3
     Hidden reason: rejected by import policy

  172.20.2.0/24 (1 entry, 0 announced)
     Nexthop: 172.22.131.37
     AS path: 65530 I
```

```
   Hidden reason: rejected by import policy

172.20.3.0/24 (1 entry, 0 announced)
   Nexthop: 172.22.131.37
   AS path: 65530 I
   Hidden reason: rejected by import policy

172.20.4.0/24 (1 entry, 0 announced)
   Nexthop: 172.22.131.37
   AS path: 65530 1 2 2 I
   Hidden reason: rejected by import policy

172.20.5.0/24 (1 entry, 0 announced)
   Nexthop: 172.22.131.37
   AS path: 65530 1 2 3 3 4 I
   Hidden reason: rejected by import policy

172.20.6.0/24 (1 entry, 0 announced)
   Nexthop: 172.22.131.37
   AS path: 65530 1 2 3 3 4 5 I
   Hidden reason: rejected by import policy

172.20.7.0/24 (1 entry, 0 announced)
   Nexthop: 172.22.131.37
   AS path: 65530 1 2 2 3 I
   Hidden reason: rejected by import policy

172.22.0.0/24 (3 entries, 1 announced)
   Nexthop: 172.22.131.37
   MED: 10
   AS path: 65530 69 86 73 76 I
   Hidden reason: rejected by import policy

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 20 destinations, 21 routes (19 active, 0 holddown, 1 hidden)

fd69:8:1::/48 (1 entry, 0 announced)
   Nexthop: ::ffff:172.22.131.37
   AS path: 65530 I
   Hidden reason: rejected by import policy
```

Question: Why are these routes hidden?

_____

_____

Answer: These routes are hidden because they have all been rejected by import policy.

Step 2.5

Review the import policies applied to the EBGP group. If there are additional parameters like `as-path` and `community` tags being used as match conditions you should investigate them also.

```
lab@mxA:R1> show configuration protocols bgp group ebgp-AS65530
import [ reject-73 tag-community1 ];
family inet {
    unicast;
}
family inet6 {
    unicast;
}
export prepend;
peer-as 65530;
neighbor 172.22.131.37;

lab@mxA:R1> show configuration policy-options policy-statement reject-73
term reject-73 {
    from as-path transit-73;
    then reject;
}

lab@mxA:R1> show configuration policy-options policy-statement tag-community1
term tag {
    then {
        community set c1;
    }
}

lab@mxA:R1> show configuration policy-options as-path transit-73
".* 73*";
```

Question: Do you see any issues with the import policies being applied as import policies?

_____

_____

Answer: The **as-path** definition of *transit-73* is incorrect. The current definition matches all AS numbers.

Step 2.6

Enter into configuration mode and navigate to the [edit policy-options] hierarchy. Correct the as-path definition for *transit-73* by altering the defined path. Use **".* 73 .*"** as the new AS path. Commit and exit to operational mode when you are finished.

```
lab@mxA:R1> configure
Entering configuration mode

[edit]
lab@mxA:R1# edit policy-options

[edit policy-options]
lab@mxA:R1# set as-path transit-73 ".* 73 .*"

[edit policy-options]
lab@mxA:R1# commit and-quit
commit complete
```

```
Exiting configuration mode

lab@mxA:R1>
```

Step 2.7

Use the **`show route receive-protocol bgp 172.22.131.37 hidden`** command to verify that the only hidden routes are routes with the transit 73 AS number in their AS path.

```
lab@mxA:R1> show route receive-protocol bgp 172.22.131.37 hidden

inet.0: 32 destinations, 42 routes (32 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.22.0.0/24           172.22.131.37        10                 65530 69 86 73 76 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 20 destinations, 21 routes (20 active, 0 holddown, 0 hidden)
```

Question: Are there any hidden routes that do not have 73 in their AS path?

_____

_____

Answer: No, you should only see one route and that route should have the AS 73 in the path.

Step 2.8

Use the **`show route receive-protocol bgp 172.22.131.37 community "65501:.*[123]" extensive`** command to review routes that are coming from the Endpoint1 router. This command is designed to look for the routes with the `prepend` community attached to it. This community is used to add our AS number based on the community definition before it is advertised to other EBGP peer on other routers.

```
lab@mxA:R1> show route receive-protocol bgp 172.22.131.37 community
   "65501:.*[123]" extensive

inet.0: 32 destinations, 42 routes (32 active, 0 holddown, 1 hidden)
* 172.20.0.0/24 (1 entry, 1 announced)
     Accepted
     Nexthop: 172.22.131.37
     AS path: 65530 I
     Communities: 65501:3

* 172.20.0.0/25 (1 entry, 1 announced)
     Accepted
     Nexthop: 172.22.131.37
     AS path: 65530 I
     Communities: 65501:3

* 172.20.0.128/25 (1 entry, 1 announced)
     Accepted
     Nexthop: 172.22.131.37
     AS path: 65530 I
```

```
     Communities: 65501:3

* 172.20.1.0/24 (1 entry, 1 announced)
     Accepted
     Nexthop: 172.22.131.37
     AS path: 65530 I
     Communities: 65501:3

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 20 destinations, 21 routes (20 active, 0 holddown, 0 hidden)
```

Question: Do you see any routes tagged with community `65501:3`?

_____

_____

Answer: Yes, you should see four routes that have the `65501:3` community applied.

Step 2.9

Use the **`show route 172.20.0.0/24 extensive | match com`** to view the community values assigned to these routes that are installed in your routing table.

```
lab@mxA:R1> show route 172.20.0.0/24 extensive | match com
     Communities: 65501:100
                Communities: 65501:100
     Communities: 65501:100
                Communities: 65501:100
     Communities: 65501:100
                Communities: 65501:100
```

Question: Do you see the `65501:3` community value assigned to any of these routes?

_____

_____

Answer: No, for some reason the routes installed in your routing table do not have the `65501:3` community tag attached to the routes.

Step 2.10

Review the import policies for the EBGP neighbor. Review the applied policies to determine why the `65501:3` community values are being removed.

```
lab@mxA:R1> show configuration protocols bgp group ebgp-AS65530
import [ reject-73 tag-community1 ];
family inet {
    unicast;
}
family inet6 {
```

```
    unicast;
}
export prepend;
peer-as 65530;
neighbor 172.22.131.37;

lab@mxA:R1> show configuration policy-options policy-statement tag-community1
term tag {
    then {
        community set c1;
    }
}

lab@mxA:R1> show configuration policy-options community c1
members 65501:100;
```

Question: Do you see any problems with the `tag-community1` policy or any of its components?

_____

_____

Answer: Yes, the policy action is `community set` which will replace all communities already applied to the route.

Question: How do you fix this so that it only adds another community value and leaves the current values attached?

_____

_____

Answer: To resolve this issue you need to use `community add` instead of `set`.

Step 2.11

Enter into configuration mode and navigate to the `[edit policy-options policy-statement tag-community1]` hierarchy. Change the policy action to add a community value instead of set the community. After making the correct changes **commit** your configuration and exit to operational mode.

```
lab@mxA:R1> configure
Entering configuration mode

[edit]
lab@mxA:R1# edit policy-options policy-statement tag-community1

[edit policy-options policy-statement tag-community1]
lab@mxA:R1# show
term tag {
    then {
        community set c1;
    }
}
```

```
[edit policy-options policy-statement tag-community1]
lab@mxA:R1# set term tag then community add c1

[edit policy-options policy-statement tag-community1]
lab@mxA:R1# delete term tag then community set c1

[edit policy-options policy-statement tag-community1]
lab@mxA:R1# show
term tag {
    then {
        community add c1;
    }
}

[edit policy-options policy-statement tag-community1]
lab@mxA:R1# commit and-quit
commit complete
Exiting configuration mode

lab@mxA:R1>
```

Step 2.12

Use the **`show route 172.20.0.0/24 extensive | match com`** to view the community values assigned to these routes that are installed in your routing table.

```
lab@mxA:R1> show route 172.20.0.0/24 extensive | match com
     Communities: 65501:3 65501:100
                Communities: 65501:3 65501:100
     Communities: 65501:3 65501:100
                Communities: 65501:3 65501:100
     Communities: 65501:3 65501:100
                Communities: 65501:3 65501:100
```

Question: Do you see the `65501:3` community tags?

_____

_____

Answer: Yes, you should see the community attached to the routes learned from your EBGP peer.

Step 2.13

Use the **`show route advertising-protocol bgp 172.22.131.37 community 65501:3`** command to verify that you are prepending your AS number three times on the routes tagged with the community *`65501:3`* that you are advertising to your EBGP peer.

```
lab@mxA:R1> show route advertising-protocol bgp 172.22.131.37 community 65501:3

inet.0: 32 destinations, 42 routes (32 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.21.0.0/24           Self                                    65501 65501 65501
   [65501] 65532 5 5 3 I
* 172.21.1.0/24           Self                                    65501 65501 65501
   [65501] 65532 I
```

```
* 172.21.3.0/24          Self                                    65501 65501 65501
   [65501] 65532 I
* 172.21.4.0/24          Self                                    65501 65501 65501
   [65501] 65532 6 7 7 I

inet6.0: 20 destinations, 21 routes (20 active, 0 holddown, 0 hidden)
```

Question: Do you see your internal AS number prepended on each route returned from the command?

_____

_____

Answer: Yes, you should see three additional instances of your internal AS number. The R1 router's routes and policies seem to be working correctly.

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> set cli logical-system R3
Logical system: R3

lab@mxB:R3>
```

Step 2.15

Use the **`show bgp summary`** command to view the BGP neighborships and routes being learned on the R3 logical system.

```
lab@mxB:R3> show bgp summary
Threading mode: BGP I/O
Groups: 3 Peers: 4 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet6.0
                       5          0          0          0          0          0
inet.0
                      50          0          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.46         65532         34         26       0       0       11:12 Establ
  inet.0: 0/12/0/0
192.168.71.2          65501         52         26       0       0       11:04 Establ
  inet.0: 0/19/0/0
  inet6.0: 0/2/0/0
192.168.71.4          65501         43         27       0       0       10:44 Establ
  inet.0: 0/19/0/0
  inet6.0: 0/2/0/0
2001:db8:101:3::46       65532         28         25       0       0       11:10
   Establ
  inet6.0: 0/1/0/0
```

Question: How many prefixes is the R3 router receiving from the Endpoint3 router?

_____

_____

Answer: You should see 12 IPv4 routes and 1 IPv6 route being learned from the Endpoint3 router.

Question: How many of the routes learned from Endpoint3 router are active in the routing table?

_____

_____

Answer: According to the output, none of the BGP routes learned from the Endpoint3 router are being used.

Step 2.16

Review the import policies applied to the EBGP group. If there are additional parameters like `as-path` and `community` tags being used as match conditions you should investigate them also.

```
lab@mxB:R3> show configuration protocols bgp
group ebgp-AS65532-ipv6 {
    type external;
    family inet6 {
        unicast;
    }
    peer-as 65532;
    neighbor 2001:db8:101:3::46;
}
group ebgp-AS65532-ipv4 {
    type external;
    family inet {
        unicast;
    }
    peer-as 65532;
    neighbor 172.22.131.46;
}
group ibgp-RR {
    type internal;
    local-address 192.168.71.3;
    family inet {
        unicast;
    }
    family inet6 {
        unicast;
    }
    export next-hop-self;
    neighbor 192.168.71.2;
    neighbor 192.168.71.4;
}
import [ no-31-32 tag-routes ];
```

```
export [ no-tagged-specifics prepend ];

lab@mxB:R3> show configuration policy-options policy-statement no-31-32
term exclude-routes {
    from as-path via-31-or-32;
}
then reject;

lab@mxB:R3> show configuration policy-options policy-statement tag-routes
term tag {
    then {
        community add c2;
    }
}

lab@mxB:R3> show configuration policy-options as-path via-31-or-32
".*3[12].*";
```

Question: Do you see any reason why the routes would not be accepted?

_____

Answer: Yes, you should notice that the `no-31-32` policy has a `reject` action that is outside the term with the conditions. This `reject` action will match all routes because there is no global condition.

Question: Did you notice anything else wrong with the policy or policy definitions.

_____

_____

Answer: You should notice that the as-path definition is wrong. It should be defined as two separate as-paths to avoid problems.

Step 2.17

Enter into configuration mode and navigate to the `[edit policy-options]` hierarchy. Make two unique as-paths called ***via-31*** and ***via-32*** for the different AS numbers. Remove the existing as-path that is incorrectly configured.

```
lab@mxB:R3> configure
Entering configuration mode

[edit]
lab@mxB:R3# edit policy-options

[edit policy-options]
lab@mxB:R3# show as-path via-31-or-32
".*3[12].*";

[edit policy-options]
lab@mxB:R3# set as-path via-31 ".* 31 .*"
```

```
[edit policy-options]
lab@mxB:R3# set as-path via-32 ".* 32 .*"

[edit policy-options]
lab@mxB:R3# delete as-path via-31-or-32

[edit policy-options]
lab@mxB:R3#
```

Step 2.18

Navigate to the `[edit policy-options policy-statement no-31-32]` hierarchy and delete the reject action that is outside the configured term. Fix the term condition to match the new as-paths that you just defined. Make sure you remove the one that was incorrect. After fixing the as-path match condition, add the action of reject to the term.

```
[edit policy-options]
lab@mxB:R3# edit policy-statement no-31-32

[edit policy-options policy-statement no-31-32]
lab@mxB:R3# show
term exclude-routes {
    from as-path via-31-or-32;
}
then reject;


[edit policy-options policy-statement no-31-32]
lab@mxB:R3# delete then reject

[edit policy-options policy-statement no-31-32]
lab@mxB:R3# set term exclude-routes from as-path via-31

[edit policy-options policy-statement no-31-32]
lab@mxB:R3# set term exclude-routes from as-path via-32

[edit policy-options policy-statement no-31-32]
lab@mxB:R3# delete term exclude-routes from as-path via-31-or-32

[edit policy-options policy-statement no-31-32]
lab@mxB:R3# set term exclude-routes then reject

[edit policy-options policy-statement no-31-32]
lab@mxB:R3# show
term exclude-routes {
    from as-path [ via-31 via-32 ];
    then reject;
}

[edit policy-options policy-statement no-31-32]
lab@mxB:R3# commit and-quit
commit complete
Exiting configuration mode

lab@mxB:R3>
```

Step 2.19

Use the **show route receive-protocol bgp 172.22.131.46 hidden** command to verify that only routes received from the Endpoint3 router that go through AS numbers 31 and/or 32 are rejected and marked as hidden.

```
lab@mxB:R3> show route receive-protocol bgp 172.22.131.46 hidden

inet.0: 37 destinations, 55 routes (34 active, 0 holddown, 9 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.21.5.0/24           172.22.131.46                           65532 6 7 31 8 I
  172.21.6.0/24           172.22.131.46                           65532 6 7 32 9 I
  172.21.7.0/24           172.22.131.46                          65532 6 7 31 32 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 19 destinations, 20 routes (19 active, 0 holddown, 0 hidden)
```

Question: Do you see any hidden routes that do not transit AS numbers 31 or 32?

_____

_____

Answer: No, you should only see hidden routes that have the 31 or 32 AS numbers in their AS path.

Step 2.20

Use the **show route advertising-protocol bgp 172.22.131.46** command to verify that you are not exporting anything to your EBGP that is longer than the required /24.

```
lab@mxB:R3> show route advertising-protocol bgp 172.22.131.46

inet.0: 37 destinations, 55 routes (34 active, 0 holddown, 9 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.20.0.0/24           Self                                   65501 65501 65501
   [65501] 65530 I
* 172.20.0.0/25           Self                                   65501 65501 65501
   [65501] 65530 I
* 172.20.0.128/25         Self                                   65501 65501 65501
   [65501] 65530 I
* 172.20.1.0/24           Self                                   65501 65501 65501
   [65501] 65530 I
* 172.20.2.0/24           Self                                     65530 I
* 172.20.3.0/24           Self                                     65530 I
* 172.20.4.0/24           Self                                     65530 1 2 2 I
* 172.20.5.0/24           Self                                    65530 1 2 3 3 4 I
* 172.20.6.0/24           Self                                  65530 1 2 3 3 4 5 I
* 172.20.7.0/24           Self                                     65530 1 2 2 3 I
```

Question: Do you see prefixes being advertised that are longer than /24?

_____

_____

Answer: Yes, you should see two routes that are /25s and they are still being advertised to your EBGP peer.

Step 2.21

Use the **show configuration protocols bgp** command to verify what policies you have applied to your BGP groups. Review the policies and determine why the /25 routes are still being advertised.

```
lab@mxB:R3> show configuration protocols bgp
group ebgp-AS65532-ipv6 {
    type external;
    family inet6 {
        unicast;
    }
    peer-as 65532;
    neighbor 2001:db8:101:3::46;
}
group ebgp-AS65532-ipv4 {
    type external;
    family inet {
        unicast;
    }
    peer-as 65532;
    neighbor 172.22.131.46;
}
group ibgp-RR {
    type internal;
    local-address 192.168.71.3;
    family inet {
        unicast;
    }
    family inet6 {
        unicast;
    }
    export next-hop-self;
    neighbor 192.168.71.2;
    neighbor 192.168.71.4;
}
import [ no-31-32 tag-routes ];
export [ no-tagged-specifics prepend ];

lab@mxB:R3> show configuration policy-options policy-statement no-tagged-specifics
term no-specifics {
    from {
        community c1;
        route-filter 0.0.0.0/24 upto /32;
    }
    then reject;
```

```
}

lab@mxB:R3> show configuration policy-options policy-statement prepend
term 3 {
    from community prepend-3;
    then as-path-prepend "65501 65501 65501";
}
term 6 {
    from community prepend-6;
    then as-path-prepend "65501 65501 65501 65501 65501 65501";
}
term 9 {
    from community prepend-9;
    then as-path-prepend "65501 65501 65501 65501 65501 65501 65501 65501 65501";
}
```

Question: Do you see any problems with the export policies?

_____

_____

Answer: Yes, the export policy *no-tagged-specifics* is using the wrong match criteria. You should use **prefix-length-range** instead.

Step 2.22

Enter into configuration mode and navigate to the [edit policy-options policy-statement *no-tagged-specifics*] hierarchy and fix the term by removing the up-to conditions and replace it with prefix-length-range conditions. Commit and exit to operational mode when you are satisfied with your changes.

```
lab@mxB:R3> configure
Entering configuration mode

[edit]
lab@mxB:R3# edit policy-options policy-statement no-tagged-specifics

[edit policy-options policy-statement no-tagged-specifics]
lab@mxB:R3# show
term no-specifics {
    from {
        community c1;
        route-filter 0.0.0.0/24 upto /32;
    }
    then reject;
}

[edit policy-options policy-statement no-tagged-specifics]
lab@mxB:R3# delete term no-specifics from route-filter 0.0.0.0/24 upto /32

[edit policy-options policy-statement no-tagged-specifics]
lab@mxB:R3# set term no-specifics from route-filter 0/0 prefix-length-range /25-/
   32
```

```
[edit policy-options policy-statement no-tagged-specifics]
lab@mxB:R3# show
term no-specifics {
    from {
        community c1;
        route-filter 0.0.0.0/0 prefix-length-range /25-/32;
    }
    then reject;
}

[edit policy-options policy-statement no-tagged-specifics]
lab@mxB:R3# commit and-quit
commit complete
Exiting configuration mode

lab@mxB:R3>
```

Step 2.23

Use the **`show route advertising-protocol bgp 172.22.131.46`** command to verify that you are not exporting anything to your EBGP that is longer than the required /24.

```
lab@mxB:R3> show route advertising-protocol bgp 172.22.131.46

inet.0: 37 destinations, 55 routes (34 active, 0 holddown, 9 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.20.0.0/24           Self                                    65501 65501 65501
   [65501] 65530 I
* 172.20.1.0/24           Self                                    65501 65501 65501
   [65501] 65530 I
* 172.20.2.0/24            Self                                     65530 I
* 172.20.3.0/24            Self                                     65530 I
* 172.20.4.0/24            Self                                     65530 1 2 2 I
* 172.20.5.0/24            Self                                    65530 1 2 3 3 4 I
* 172.20.6.0/24           Self                                    65530 1 2 3 3 4 5 I
* 172.20.7.0/24            Self                                     65530 1 2 2 3 I
```

Question: Are you still advertising the routes with a /25 subnet mask?

_____

Answer: No, you should not be advertising any routes that are longer than /24. The R3 router's routes and policies seem to be working correctly now.

Step 2.24

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 R5 logical system by issuing the **`set cli logical-system R5`** command.

```
lab@mxC> set cli logical-system R5
Logical system: R5

lab@mxC:R5>
```

Step 2.25

Use the **show bgp summary** command to view the BGP neighborships and routes being learned by the R5 router.

```
lab@mxC:R5> 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
                      45         22          0          0          0          0
inet6.0
                       5          2          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.22.131.42         65532         52         49       0       0       18:00 Establ
  inet.0: 8/9/9/0
  inet6.0: 1/1/1/0
192.168.71.2          65501         64         49       0       0       16:51 Establ
  inet.0: 14/18/18/0
  inet6.0: 1/2/2/0
192.168.71.4          65501         51         50       0       0       16:58 Establ
  inet.0: 0/18/18/0
  inet6.0: 0/2/2/0
```

Question: How many prefixes is the R5 router receiving from the Endpoint2 router?

_____

_____

Answer: You should see 9 IPv4 routes and 1 IPv6 route being learned from the Endpoint2 router.

Question: How many of the routes learned from the Endpoint2 router are active in the routing table?

_____

_____

Answer: According to the output, 8 of the 9 BGP IPv4 routes learned from the Endpoint2 router are active in the routing table. Also, the only received IPv6 route from the Endpoint2 router is active in the routing table.

Step 2.26

Use the **show route receive-protocol bgp 172.22.131.42** command to review the routes learned from the Endpoint2 router to ensure that the R5 router is not accepting routes that have transited AS number 3.

```
lab@mxC:R5> show route receive-protocol bgp 172.22.131.42

inet.0: 35 destinations, 58 routes (35 active, 0 holddown, 0 hidden)
```

```
  Prefix                  Nexthop              MED     Lclpref    AS path
  172.21.0.0/24           172.22.131.42                           65532 5 5 3 I
* 172.21.1.0/24           172.22.131.42                           65532 I
* 172.21.2.0/24           172.22.131.42                           65532 I
* 172.21.3.0/24           172.22.131.42                           65532 I
* 172.21.4.0/24           172.22.131.42                           65532 6 7 7 I
* 172.21.5.0/24           172.22.131.42                           65532 6 7 31 8 I
* 172.21.6.0/24           172.22.131.42                           65532 6 7 32 9 I
* 172.21.7.0/24           172.22.131.42                          65532 6 7 31 32 I
* 172.22.0.0/24          172.22.131.42                         65532 69 86 73 76 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 20 destinations, 23 routes (20 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* fd69:8:2::/48           ::ffff:172.22.131.42                    65532 I
```

Question: Do you see routes that have transited AS number 3?

_____

Answer: Yes, there is a route that transited AS number 3. The policy must not be rejecting the correct routes.

Step 2.27

Review the BGP configuration and import policies to determine the reason why you are still accepting routes that have traversed the AS 3.

```
lab@mxC:R5> show configuration protocols bgp group ebgp-AS65532
type external;
import [ tag-routes reject-3 ];
family inet {
    unicast;
}
family inet6 {
    unicast;
}
export prepend;
peer-as 65532;
neighbor 172.22.131.42;

lab@mxC:R5> show configuration policy-options policy-statement tag-routes
then {
    community add c2;
    accept;
}

lab@mxC:R5> show configuration policy-options policy-statement reject-3
term reject-3 {
    from as-path as-3;
    then reject;
}

lab@mxC:R5> show configuration policy-options community c2
members 65501:200;
```

```
lab@mxC:R5> show configuration policy-options as-path as-3
".* 3";
```

Question: Do you see any problems with the current *reject-3* policy or as-path definition?

_____

Answer: The policy seems okay and the as-path definition should work fine as it is.

Question: Do you see any problem with the current *tag-routes* policy?

_____

Answer: The policy will accept all routes that are evaluated by this policy.

Question: How do you fix this issue?

_____

_____

Answer: You have a couple options to fix this issue. You can re-order the policies so that the routes that have the AS 3 in their AS path are rejected before being evaluated and accepted by the *tag-route* policy. Alternately, you can remove the accept statement from the *tag-route* policy.

Step 2.28

Enter into configuration mode and navigate to the [edit policy-options policy-statement *tag-routes*] hierarchy and remove the accept action from the policy. Commit and exit to operational mode when you are satisfied with your changes.

```
lab@mxC:R5> configure
Entering configuration mode

[edit]
lab@mxC:R5# edit policy-options policy-statement tag-routes

[edit policy-options policy-statement tag-routes]
lab@mxC:R5# show
then {
    community add c2;
    accept;
}

[edit policy-options policy-statement tag-routes]
lab@mxA-1:R5# delete then accept

[edit policy-options policy-statement tag-routes]
lab@mxA-1:R5# commit and-quit
```

```
commit complete
Exiting configuration mode

lab@mxA-1:R5>
```

Step 2.29

Use the **show route receive-protocol bgp 172.22.131.42** command to review the routes learned from the Endpoint2 router to ensure that the R5 router is not accepting routes that have transited AS number 3.

```
lab@mxC:R5> show route receive-protocol bgp 172.22.131.42

inet.0: 35 destinations, 58 routes (35 active, 0 holddown, 1 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 172.21.1.0/24           172.22.131.42                           65532 I
* 172.21.2.0/24           172.22.131.42                           65532 I
* 172.21.3.0/24           172.22.131.42                           65532 I
* 172.21.4.0/24           172.22.131.42                           65532 6 7 7 I
* 172.21.5.0/24           172.22.131.42                           65532 6 7 31 8 I
* 172.21.6.0/24           172.22.131.42                           65532 6 7 32 9 I
* 172.21.7.0/24           172.22.131.42                           65532 6 7 31 32 I
* 172.22.0.0/24           172.22.131.42                           65532 69 86 73 76 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

inet6.0: 20 destinations, 23 routes (20 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* fd69:8:2::/48           ::ffff:172.22.131.42                    65532 I
```

Question: Do you see any routes that have transited through AS 3?

_____

_____

Answer: No, you should not see any routes that have transited through AS 3.

Step 2.30

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

```
lab@mxC:R5> exit


mxC (ttyu0)

login:
```

Tell your instructor that you have completed this lab.

Network Management Diagram

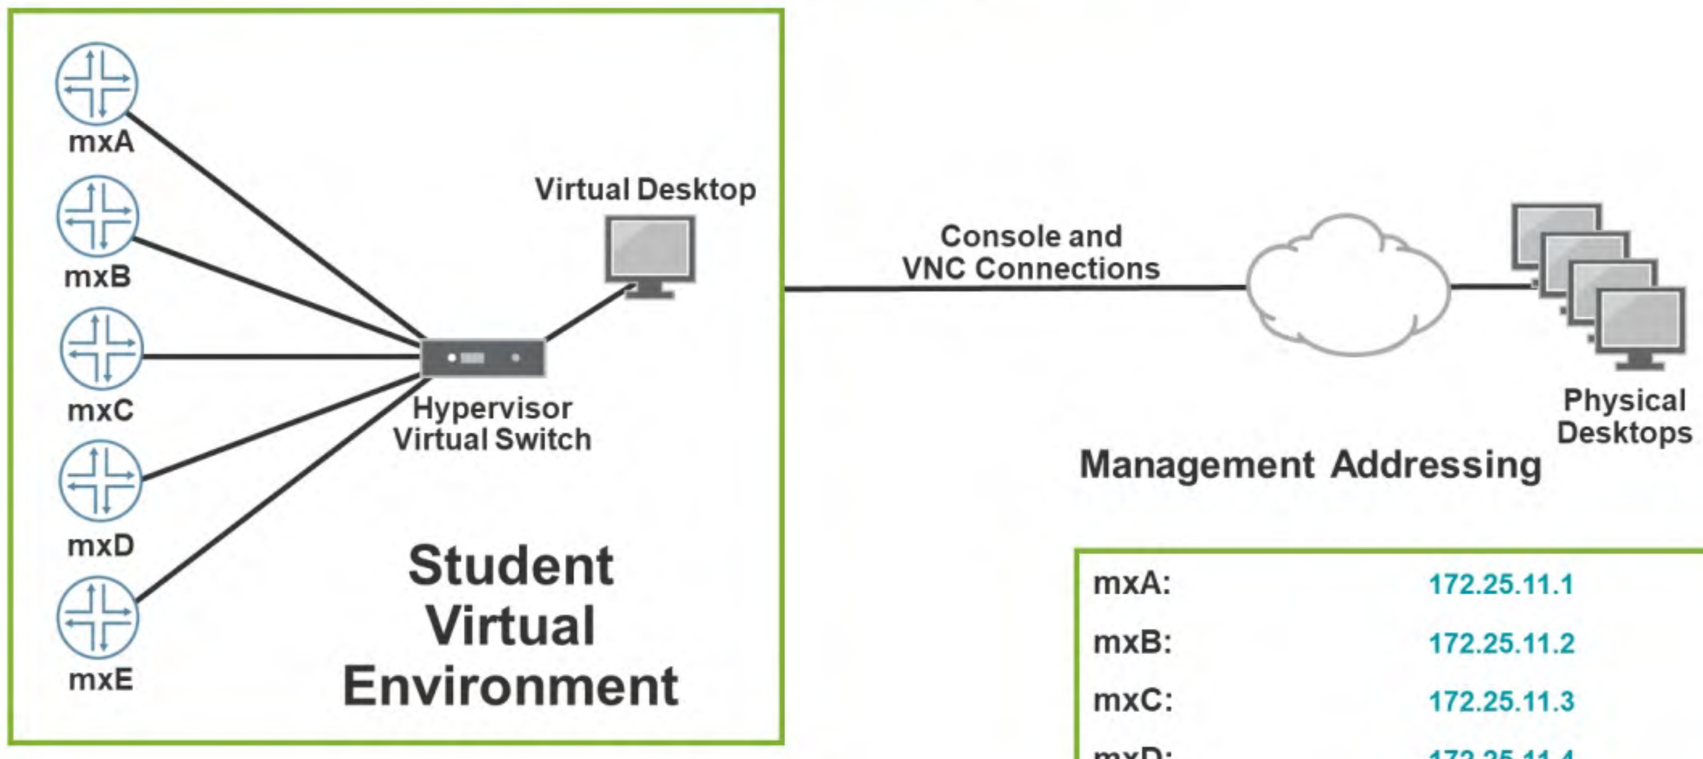

Management Addressing

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: Troubleshooting Routing Policy

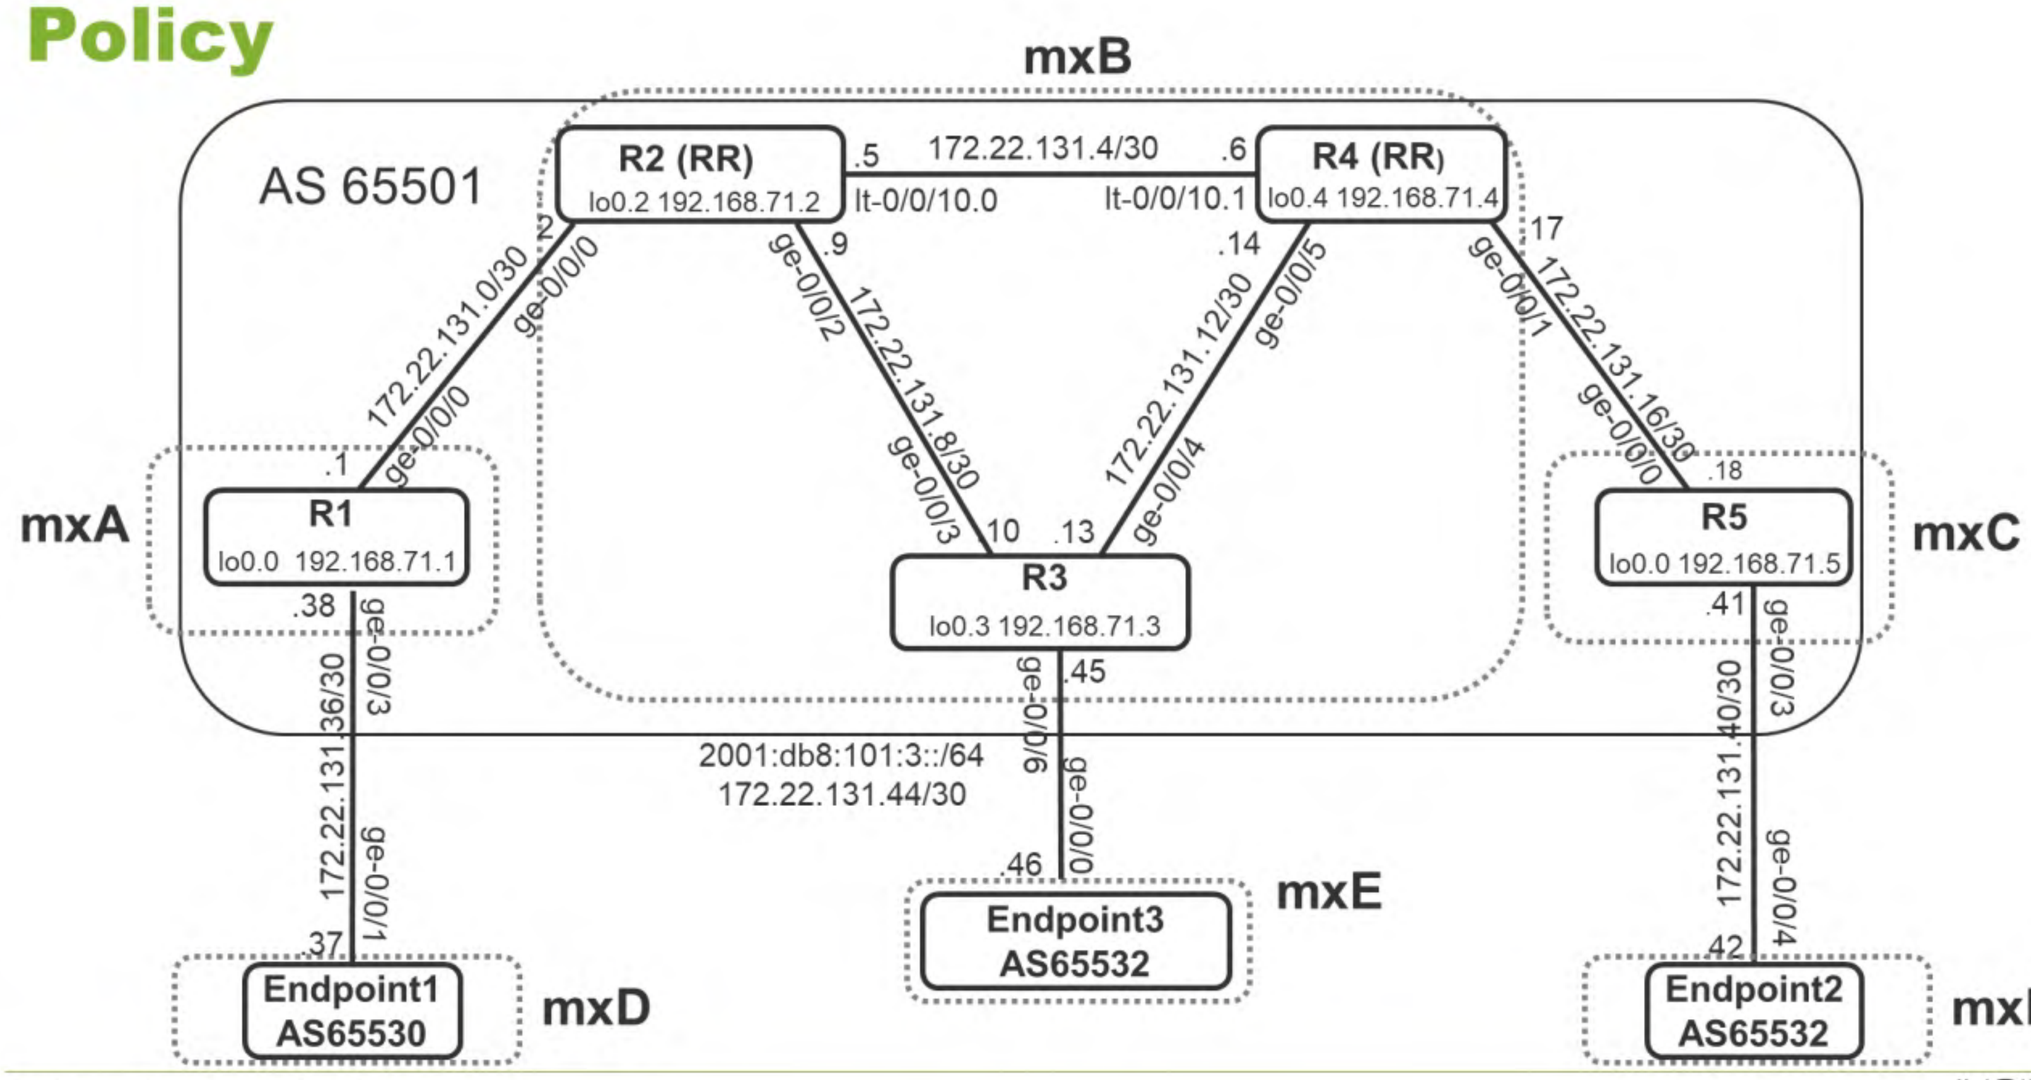

© 2020 Juniper Networks, Inc. All Rights Reserved.

Lab

BGP Route Damping

Overview

In this lab, you use the lab diagram titled “Lab Network Diagram: BGP Route Damping” to monitor the EBGP-received routes for any link flapping that might occur within the network. Route damping monitors the behavior of EBGP-received routes being withdrawn and readvertised. It uses a point system known as figure-of-merit to determine whether routes should be installed and advertised into the IBGP topology, or suppressed at the edge.

By completing this lab, you perform the following tasks:

- Create a static route.
- Modify the export policies to advertise the static route.
- Configure damping.
- Flap the advertised static route.
- Configure policy to alter the default damping parameters.
- Apply the policy as an EBGP import policy.

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/lab16-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 27 17:20:28 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/lab16-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: 2019-12-06 15:11:54 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;
            }
        }
        bgp {
            group ibgp {
                type internal;
                local-address 172.16.1.2;
                export [ redistribute-statics next-hop-self ];
                neighbor 172.16.1.1;
            }
            group P3 {
                type external;
                multihop;
                local-address 172.16.1.2;
                export export-aggregate;
                peer-as 65020;
                neighbor 172.31.102.1;
            }
        }
    }
    policy-options {
        policy-statement export-aggregate {
            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;
            }
        }
        policy-statement next-hop-self {
            term 1 {
                from {
                    protocol bgp;
                    route-type external;
                }
                then {
                    next-hop self;
                }
            }
        }
        policy-statement redistribute-statics {
            term 1 {
                from protocol static;
                then accept;
            }
        }
    }
    routing-options {
```

```
        static {
            route 172.16.1.128/26 reject;
            route 172.16.1.192/26 reject;
            route 172.31.102.1/32 {
                next-hop 172.22.125.2;
                no-readvertise;
            }
        }
        autonomous-system 65001;
        aggregate {
            route 172.16.1.0/24;
        }
    }
}
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 next-hop-self ];
                neighbor 172.16.2.1;
            }
            group P3 {
                type external;
                multihop;
```

```
                    local-address 172.16.2.2;
                    export export-aggregate;
                    neighbor 172.31.102.1 {
                        peer-as 65020;
                    }
                }
            }
        }
        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 export-p3 {
                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 next-hop-self {
                term 1 {
                    from {
                        protocol bgp;
                        route-type external;
                    }
                    then {
                        next-hop self;
                    }
                }
            }
            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: Based on the configuration, which routers (logical systems) will mxB represent in this lab?

_____

_____

Answer: The mxB device will, using logical systems, represent the P3, R3-1, and R3-2 routers.

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 any 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 configured for an IBGP session with the mxA router and an EBGP session with the P3 router.

Question: What routes will the R3-1 router be advertising to its EBGP peer?

_____

_____

Answer: The R3-1 router is configured with an export policy called *export-aggregate* that is applied to its EBGP peer. This policy causes the R3-1 router to advertise the configured aggregate route of 172.16.1.0/24.

Question: Which routers are R3-2 configured to peer with using BGP?

_____

_____

Answer: The R3-2 router is configured for an EBGP session with the P3 router and an IBGP session with the mxC router.

Question: What routes will the R3-2 router be advertising to its EBGP peer?

_____

_____

Answer: The R3-2 router is configured with an export policy called *export-aggregate* that is applied to its EBGP peer. This policy causes the R3-2 router to advertise the configured aggregate route of 172.16.2.0/24.

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
Threading mode: BGP I/O
Groups: 2 Peers: 2 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          5          5       0       0          45 Establ
  inet.0: 1/1/1/0
172.16.2.2            65002          5          5       0       0          45 Establ
  inet.0: 1/1/1/0
```

Question: What is the state of the P3 router's BGP neighbor relationships?

_____

_____

Answer: The P3 router should have an established EBGP session with both the R3-1 and R3-2 routers.

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/lab16-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 27 17:20:28 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/lab16-start.config
load complete

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

lab@mxC>
```

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: 2019-12-08 22:07:24 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 next-hop-self {
        term 1 {
            from {
```

```
                protocol bgp;
                route-type external;
            }
            then {
                next-hop self;
            }
        }
    }
    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;
                route-filter 172.16.0.0/16 orlonger;
            }
            then accept;
        }
    }
    as-path partner-as ".* 65001";
    as-path internal-as "()";
}
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 next-hop-self ];
            neighbor 172.16.2.2;
        }
        group P1-P2 {
            type external;
            export export-aggregate;
            peer-as 65412;
            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: What routes will the mxC router be advertising to its EBGP peers?

_____

_____

Answer: The mxC router is configured with an export policy called *export-aggregate* that is applied to its EBGP peers. This policy causes the mxC router to advertise the configured aggregate route of 172.16.2.0/24.

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: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      17         11          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        1:09 Establ
  inet.0: 7/7/7/0
172.22.122.2          65412          0          0       0       0       1:20 Connect
172.22.124.2          65412          0          0       0       0       1:20 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. Because of the state of the lab, you may see an established session with both the P1 and P2 routers. If you do, there is no problem.

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/lab16-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: Wed Nov 27 17:20:28 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/lab16-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 22:15:06 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
                       6          6          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        1:50 Connect
172.22.124.1          65002         17         16       0       0        5:30 Establ
  inet.0: 6/6/6/0
172.31.100.1          65412          0          0       0       0        1:50 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. Depending on the topology state, you might see the BGP sessions with the P1 and mxA router in the established state. If so, this is not a problem.

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/lab16-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 27 17:20:28 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/lab16-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: 2019-12-27 22:14:55 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
                      12          6          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        1:20 Connect
172.22.122.1          65002         19         18       0       0        6:44 Establ
  inet.0: 6/6/6/0
172.31.101.1          65412         18         18       0       0        5:58 Establ
  inet.0: 0/6/6/0
```

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. Depending on the state of the lab, you may see that the BGP session with the mxA router is in the established state. If so, this is not a problem.

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/lab16-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: Wed Nov 27 17:20:28 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/lab16-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, OSPF, and BGP configuration.

```
lab@mxA> show configuration
## Last commit: 2019-12-08 22:07:17 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;
            }
        }
    }
}
policy-options {
    policy-statement export-aggregate {
        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;
        }
    }
    policy-statement next-hop-self {
        term 1 {
            from {
                protocol bgp;
                route-type external;
            }
            then {
                next-hop self;
            }
        }
    }
    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 172.16.1.0/26 reject;
        route 172.16.1.64/26 reject;
    }
    forwarding-table {
        export pfe-load-balance;
    }
    autonomous-system 65001;
    aggregate {
        route 172.16.1.0/24;
    }
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface lo0.0;
            interface ge-0/0/0.0;
        }
    }
    bgp {
        group ibgp {
            type internal;
            local-address 172.16.1.1;
            export [ redistribute-statics next-hop-self ];
            neighbor 172.16.1.2;
        }
        group P1-P2 {
            type external;
            export export-aggregate;
            peer-as 65412;
            neighbor 172.22.121.2;
            neighbor 172.22.123.2;
        }
    }
}
```

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: Which routers are mxA configured to peer with using BGP?

_____

_____

Answer: The mxA router is configured with an IBGP session with the R3-1 router and an EBGP session with the P1 and P2 routers.

Question: What routes will the mxA router be advertising to its EBGP peers?

_____

_____

Answer: The mxA router is configured with an export policy called *export-aggregate* that is applied to its EBGP peers. This policy causes the mxA router to advertise the configured aggregate route of 172.16.1.0/24.

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
Threading mode: BGP I/O
Groups: 2 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0
                      17         11          0          0          0          0
Peer                     AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
   State|#Active/Received/Accepted/Damped...
172.16.1.2            65001         27         25       0       0        9:36 Establ
  inet.0: 6/7/7/0
172.22.121.2          65412         28         26       0       0       10:24 Establ
  inet.0: 5/5/5/0
172.22.123.2          65412         27         26       0       0       10:18 Establ
  inet.0: 0/5/5/0
```

Question: What is the status of the mxA router's BGP neighbor relationships?

_____

Answer: The mxA router has established BGP sessions with the R3-1, P1, and P2 routers.

Part 2: Configuring BGP Damping

In this lab part, you create and advertise a static route from the mxC router to the P1 and P2 routers which will propagate the route through EBGP to the mxA router. After damping is enabled in the mxA router, you will *flap* the route by deleting the static route and adding it back on the mxC router.

Step 2.1

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 routing-options]` hierarchy. Configure the 172.22.2.0/24 static route with a next hop of `reject`.

```
lab@mxC> configure
Entering configuration mode

[edit]
lab@mxC# edit routing-options

[edit routing-options]
lab@mxC# set static route 172.22.2.0/24 reject

[edit routing-options]
lab@mxC#
```

Step 2.2

Navigate to the `[edit policy-options policy-statement export-aggregate]` hierarchy and configure a third term in the policy to advertise the new static route. Use the **show** command to display the policy. Commit your configuration when finished.

```
[edit routing-options]
lab@mxC# top edit policy-options policy-statement export-aggregate

[edit policy-options policy-statement export-aggregate]
lab@mxC# set term 3 from protocol static

[edit policy-options policy-statement export-aggregate]
lab@mxC# set term 3 from route-filter 172.22.2.0/24 exact

[edit policy-options policy-statement export-aggregate]
lab@mxC# set term 3 then accept

[edit policy-options policy-statement export-aggregate]
lab@mxC# show
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;
}
term 3 {
    from {
        protocol static;
        route-filter 172.22.2.0/24 exact;
    }
    then accept;
}

[edit policy-options policy-statement export-aggregate]
lab@mxC# commit
commit complete

[edit policy-options policy-statement export-aggregate]
lab@mxC#
```

Step 2.3

Issue the **`run show route advertising-protocol bgp 172.22.122.2`** and **`run show route advertising-protocol bgp 172.22.122.2`** commands to view the routes that are being advertised from the mxC router to the P1 and P2 routers.

```
[edit policy-options policy-statement export-aggregate]
lab@mxC# run show route advertising-protocol bgp 172.22.122.2

inet.0: 26 destinations, 32 routes (26 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 40.40.0.0/24            Self                                    65020 I
* 40.40.1.0/24            Self                                    65020 I
* 40.40.2.0/24            Self                                    65020 I
* 40.40.3.0/24            Self                                    65020 I
* 172.16.2.0/24           Self                                    I
* 172.22.2.0/24           Self                                    I

[edit policy-options policy-statement export-aggregate]
lab@mxC# run show route advertising-protocol bgp 172.22.124.2

inet.0: 26 destinations, 32 routes (26 active, 0 holddown, 0 hidden)
  Prefix                  Nexthop              MED     Lclpref    AS path
* 40.40.0.0/24            Self                                    65020 I
* 40.40.1.0/24            Self                                    65020 I
* 40.40.2.0/24            Self                                    65020 I
* 40.40.3.0/24            Self                                    65020 I
* 172.16.2.0/24           Self                                    I
* 172.22.2.0/24           Self                                    I
```

Question: Is the 172.22.2.0/24 route being advertised to P1 and P2?

_____

Answer: Yes, the 172.22.2.0/24 route is being advertised to the P1 and P2 routers.

Step 2.4

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

From the SSH session with the mxA router, issue the **`show route 172.22.2.0/24`** command to display the new route being advertised by the mxC router.

```
lab@mxA> show route 172.22.2.0/24

inet.0: 26 destinations, 33 routes (26 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.22.2.0/24       *[BGP/170] 00:02:51, localpref 100
                       AS path: 65412 65002 I, validation-state: unverified
                     > to 172.22.121.2 via ge-0/0/3.0
                     [BGP/170] 00:02:51, localpref 100
                       AS path: 65412 65002 I, validation-state: unverified
                     > to 172.22.123.2 via ge-0/0/1.0
```

Question: Is the 172.22.2.0/24 route an active BGP route in the mxA router's routing table?

_____

Answer: Yes, the route should be an active BGP route in the mxA router's routing table. If the route is not active, check your configuration or ask your instructor for help.

Step 2.5

Enter configuration mode and navigate to the `[edit protocols bgp]` hierarchy. Enable BGP `damping` as a global command. Commit the change.

```
lab@mxA> configure
Entering configuration mode

[edit]
lab@mxA# edit protocols bgp

[edit protocols bgp]
lab@mxA# set damping

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

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

Step 2.6

You will now cause the 172.22.2.0/24 route to flap on the mxC router.

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 routing-options]` hierarchy. Cause the newly-configured static route to flap by deleting the static route. Commit the change when you are ready.

```
[edit policy-options policy-statement export-aggregate]
lab@mxC# top edit routing-options

[edit routing-options]
lab@mxC# delete static route 172.22.2.0/24

[edit routing-options]
lab@mxC# commit
commit complete

[edit routing-options]
lab@mxC#
```

Step 2.7

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 damping history detail`** command to display routes that are withdrawn but have a history of figure-of-merit in the mxA router's routing table.

```
[edit protocols bgp]
lab@mxA# run show route damping history detail

inet.0: 26 destinations, 33 routes (25 active, 0 holddown, 2 hidden)
172.22.2.0/24 (2 entries, 0 announced)
         BGP                    /-101
                Next hop type: Router, Next hop index: 591
                Address: 0xb39d510
                Next-hop reference count: 6
                Source: 172.22.123.2
                Next hop: 172.22.123.2 via ge-0/0/1.0, selected
                Session Id: 0x16b
                State: <Hidden Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 56
                Validation State: unverified
                Task: BGP_65412.172.22.123.2
                AS path: 65412 65002 I
                Accepted
                Localpref: 100
                Router ID: 172.31.101.1
                Merit (last update/now): 1000/962
                Default damping parameters used
                Last update: 00000000:00:56 First update: 00000000:00:56
                Flaps: 1
                History entry.  Expires in: 00000000:34:00
                Hidden reason: suppressed by damping
         BGP                    /-101
                Next hop type: Router, Next hop index: 590
                Address: 0xb39d630
                Next-hop reference count: 11
                Source: 172.22.121.2
                Next hop: 172.22.121.2 via ge-0/0/3.0, selected
                Session Id: 0x16a
                State: <Hidden Ext>
                Local AS: 65001 Peer AS: 65412
```

```
                    Age: 56
                    Validation State: unverified
                    Task: BGP_65412.172.22.121.2
                    AS path: 65412 65002 I
                    Accepted
                    Localpref: 100
                    Router ID: 172.31.100.1
                    Merit (last update/now): 1000/962
                    Default damping parameters used
                    Last update: 00000000:00:56 First update: 00000000:00:56
                    Flaps: 1
                    History entry.  Expires in: 00000000:34:00
                    Hidden reason: suppressed by damping

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Do any withdrawn routes have a history of figure-of-merit?

_____

Answer: A withdrawn route should have figure-of-merit. If no routes have been withdrawn, check your configuration on the mxC router or ask your instructor for help.

Question: What is the current figure-of-merit for this withdrawn route?

_____

_____

Answer: In this example, the current figure-of-merit is 962. This value changes based on the half-life parameter and the amount of time that has elapsed.

Step 2.8

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

From the SSH session with the mxC router, navigate to the top of the configuration and perform a **rollback 1** to readvertise the static route. Commit the configuration when completed.

```
[edit routing-options]
lab@mxC# top

[edit]
lab@mxC# rollback 1
load complete

[edit]
lab@mxC# commit
commit complete

[edit]
lab@mxC#
```

Step 2.9

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 damping decayed detail`** command to show active routes that are decaying but not suppressed.

```
[edit protocols bgp]
lab@mxA# run show route damping decayed detail

inet.0: 26 destinations, 33 routes (26 active, 0 holddown, 0 hidden)
172.22.2.0/24 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 591
                Address: 0xb39d510
                Next-hop reference count: 7
                Source: 172.22.123.2
                Next hop: 172.22.123.2 via ge-0/0/1.0, selected
                Session Id: 0x16b
                State: <Active Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 47
                Validation State: unverified
                Task: BGP_65412.172.22.123.2
                Announcement bits (3): 0-KRT 6-BGP_RT_Background 7-Resolve tree 1
                AS path: 65412 65002 I
                Accepted
                Localpref: 100
                Router ID: 172.31.101.1
                Merit (last update/now): 1688/1636
                Default damping parameters used
                Last update: 00000000:00:47 First update: 00000000:04:31
                Flaps: 2
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 590
                Address: 0xb39d630
                Next-hop reference count: 11
                Source: 172.22.121.2
                Next hop: 172.22.121.2 via ge-0/0/3.0, selected
                Session Id: 0x16a
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65001 Peer AS: 65412
                Age: 47
                Validation State: unverified
                Task: BGP_65412.172.22.121.2
                AS path: 65412 65002 I
                Accepted
                Localpref: 100
                Router ID: 172.31.100.1
                Merit (last update/now): 1688/1636
                Default damping parameters used
                Last update: 00000000:00:47 First update: 00000000:04:31
                Flaps: 2

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Are there routes decaying but not suppressed?

_____

Answer: The route should be an active route in the default routing table with decaying figure-of-merit.

Question: How many flaps have occurred on the active route?

_____

_____

Answer: Only 2 flaps have occurred for this route: a withdrawn route and readvertised route.

Step 2.10

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

From the SSH session with the mxC router, flap the route 4 times using the **rollback 1** and **commit** commands.

```
[edit]
lab@mxC# rollback 1
load complete

[edit]
lab@mxC# commit
commit complete

[edit]
lab@mxC# rollback 1
load complete

[edit]
lab@mxC# commit
commit complete

[edit]
lab@mxC# rollback 1
load complete

[edit]
lab@mxC# commit
commit complete

[edit]
lab@mxC# rollback 1
load complete

[edit]
lab@mxC# commit
commit complete

[edit]
lab@mxC#
```

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 route damping suppressed detail`** command to display routes that have been suppressed due to damping in the default routing table.

```
[edit protocols bgp]
lab@mxA# run show route damping suppressed detail

inet.0: 26 destinations, 33 routes (25 active, 0 holddown, 2 hidden)
172.22.2.0/24 (2 entries, 0 announced)
         BGP                    /-101
                Next hop type: Router, Next hop index: 590
                Address: 0xb39d630
                Next-hop reference count: 11
                Source: 172.22.121.2
                Next hop: 172.22.121.2 via ge-0/0/3.0, selected
                Session Id: 0x16a
                State: <Hidden Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 56
                Validation State: unverified
                Task: BGP_65412.172.22.121.2
                AS path: 65412 65002 I
                Localpref: 100
                Router ID: 172.31.100.1
                Merit (last update/now): 5394/5190
                Default damping parameters used
                Last update: 00000000:00:56 First update: 00000000:07:36
                Flaps: 6
                Suppressed. Reusable in: 00000000:42:00
                Preference will be: 170
                Hidden reason: suppressed by damping
         BGP                    /-101
                Next hop type: Router, Next hop index: 591
                Address: 0xb39d510
                Next-hop reference count: 6
                Source: 172.22.123.2
                Next hop: 172.22.123.2 via ge-0/0/1.0, selected
                Session Id: 0x16b
                State: <Hidden Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 56
                Validation State: unverified
                Task: BGP_65412.172.22.123.2
                AS path: 65412 65002 I
                Localpref: 100
                Router ID: 172.31.101.1
                Merit (last update/now): 5394/5190
                Default damping parameters used
                Last update: 00000000:00:56 First update: 00000000:07:36
                Flaps: 6
                Suppressed. Reusable in: 00000000:42:00
                Preference will be: 170
```

```
                Hidden reason: suppressed by damping

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Have any routes in the default routing table been suppressed because of damping?

_____

Answer: After you have completed the 4 flaps in the previous step, the static route should be suppressed in the mxA router's routing table.

Question: If no other flaps occur, what is the estimated reuse time for this route?

_____

Answer: In this example, the estimated reuse time is 00:42:00 minutes, which is calculated based on the amount of figure-of-merit and the half-life time.

Part 3: Modifying the BGP Damping Parameters

In this lab part, you use routing policy to modify the default damping parameters. You create an import policy to disable damping on routes received from the P1 router, and another policy to aggressively damp routes received from the P2 router.

Step 3.1

Navigate to the `[edit policy-options]` hierarchy. Create and configure two damping profiles named *`disable`* and *`aggressive`*. The *`disable`* profile should use the *`disable`* option. The *`aggressive`* profile should set the **`suppress`** parameter to *`1500`*, `half-life` to *`10`* minutes, and the `reuse` parameter to *`500`*.

```
[edit protocols bgp]
lab@mxA# top edit policy-options

[edit policy-options]
lab@mxA# set damping disable disable

[edit policy-options]
lab@mxA# set damping aggressive suppress 1500

[edit policy-options]
lab@mxA# set damping aggressive half-life 10

[edit policy-options]
lab@mxA# set damping aggressive reuse 500

[edit policy-options]
lab@mxA#
```

Step 3.2

Navigate to the `[edit policy-options policy-statement modify-damping]` hierarchy. Configure the first term of the *modify-damping* import policy to use the *disable* profile on all BGP routes received from the P1 neighbor. Configure the second term to use the *aggressive* profile on all routes received from the P2 neighbor. Use the **show** command to display the policy.

```
[edit policy-options]
lab@mxA# edit policy-statement modify-damping

[edit policy-options policy-statement modify-damping]
lab@mxA# set term 1 from protocol bgp

[edit policy-options policy-statement modify-damping]
lab@mxA# set term 1 from neighbor 172.22.121.2

[edit policy-options policy-statement modify-damping]
lab@mxA# set term 1 then damping disable

[edit policy-options policy-statement modify-damping]
lab@mxA# set term 2 from protocol bgp

[edit policy-options policy-statement modify-damping]
lab@mxA# set term 2 from neighbor 172.22.123.2

[edit policy-options policy-statement modify-damping]
lab@mxA# set term 2 then damping aggressive

[edit policy-options policy-statement modify-damping]
lab@mxA# show
term 1 {
    from {
        protocol bgp;
        neighbor 172.22.121.2;
    }
    then damping disable;
}
term 2 {
    from {
        protocol bgp;
        neighbor 172.22.123.2;
    }
    then damping aggressive;
}

[edit policy-options policy-statement modify-damping]
lab@mxA#
```

Step 3.3

Navigate to the `[edit protocols bgp group P1-P2]` hierarchy and apply the *modify-damping* policy as a group import policy. Commit the changes when completed.

```
[edit policy-options policy-statement modify-damping]
lab@mxA# top edit protocols bgp group P1-P2
```

```
[edit protocols bgp group P1-P2]
lab@mxA# set import modify-damping

[edit protocols bgp group P1-P2]
lab@mxA# commit
commit complete

[edit protocols bgp group P1-P2]
lab@mxA#
```

Step 3.4

Use the **run clear bgp damping** command to reset the figure-of-merit to zero on all routes. Use the **run show route damping suppressed** command to verify that all routes are active and no routes are suppressed in the default routing table.

```
[edit protocols bgp group P1-P2]
lab@mxA# run clear bgp damping

[edit protocols bgp group P1-P2]
lab@mxA# run show route damping suppressed

inet.0: 26 destinations, 33 routes (26 active, 0 holddown, 0 hidden)

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

[edit protocols bgp group P1-P2]
lab@mxA#
```

Question: Are suppressed routes in the routing table?

_____

_____

Answer: No suppressed routes are in the routing table.

Step 3.5

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

From the SSH session with the mxC router, flap the route 2 times using the **rollback 1** and **commit** commands.

```
[edit]
lab@mxC# rollback 1
load complete

[edit]
lab@mxC# commit
commit complete

[edit]
lab@mxC# rollback 1
load complete

[edit]
lab@mxC# commit
```

```
commit complete

[edit]
lab@mxC#
```

Step 3.6

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 damping suppressed detail`** command to display any routes suppressed due to damping in the mxA router's routing table.

```
[edit protocols bgp group P1-P2]
lab@mxA# run show route damping suppressed detail

inet.0: 26 destinations, 33 routes (26 active, 0 holddown, 1 hidden)
172.22.2.0/24 (2 entries, 1 announced)
         BGP                    /-101
                Next hop type: Router, Next hop index: 591
                Address: 0xb39d510
                Next-hop reference count: 6
                Source: 172.22.123.2
                Next hop: 172.22.123.2 via ge-0/0/1.0, selected
                Session Id: 0x16b
                State: <Hidden Ext>
                Inactive reason: Unusable path
                Local AS: 65001 Peer AS: 65412
                Age: 1:07
                Validation State: unverified
                Task: BGP_65412.172.22.123.2
                AS path: 65412 65002 I
                Localpref: 100
                Router ID: 172.31.101.1
                Merit (last update/now): 2000/1865
                damping-parameters: aggressive
                Last update: 00000000:01:07 First update: 00000000:01:11
                Flaps: 2
                Suppressed. Reusable in: 00000000:19:00
                Preference will be: 170
                Hidden reason: suppressed by damping

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Are any suppressed routes in the routing table?

_____

Answer: Yes. One suppressed route is in the routing table.

Question: In the suppressed route, what is the damping parameter being used to suppress this route?

_____

_____

Answer: The damping parameter being used to suppress this route is *aggressive*. This parameter is configured in the *modify-damping* policy for the P2 router.

Step 3.7

Using the data in the following table, issue the **run show route 172.22.2.0/24 detail** command to verify that the damping policy on the P1 router is disabling damping on received routes.

```
[edit protocols bgp group P1-P2]
lab@mxA# run show route 172.22.2.0/24 detail

inet.0: 26 destinations, 33 routes (26 active, 0 holddown, 1 hidden)
172.22.2.0/24 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 590
                Address: 0xb39d630
                Next-hop reference count: 12
                Source: 172.22.121.2
                Next hop: 172.22.121.2 via ge-0/0/3.0, selected
                Session Id: 0x16a
                State: <Active Ext>
                Local AS: 65001 Peer AS: 65412
                Age: 3:16
                Validation State: unverified
                Task: BGP_65412.172.22.121.2
                Announcement bits (3): 0-KRT 6-BGP_RT_Background 7-Resolve tree 1
                AS path: 65412 65002 I
                Accepted
                Localpref: 100
                Router ID: 172.31.100.1
```

Question: Which EBGP peer is advertising the active route in the default routing table?

_____

_____

Answer: In this example, the source of the active route is 172.22.121.2, which is the BGP peer address of the P1 router.

Step 3.8

Use the **run show route damping decayed detail** command to display any active routes with figure-of-merit in the default routing table.

```
[edit protocols bgp group P1-P2]
lab@mxA# run show route damping decayed detail
```

```
inet.0: 26 destinations, 33 routes (26 active, 0 holddown, 1 hidden)

inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Question: Is the EBGP import policy that disables damping on the P1 router working?

_____

_____

Answer: Because the route is active and no figure-of-merit is present, it appears that the import policy applied to the P1 router is disabling damping.

Step 3.9

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.

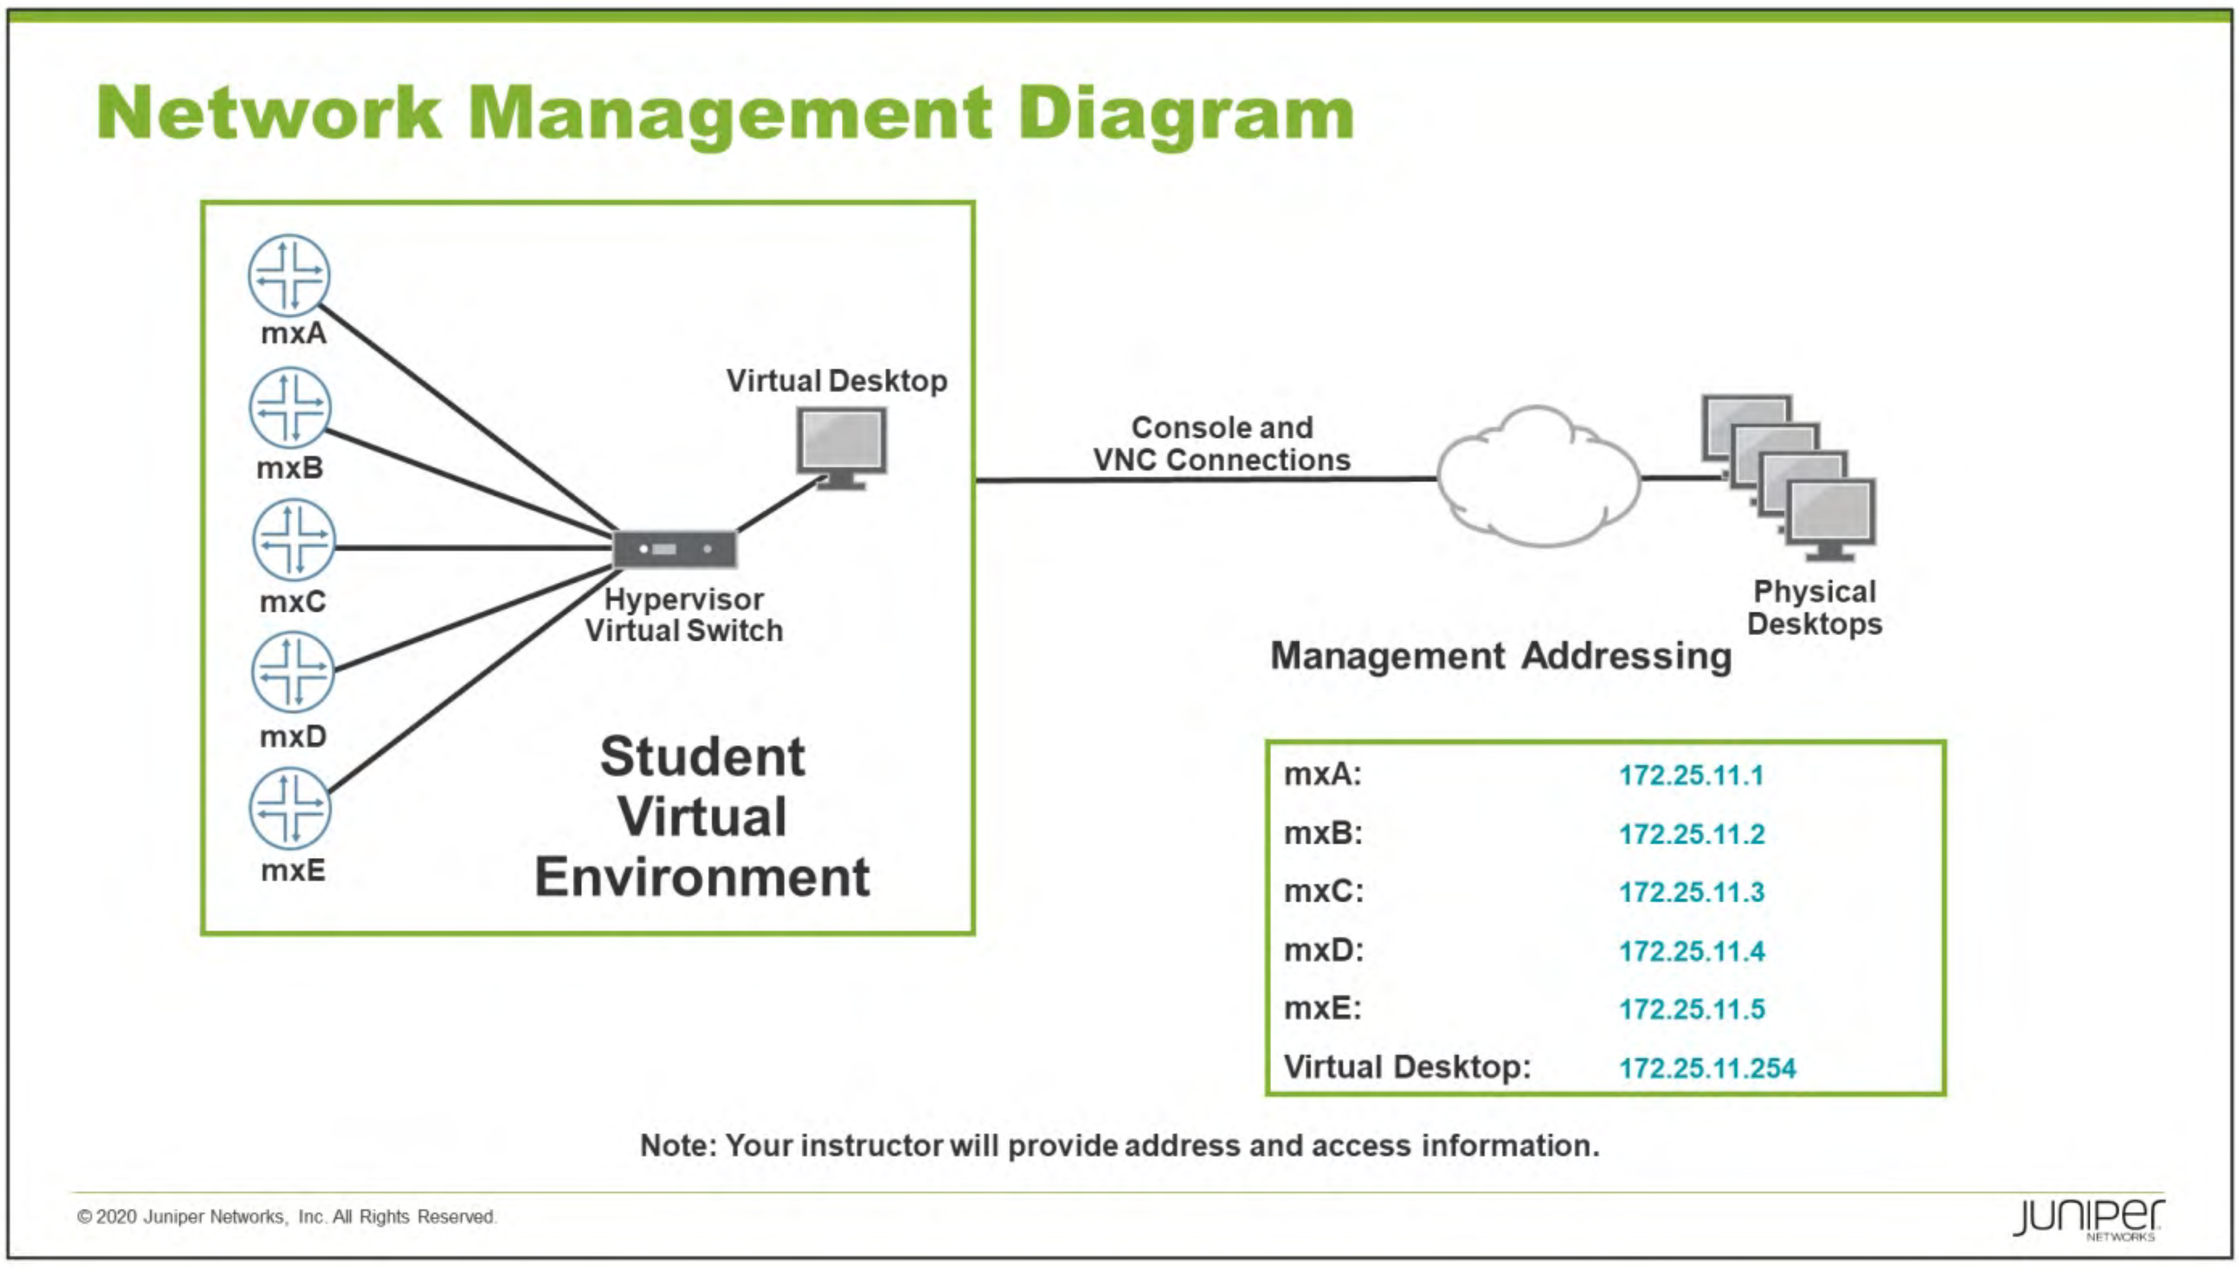
Network Management Diagram
mxA
mxB
mxC
mxD
mxE
Virtual Desktop
Hypervisor
Virtual Switch
Student
Virtual
Environment
Console and
VNC Connections
Physical
Desktops
Management Addressing
mxA: 172.25.11.1
mxB: 172.25.11.2
mxC: 172.25.11.3
mxD: 172.25.11.4
mxE: 172.25.11.5
Virtual Desktop: 172.25.11.254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

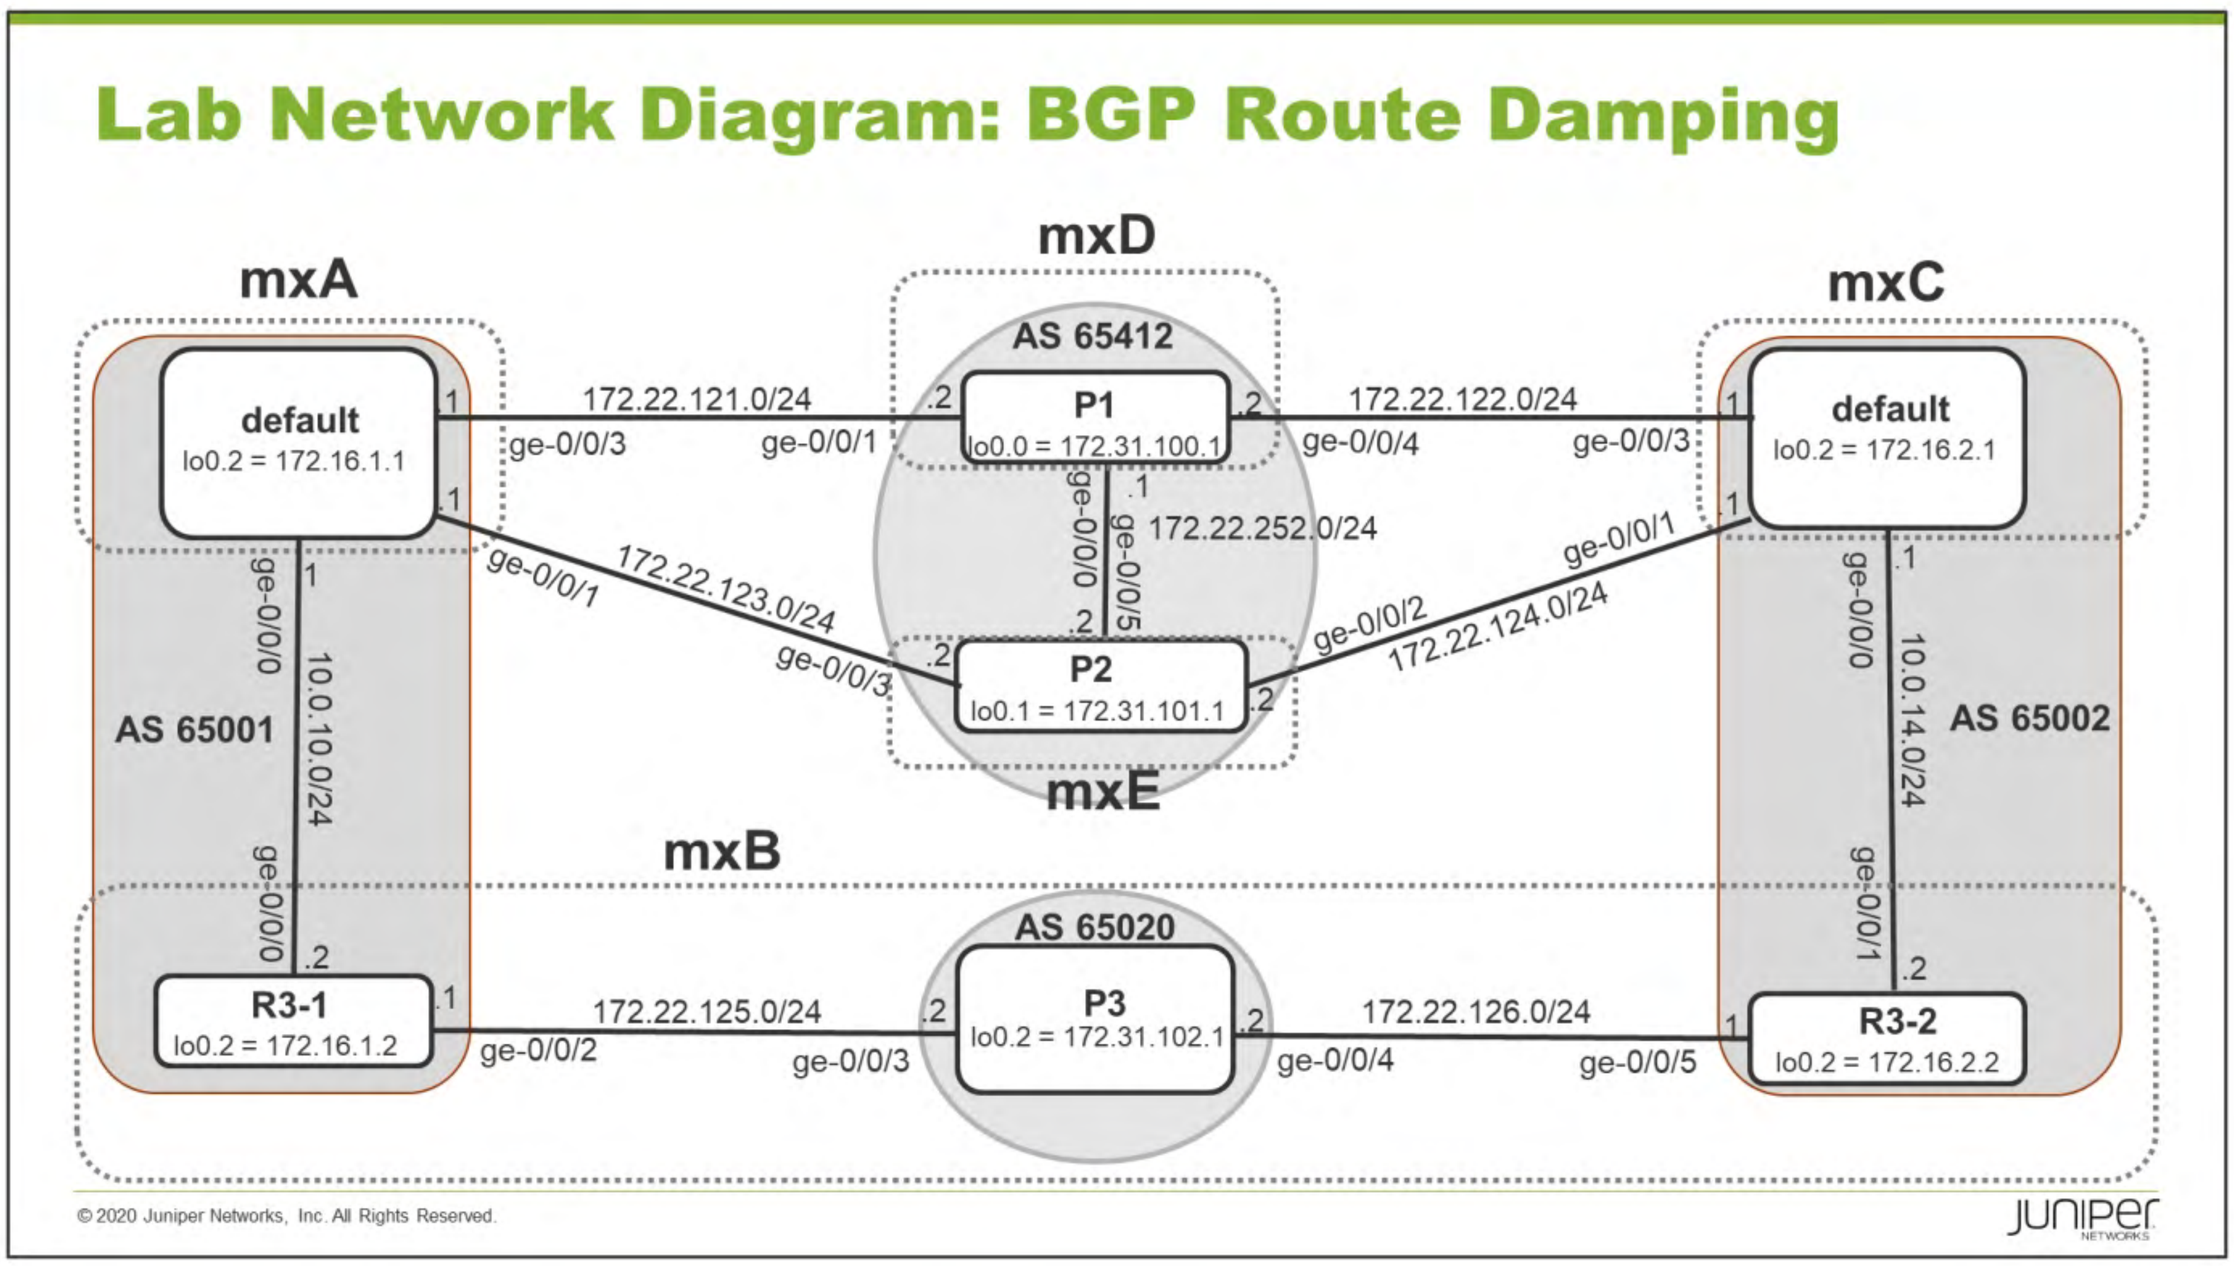
Lab Network Diagram: BGP Route Damping
mxA
mxD
mxC
mxB
mxE
AS 65412
AS 65001
AS 65002
AS 65020
default
lo0.2 = 172.16.1.1
P1
lo0.0 = 172.31.100.1
default
lo0.2 = 172.16.2.1
P2
lo0.1 = 172.31.101.1
R3-1
lo0.2 = 172.16.1.2
P3
lo0.2 = 172.31.102.1
R3-2
lo0.2 = 172.16.2.2
172.22.121.0/24
172.22.122.0/24
172.22.123.0/24
172.22.124.0/24
172.22.125.0/24
172.22.126.0/24
172.22.252.0/24
10.0.10.0/24
10.0.14.0/24
ge-0/0/3
ge-0/0/1
ge-0/0/4
ge-0/0/3
ge-0/0/0
ge-0/0/5
ge-0/0/1
ge-0/0/1
ge-0/0/3
ge-0/0/2
ge-0/0/0
ge-0/0/0
ge-0/0/1
ge-0/0/2
ge-0/0/3
ge-0/0/4
ge-0/0/5
© 2020 Juniper Networks, Inc. All Rights Reserved.
juniper NETWORKS

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